

CDP Private Cloud Data Services 1.5.0

# CDP Private Cloud Data Services Management Console Release Notes

Date published: 2020-12-16

Date modified: 2023-01-25

# CLUSTERA

<https://docs.cloudera.com/>

# Legal Notice

© Cloudera Inc. 2024. All rights reserved.

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

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

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

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

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

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

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

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

# Contents

<b>What's new.....</b>	<b>4</b>
What's new June 14, 2023 (CDP Private Cloud Data Services 1.5.1).....	4
January 25, 2023.....	4
November 18, 2022.....	5
June 22, 2022.....	5
March 25, 2022.....	5
December 13, 2021.....	6
November 8, 2021.....	6
October 4, 2021.....	6
April 27, 2021.....	7
December 16, 2020.....	7
August 17, 2020.....	8
<b>Known issues.....</b>	<b>8</b>
<b>Fixed issues.....</b>	<b>21</b>

## What's new

This section lists major features and updates for the CDP Private Cloud Management Console service.

### What's new June 14, 2023 (CDP Private Cloud Data Services 1.5.1)

New features in the 1.5.1 release of the CDP Private Cloud Management Console service.

#### Security

- CVEs

- DirName Kerberos support for Data Services (CDW and CML)

#### Data Services Enablements

- Cadence support to enable CML DRS capability

- Client configuration for enabling CDE with Kafka

- ServiceDiscovery API to enable Data Services access to Base Storage services

- Provide Spark3 client configurations for CML

- Elastic Quota Management capabilities for CDE

#### Certifications

- FreeIPA

- OCP 4.11

- RKE2 (v1.24) and Longhorn (1.4.2) version upgrades

#### Backup and Restore Manager

You can use the Backup and Restore Manager in the CDP Private Cloud Data Services Management Console to back up and restore Kubernetes namespaces and resources for Control Plane on Embedded Container Service (ECS) and OpenShift Container Platform (OCP). Backup and Restore Manager leverages the Data Recovery Service (DRS) capabilities to backup and restore the namespaces and resources.

#### External metadata databases are no longer supported on OCP

As of CDP Private Cloud Data Services 1.5.1, external Control Plane metadata databases are no longer supported. New installs require the use of an embedded Control Plane database. Upgrades from CDP Private Cloud Data Services 1.4.1 or 1.5.0 to 1.5.1 are supported, but there is currently no migration path from a previous external Control Plane database to the embedded database. Upgrades from 1.4.0 or 1.5.0 with external Control Plane metadata databases also require additional steps, which are described in the CDP Private Cloud Data Services 1.5.1 upgrade topics.

### January 25, 2023

New features in the 1.5.0 release of the CDP Private Cloud Management Console service.

#### Certification

- Platform certification of OCP 4.10, CentOS 8.6, RHEL 8.6

- K8s 1.23 platform-wide support (OCP 4.10 & ECS with K8s 1.23)

#### Enterprise Readiness

- Data Recovery Service for Platform and CDW

- Longhorn Logs are now included as part of the Diagnostics bundle

- Longhorn stability improvements

Ability to silence alert notifications via the CDP Management Console dashboard  
Metric replication with Prometheus Remote Write Protocol (Preview)

## November 18, 2022

New features in the 1.4.1 release of the CDP Private Cloud Management Console service.

### **New screen for editing external database configurations for Embedded Container Service (ECS) clusters.**

You can now edit the following database configurations for external databases in ECS clusters. (If you have deployed using the OpenShift Container Platform (OCP), you can view these configurations, but not change them:

- Username
- Password
- Database Host
- Database Port
- Database Name

An option to test the connection is also available.

[See Modifying database properties.](#)

## June 22, 2022

New features in the 1.4.0 release of the CDP Private Cloud Management Console service.

There are no new features in the Management Console.

## March 25, 2022

New features in the 1.3.4 release of the CDP Private Cloud Management Console service.

### **Terminology Changes**

The following terminology has been changed:

- "Private Cloud Experiences" is now "Private Cloud Data Services"
- "Experiences Cluster" is now "Containerized Cluster"
- "Experiences Compute Service" is now "Embedded Container Service"

### **External database names, hosts, ports, usernames, and passwords are now exposed through an Environment Service configuration**

External database names, hosts, ports, usernames and passwords are now exposed through an Environment Service configuration called `\unifiedDbDetails`. It can be fetched or updated using the `get-environment-setting` or `set-environment-setting` command of the CDP Environments CLI, respectively. When these configurations are set, any new values provided will be merged with existing values, with new values taking precedence.

### **When installing CDP Private Cloud Data Services, you can now select which Docker images to download**

Prior to this change, all images were downloaded, regardless of which Data Services are deployed. With this change, you can select which images to download during the installation process.

### **Rancher Kubernetes Engine (RKE) has been updated to version 1.21.8**

### **Longhorn has been upgraded to version 1.2.2**

### **Cloudera Manager now prevents ECS Server hosts from running workloads.**

ECS Server nodes will now automatically be configured by Cloudera Manager to prevent workloads from running on them.

**ECS hosts can now be configured to reserve then for workloads that require GPU drivers**

You can configure this in the following ways:

- During ECS installation. After adding the GPU host(s) to Cloudera Manager but prior to creation of the ECS cluster, visit the Host Configuration page, select the "Dedicated GPU Node for Data Services" checkbox and Save the configuration. Repeat for all hosts on which the taint is desired. Then, proceed with installation via the Add Cluster wizard.
- During ECS upgrade. After upgrading Cloudera Manager (if applicable), set the host configuration as described above on one or more hosts in the ECS cluster. Then, proceed with upgrade via the Upgrade Cluster wizard.
- Independently of ECS install or upgrade. Set the host configuration as described above on one or more hosts in the ECS cluster. Redeploy the client configuration on the ECS cluster. Finally, run the "Reapply All Settings to Cluster" command on the ECS service, which can be found in the Service Actions menu.

**SELinux is now supported for ECS clusters.**

See [Configuring a Containerized cluster with SELinux](#) for the steps to configure SELinux.

**FreeIPA is now supported for Kerberos configurations**

See [Configuring LDAP authentication for CDP Private Cloud](#)

## December 13, 2021

There are no new features in the 1.3.3 release of the CDP Private Cloud Management Console service.

## November 8, 2021

There are no new features in the 1.3.2 release of the CDP Private Cloud Management Console service.

## October 4, 2021

The 1.3.1 release of the CDP Private Cloud Management Console service provides the following new features:

**Support for new resource roles**

CDP Private Cloud Management Console introduces two new resource roles for managing the Cloudera Data Engineering (CDE) services: DEAdmin and DEUser.

- The DEAdmin role grants a CDP user/group the permission to create, delete and administer Cloudera Data Engineering services for a given CDP environment.
- The DEUser role grants a CDP user/group the permission to list and use Cloudera Data Engineering services for a given CDP environment.

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

**Support for CDP CLI**

CDP Private Cloud Data Services enables you to configure the CDP client that gives you access to the CDP CLI tool. The CDP CLI allows you to perform the same actions as can be performed from the Management Console.

For more information, see [CDP Private Cloud CLI](#).

**Configuring alert rules**

CDP Private Cloud Management Console enables you to define alert rules based on [PromQL](#) expressions. The alerts are automatically triggered when specific events occur in your CDP Private Cloud Data Services deployment.

For more information, see [Configuring alert rules](#).

## April 27, 2021

The 1.2 release of the CDP Private Cloud Management Console service provides the following new features:

### **Managing user groups**

CDP Private Cloud Management Console allows you to manage user groups. As a CDP administrator, you can create a group and manage the group membership. You can also manage the roles and resources assigned to the group.

For more information, see [Understanding CDP groups](#).

### **Uploading multiple types of TLS certificates to the CDP trust store**

CDP Private Cloud Management Console enables you to update the TLS certificates that CDP uses to make secure connections with different types of services and workloads such as external databases, external vaults, Docker registries, services used during CDP Private Cloud installation and runtime, and so on.

For more information, see [Update TLS certificates](#).

### **Configuring alert receivers**

CDP Private Cloud Management Console enables you to configure alert receivers to trigger automated system-specific event notifications through external services such as emails, Slack channel messages, webhook notifications, or PagerDuty messages.

For more information, see [Configuring alert receivers](#).

### **Updated options for collecting diagnostic data**

You can collect and download CDP Private Cloud diagnostic data for different components and services by specifying various criteria in the Collect and Send Diagnostic Data pop-up window.

For more information, see [Private Cloud Monitoring and Alerts](#).

## December 16, 2020

The 1.1 release of the CDP Private Cloud Management Console service provides the following new features:

### **Support for Red Hat OpenShift version 4.5**

This release of CDP Private Cloud now supports Red Hat OpenShift Container Platform version 4.5.x or later.

For more information, see [OpenShift Container Platform requirements](#).

### **Viewing the Platform Management Dashboard**

You can get insights into the resource utilization and health of the CDP Private Cloud Management Console components and the active environments through the new Dashboard page.

For more information, see [Management Console Dashboard](#).

### **Updating TLS certificates**

You can now update TLS certificates that the Management Console uses for secure connections with an external database, an external vault, and the Cloudera Manager associated with the CDP Private Cloud base cluster.

For more information, see the following:

- [Update a TLS certificate for a secure database connection](#)
- [Update a TLS certificate for a secure vault connection](#)
- [Update a TLS certificate for a secure Cloudera Manager connection](#)

### Support for OpenLDAP

In addition to authenticating users through Microsoft Active Directory LDAP, you can now use OpenLDAP for authenticating users.

For more information, see [User Management](#).

### Importing users in bulk

You can now perform a bulk import of users to CDP Private Cloud and assign them rights and roles. This improves the experience from the previous version where each user was required to log in at least once before access rights could be configured.

For more information, see [Importing or uploading users](#).

## August 17, 2020

This is the first release of the CDP Private Cloud Management Console service.

The Management Console service provides the following capabilities:

### Registering environments

In a CDP Private Cloud deployment, an environment represents that the association between a Data Lake and multiple compute resources using which you can provision and manage workloads for services such as Data Warehouse and Machine Learning. You can register as many environments as you require.

For more information, see [Private Cloud Environments](#).

### Managing users

The CDP Private Cloud Management Console service allows you to perform different type of user management tasks such as creating and onboarding different types of users, configuring identity providers, adding users, assigning roles to users, generating access keys, and removing roles assigned to users.

For more information, see [Private Cloud User Management](#).

### Accessing resource utilization and health monitoring dashboards

The CDP Private Cloud Management Console service contains dashboards that help you track the consumption of compute resources and monitor health information. The resource utilization dashboard provides an overview of the resources consumed by the CDP workloads while the monitoring dashboards provide health information for both the control plane and specific environments.

For more information, see [Private Cloud Resource Utilization](#) and [Private Cloud Monitoring and Alerts](#).

## Known issues for the CDP Private Cloud Data Services Management Console

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

### Known Issues in Management Console 1.5.1

#### OPSX-4560: ECS server restart failed in an air gap environment as it requires "yum install" from repo

ECS service restart may fail in an air gap environment with a yum download error.

Workaround:



Open the `/opt/cloudera/cm-agent/service/ecs/rke.sh` file with a text editor and remove the following line:

```
yum -y install iscsi-initiator-utils nfs-utils
```

### **OPSAPS-68558: [7.9.5->7.11.3.2] CM upgrade failed with BeanCreationException: Error creating bean with name 'com.cloudera.server.cmf.TrialState'**

After upgrading the Cloudera Manager package, the Cloudera Manager Server does not start. An error about "Active Commands" is shown in the Cloudera Manager Server log.

This may happen when the Private Cloud Data Services Control Plane is actively issuing requests to Cloudera Manager while an upgrade is being performed.

Workaround:

Before upgrading Cloudera Manager make sure there are no active commands. If there are any active commands, wait for them to complete before starting a Cloudera Manager upgrade.

If Cloudera Manager restart fails after upgrade due to an active `getClientConfig` command, check the Cloudera Manager server log for a "There are 1 active commands of type `GetClientConfigFiles`" error. This may block a Cloudera Manager restart after upgrade. Use the following steps to resolve this issue:

1. Login to Cloudera Manager database.
2. Search for any active `GetClientConfigFiles` command in the `COMMANDS` table.

```
UPDATE COMMANDS SET active=0,success=false,state='CANCELLED'
where command_id=<command_id>;
```

3. Delete these entries, including foreign key dependencies, in the following tables:

- `PROCESSES`
- `PROCESSES_DETAIL`
- `COMMANDS_DETAIL`

```
cm=> DELETE FROM COMMANDS where command_id=1546340765;
ERROR: update or delete on table "commands" violates foreign
key constraint "fk_process_command" on table "processes"
DETAIL: Key (command_id)=(1546340765) is still referenced fro
m table "processes".
cm=>
cm=> DELETE FROM processes where command_id=1546340765;
ERROR: update or delete on table "processes" violates foreign
key constraint "fk_processes_detail_process" on table "proc
esses_detail"
DETAIL: Key (process_id)=(1546340766) is still referenced fro
m table "processes_detail".
cm=>
cm=>
cm=> DELETE FROM processes_detail where process_id=1546340766;
DELETE 1
cm=> DELETE FROM processes where command_id=1546340765;
DELETE 1
cm=> DELETE FROM COMMANDS where command_id=1546340765;
ERROR: update or delete on table "commands" violates foreign
key constraint "fk_commands_detail_command" on table "comma
nds_detail"
DETAIL: Key (command_id)=(1546340765) is still referenced f
rom table "commands_detail".
cm=>
cm=> DELETE FROM commands_detail where command_id=1546340765;
DELETE 1
```

```
cm=> DELETE FROM COMMANDS where command_id=1546340765;  
DELETE 1
```

4. Restart the Cloudera Manager server.

### External metadata databases are no longer supported on OCP

As of CDP Private Cloud Data Services 1.5.1, external Control Plane metadata databases are no longer supported. New installs require the use of an embedded Control Plane database. Upgrades from CDP Private Cloud Data Services 1.4.1 or 1.5.0 to 1.5.1 are supported, but there is currently no migration path from a previous external Control Plane database to the embedded Control Plane database. Upgrades from 1.4.0 or 1.5.0 with external Control Plane metadata databases also require additional steps, which are described in the CDP Private Cloud Data Services 1.5.1 upgrade topics.

### DOCS-18031: Nodes are in "Not Ready" status during Rolling Restart of ECS

During a rolling restart of ECS, nodes are in a "Not Ready" state, and the dmesg command returns the following error message on the applicable nodes.

```
[Tue Aug 8 16:30:50 2023] nfs: server 10.46.157.145 not responding, timed out  
[Tue Aug 8 16:31:16 2023] nfs: server 10.46.157.145 not responding, timed out  
[Tue Aug 8 16:31:30 2023] nfs: server 10.46.157.145 not responding, timed out
```

Also, the df command may hang on these hosts.

Workaround:

Run the following commands on each unavailable node:

1. Find the NFS mounts:

```
mount | grep "nfs"
```

2. Force unmount:

```
umount -f <mount points found in Step 1 separated with a space>
```

### OPSAPS-67214: Single Node | Restart Stability | Rolling start is failing with "global timeout reached: 10m0s, error when evicting pods"

For the ECS service, rolling restart is not applicable to a single node cluster.

Workaround:

Instead of a rolling restart, you should stop and start the ECS Service.

### CDPVC-1098: How to refresh the YuniKorn configuration

Sometimes it is possible for the scheduler state to go out of sync from the cluster state. This may result in pods in Pending and ApplicationRejected states, with pod events showing Placement Rule related errors. To recover from this, you may need to refresh the YuniKorn configuration.

Workaround:

Follow the steps in [Refreshing the YuniKorn configuration](#).

### OPSAPS-67340: L1 runs failing as service monitor is in bad health state

Service Monitor (SMON) ends up in a bad health state after restarting the Cloudera Manager (CM) server, reporting problems with descriptor and metric schema age, when Kerberos and CM SPNEGO authentication are both enabled.

Workaround:

Use the following steps to restart SMON each time a CM server restart is required:

1. Stop SMON
2. Restart the CM server
3. Start SMON

If the health status is already bad, a simple restart of SMON is sufficient.

#### **DOCS-15855: Networking API is deprecated after upgrade to CDP Private Cloud Data Services 1.5.1 (K8s 1.24)**

When the control plane is upgraded from 1.4.1 to 1.5.1, the Kubernetes version changes to 1.24. The Livy pods running in existing Virtual Clusters (VCs) use a deprecated networking API for creating ingress for Spark driver pods. Because the old networking API is deprecated and does not exist in Kubernetes 1.24, any new job run will not work for the existing VCs.

#### **CDPQE-24295: Update docker client on docker.lab.eng.hortonworks machine**

When you attempt to execute the Docker command to fetch the Cloudera-provided images into your air-gapped environment, you may encounter an issue where Docker pulls an incorrect version of the HAProxy image, especially if you are using an outdated Docker client. This situation arises due to the Cloudera registry containing images with multiple platform versions. Unfortunately, older Docker clients may lack the capability to retrieve the appropriate architecture version, such as amd64.

Workaround:

Update the Docker client. It has been demonstrated that Docker 20.10.5 and later versions have been successful in resolving this problem.

#### **OPSX-4326: OCP upgrade from 1.5.0 to 1.5.1 – Restore is unsuccessful after upgrade**

After upgrading CDP Private Cloud Data Services on OCP from 1.5.0 to 1.5.1, Restore using a 1.5.0 backup could not be performed successfully.

Workaround:

Make a backup of the OpenShift routes before upgrading to 1.5.1. If you need to restore the control plane on a CDP Private Cloud Data Services 1.5.1 OpenShift cluster using a 1.5.0 backup, contact Cloudera Customer Support.

#### **OPSX-4266: ECS upgrade from 1.5.0 to 1.5.1 is failing in Cadence schema update job**

When upgrading from ECS 1.5.0 to 1.5.1, the CONTROL\_PLANE\_CANARY fails with the following error:

```
Firing alerts for Control Plane: Job did not complete in time, Job failed to complete.
```

And the cdp-release-cdp-cadence-schema-update job fails.

Workaround:

Use the following steps to manually execute the job:

1. Export the job manifest into a file:

```
kubectl get job cdp-release-cdp-cadence-schema-update -n <cdp> -o yaml > job.yaml
```

2. Delete the cdp-release-cdp-cadence-schema-update job:

```
kubectl delete job cdp-release-cdp-cadence-schema-update -n <cdp>
```

3. Remove runtime information from the manifest, such as:

```
resourceVersion
uid
selector
  matchLabels
    controller-uid
labels
  controller-uid
status section
```

4. Create the job:

```
kubectl apply -f job.yaml
```

If the job still fails, contact Cloudera Support.

**OPSX-4076:**

When you delete an environment after the backup event, the restore operation for the backup does not bring up the environment.

Create the environment manually.

**OPSX-4295:**

The logs for the backups created in CDP Private Cloud Data Services 1.5.0 version do not appear after you upgrade to version 1.5.1.

**OPSX-4024: CM truststore import into unified truststore should handle duplicate CommonNames**

If multiple CA certificates with the exact same value for the Common Name field are present in the Cloudera Manager truststore when a Private Cloud Data Services cluster is installed, only one of them may be imported into the Data Services truststore. This may cause certificate errors if an incorrect/old certificate is imported.

Workaround:

Remove old certificates from the Cloudera Manager truststore, and ensure certificates have unique Common Names.

**OPSX-2713: PVC ECS Installation: Failed to perform First Run of services**

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

Workaround:

Because 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 then retry from the beginning of the Add Cluster wizard. Any nodes that are reused must be cleaned before re-attempting installation.

**OPSX-3666: mlx\_crud\_app DB connection fails with error "unable to create connection: x509: certificate relies on legacy Common Name field, use SANs instead"**

After upgrade, the mlx-crud-app fails with the error "unable to create connection: x509: certificate relies on legacy Common Name field, use SANs instead"

Workaround:

If you are upgrading from CDP Private Cloud Data Services 1.4.1 or 1.5.0 to 1.5.1, and you were previously using an external database, you must regenerate the DB certificate with SAN before upgrading to CDP Private Cloud Data Services 1.5.1.

**OPSAPS-66166: FreeIPA cadminrole needs more privileges for PvC+ after upgrade**

After upgrade, the Cloudera Manager admin role may be missing the Host Administrators privilege in an upgraded cluster.

Workaround:

The cluster administrator should run the following command to manually add this privilege to the role.

```
ipa role-add-privilege <cmadminrole> --privileges="Host Administrators"
```

### COMOPS-2822: OCP error x509: certificate signed by unknown authority

The error x509: certificate signed by unknown authority usually means that the Docker daemon that is used by Kubernetes on the managed cluster does not trust the self-signed certificate.

Workaround:

Usually the fix is to copy the certificate to the path below on all of the worker nodes in the cluster:

```
/etc/docker/certs.d/<your_registry_host_name>:<your_registry_host_port>/ca.crt
```

### OPSX-4225: Upgrade failed as cadence pods are crashlooping post upgrade

When doing a fresh install of CDP Private Cloud Data Services 1.5.1, external metadata databases are no longer supported. Instead, the CDP Private Cloud Data Services installer will create an embedded database pod by default, which runs inside the Kubernetes cluster to host the databases required for installation.

If you are upgrading from CDP Private Cloud Data Services 1.4.1 or 1.5.0 to 1.5.1, and you were previously using an external database, you must run the following psql commands to create the required databases. You should also ensure that the two new databases are owned by the common database users known by the control plane.

```
CREATE DATABASE db-cadence;  
CREATE DATABASE db-cadence-visibility;
```

### OPSAPS-67046: Docker Server role fails to come up and returns a connection error during ECS upgrade

When upgrading from 1.4.1 to 1.5.1, a Docker server role can sometimes fail to come up and return the following error:

```
grpc: addrConn.createTransport failed to connect to {unix:///var/run/docker/containerd/containerd.sock <nil> 0 <nil>}.  
Err :connection error: desc = "transport: error while dialing: dial unix:///var/run/docker/containerd/containerd.sock: timeout".  
Reconnecting... module=grpc  
failed to start containerd: timeout waiting for containerd to start
```

This error appears in the stderr file of the command, and can be caused by a mismatch in the pid of containerd.

Workaround:

1. Ensure that the problematic Docker server role has been stopped.
2. Log in to the failing Docker server host.
3. Run the following commands:

```
cd /var/run/docker/containerd/  
rm containerd.pid
```

- Restart the Docker server role.

#### Longhorn-4212 Somehow the Rebuilding field inside volume.meta is set to true causing the volume to get stuck in attaching/detaching loop

This is a condition that can occur in ECS Longhorn storage.

Since the volume has only 1 replica in this case, we can:

- Scale down the workload. The Longhorn volume will be detached.
- Wait for the Longhorn volume to be detached.
- SSH into the node that has the replica.
- cd into the replica folder (for example, /longhorn/replicas/pvc-126d40e2-7bff-4679-a310-e444e84df267-1a5dc941).
- Change the "Rebuilding" field from true to false in the volume.meta file.
- Scale up the workload to attach the volume.

#### OPX-3073 [ECS] First run command failed at setup storage step with error "Timed out waiting for local path storage to come up"

Pod stuck in pending state on host for a long time. Error in Role log related to CNI plugin:

Events:

Type	Reason	Age	From
Warning	FailedCreatePodSandBox	3m5s (x269 over 61m)	kubelet
(combined from similar events): Failed to create pod sandbox: rpc error: code = Unknown desc = failed to setup network for sandbox "70427e9b26fb014750dfe4441fdfae96cb4d73e3256ff5673217602d503e806f": failed to find plugin "calico" in path [/opt/cni/bin]			

Delete the cni directory on the host failing to launch pods:

```
ssh root@ecs-hal-p-7.vpc.cloudera.com rm -rf /var/lib/cni
```

Restart the canal pod running on that host:

```
kubectl get pods -n kube-system -o wide | grep ecs-hal-p-7.vpc.cloudera.com
kube-proxy-ecs-hal-p-7.vpc.cloudera.com          1/1
Running      0          11h    10.65.52.51    ecs-hal-p-7.vpc.cloudera.com    <none>    <none>
rke2-canal-1lkc9                                2/2
Running      0          11h    10.65.52.51    ecs-hal-p-7.vpc.cloudera.com    <none>    <none>
rke2-ingress-nginx-controller-dqtz8             1/1
Running      0          11h    10.65.52.51    ecs-hal-p-7.vpc.cloudera.com    <none>    <none>
kubectl delete pod rke2-canal-1lkc9 -n kube-system
```

#### OPX-3528: [Pulse] Prometheus config reload fails if multiple remote storage configurations exist with the same name

It is possible to create multiple remote storage configurations with the same name. However, if such a situation occurs, the metrics will not flow to the remote storage as the config reload of the original prometheus will fail.

At any point in time, there should never be multiple remote storage configurations existing that have the same name.

**OPSX-1405: Able to create multiple CDP PVC Environments with the same name**

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

Delete the environment and try again with only one user trying to create the environment.

**OPSX-1412: Creating a new environment through the CDP CLI reports intermittently that "Environment name is not unique" even though it is unique**

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

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

**OPSX-3323: Custom Log Redaction | String is not getting redacted from all places in diagnostic bundle**

Custom redaction rule for URLs does not work.

**Cloudera Data Engineering service fails to start due to Ozone**

If the Ozone service is missing, misconfigured, or having other issues when an Environment is registered in the Management Console, CDE fails to start.

Workaround:

1. Correct the issues with the Ozone service.
2. Ensure that Ozone is running as expected.
3. Re-create the environment.
4. Create a new Cloudera Data Engineering service.

**Known Issues in Management Console 1.5.0****Longhorn-4212 Somehow the Rebuilding field inside volume.meta is set to true causing the volume to get stuck in attaching/detaching loop**

This is a condition that can occur in ECS Longhorn storage.

Since the volume has only 1 replica in this case, we can:

1. Scale down the workload. The Longhorn volume will be detached.
2. Wait for the Longhorn volume to be detached.
3. SSH into the node that has the replica.
4. cd into the replica folder (for example, /longhorn/replicas/pvc-126d40e2-7bff-4679-a310-e444e84df267-1a5dc941).
5. Change the "Rebuilding" field from true to false in the volume.meta file.
6. Scale up the workload to attach the volume.

**COMPX-13185 Upgrade from 1.4.1 to 1.5.0 failed - error: timed out waiting for the condition on jobs/helm-install-longhorn**

Before ECS upgrade, you must update a specific ECS server node toleration explicitly to ensure a cleaner upgrade process.

Delete the cni directory on the host failing to launch pods:

```
ssh root@ecs-ha1-p-7.vpc.cloudera.com rm -rf /var/lib/cni
```

Before ECS upgrade, run the following commands on the ECS Server hosts:

```
TOLERATION='{ "spec": { "template": { "spec": { "tolerations": [ {
"effect": "NoSchedule", "key": "node-role.kubernetes.io/control-p
lane", "operator": "Exists" } ] } } } }'

kubectl patch deployment/yunikorn-admission-controller -n yuniko
rn -p "$TOLERATION"
kubectl patch deployment/yunikorn-scheduler -n yunikorn -p "$TO
LERATION"
```

### OPSX-3073 [ECS] First run command failed at setup storage step with error "Timed out waiting for local path storage to come up"

Pod stuck in pending state on host for a long time. Error in Role log related to CNI plugin:

Events:

Type	Reason	Age	From
Warning	FailedCreatePodSandBox	3m5s (x269 over 61m)	kubelet
(combined from similar events):			
Failed to create pod sandbox: rpc error: code = Unknown desc = f			
ailed to setup network for sandbox			
"70427e9b26fb014750dfe4441fdfae96cb4d73e3256ff5673217602d503e806			
f":			
failed to find plugin "calico" in path [/opt/cni/bin]			

Delete the cni directory on the host failing to launch pods:

```
ssh root@ecs-hal-p-7.vpc.cloudera.com rm -rf /var/lib/cni
```

Restart the canal pod running on that host:

```
kubectl get pods -n kube-system -o wide | grep ecs-hal-p-7.vpc.
cloudera.com
kube-proxy-ecs-hal-p-7.vpc.cloudera.com 1/1
Running 0 11h 10.65.52.51 ecs-hal-p-7.vpc.clo
udera.com <none> <none>
rke2-canal-1lkc9 2/2
Running 0 11h 10.65.52.51 ecs-hal-p-7.vpc.clou
dera.com <none> <none>
rke2-ingress-nginx-controller-dqtz8 1/1 R
unning 0 11h 10.65.52.51 ecs-hal-p-7.vpc.cloud
era.com <none> <none>
kubectl delete pod rke2-canal-1lkc9 -n kube-system
```

### OPSX-3528: [Pulse] Prometheus config reload fails if multiple remote storage configurations exist with the same name

It is possible to create multiple remote storage configurations with the same name. However, if such a situation occurs, the metrics will not flow to the remote storage as the config reload of the original prometheus will fail.

At any point in time, there should never be multiple remote storage configurations existing that have the same name.

### OPSX-2062: Platform not shown on the Compute Cluster UI tab

On the CDP Private Cloud Management Console UI in ECS, when listing the compute clusters, the Platform field is empty (null) instead of displaying RKE as the Platform.

### OPSX-1405: Able to create multiple CDP PVC Environments with the same name



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

Delete the environment and try again with only one user trying to create the environment.

**OPsx-1412: Creating a new environment through the CDP CLI reports intermittently that "Environment name is not unique" even though it is unique**

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

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

**OPsx-2062: Platform not shown on the Compute Cluster UI tab**

On CDP Private Console UI in ECS, when listing the compute clusters, the Platform field is empty (null) instead of displaying RKE as the Platform.

**OPsx-3323: Custom Log Redaction | String is not getting redacted from all places in diagnostic bundle**

Custom redaction rule for URLs does not work.

**Cloudera Data Engineering service fails to start due to Ozone**

If the Ozone service is missing, misconfigured, or having other issues when an Environment is registered in the Management Console, CDE fails to start.

Workaround:

1. Correct the issues with the Ozone service.
2. Ensure that Ozone is running as expected.
3. Re-create the environment.
4. Create a new Cloudera Data Engineering service.

**Known Issues in Management Console 1.4.1**

**INSIGHT-2469: COE Insight from case 922848: Not able to connect to bit bucket**

After installing CML on an ECS cluster, users were not able to connect the internal bitbucket repo.

Workaround:

In this case the MTU of the ECS virtual network interfaces were larger than that of host external interface, which may cause the network requests from ECS containers to get truncated.

The Container Network Interface (CNI) is a framework for dynamically configuring networking resources. CNI integrates smoothly with Kubernetes to enable the use of an overlay or underlay network to automatically configure the network between pods. Cloudera ECS uses Calico as the CNI network provider.

The MTU of the pods' virtual network interface can be seen by running the `ifconfig` command.

The default MTU of the virtual network interfaces is 1450.

The MTU setting of the virtual interfaces is stored as a configmap in the `kube-system` namespace. To modify the MTU, edit the `rke2-canal-config` configmap.

```
$ /var/lib/rancher/rke2/bin/kubectl --kubeconfig  
/etc/rancher/rke2/rke2.yaml --namespace kube-system  
edit cm rke2-canal-config
```

Find the `veth_mtu` parameter in the YAML content. Modify the default value of 1450 to the required MTU size.

Next, restart the `rke2-canal` pods from the `kube-system` namespace. There will be `rke2-canal` pods for each ECS node.

After the pods are restarted, all subsequent new pods will use the new MTU setting. However, existing pods that are already running will remain on the old MTU setting. Restart all of the pods to apply the new MTU setting.

#### **OPSAPS-67046: Docker Server role fails to come up and returns a connection error**

A Docker server role can sometimes fail to come up and return the following error:

```
grpc: addrConn.createTransport failed to connect to {unix:///var/run/docker/containerd/containerd.sock <nil> 0 <nil>}.  
Err :connection error: desc = "transport: error while dialing: dial unix:///var/run/docker/containerd/containerd.sock: timeout".  
Reconnecting... module=grpc  
failed to start containerd: timeout waiting for containerd to start
```

This error appears in the stderr file of the command, and can be caused by a mismatch in the pid of containerd.

Workaround:

1. Ensure that the problematic Docker server role has been stopped.
2. Log in to the failing Docker server host.
3. Run the following commands:

```
cd /var/run/docker/containerd/  
rm containerd.pid
```

4. Restart the Docker server role.

#### **OPSX-1405: Able to create multiple CDP PVC Environments with the same name**

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

Delete the environment and try again with only one user trying to create the environment.

#### **OPSX-1412: Creating a new environment through the CDP CLI reports intermittently that "Environment name is not unique" even though it is unique**

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

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

#### **OPSX-3323: Custom Log Redaction | String is not getting redacted from all places in diagnostic bundle**

Custom redaction rule for URLs does not work.

#### **Cloudera Data Engineering service fails to start due to Ozone**

If the Ozone service is missing, misconfigured, or having other issues when an Environment is registered in the Management Console, CDE fails to start.

Workaround:

1. Correct the issues with the Ozone service.
2. Ensure that Ozone is running as expected.
3. Re-create the environment.
4. Create a new Cloudera Data Engineering service.

### **Known Issues in Management Console 1.4.0**

#### **Cloudera Data Engineering service fails to start due to Ozone**

If the Ozone service is missing, misconfigured, or having other issues when an Environment is registered in the Management Console, CDE fails to start.

Workaround:

1. Correct the issues with the Ozone service.
2. Ensure that Ozone is running as expected.
3. Re-create the environment.
4. Create a new Cloudera Data Engineering service.

**OPSX-2062: Platform not shown on the Compute Cluster UI tab**

On CDP Private Console UI in ECS, when listing the compute clusters, the Platform field is empty (null) instead of displaying RKE as the Platform.

None.

**OPSX-2713: ECS Installation: Failed to perform First Run of services.**

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

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-735: Kerberos service should handle CM downtime**

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

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: Able to create multiple CDP PVC Environments with the same name**

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

Delete the environment and try again with only one user trying to create the environment.

**OPSX-1412: Creating a new environment through the CDP CLI intermittently reports that, "Environment name is not unique" even though it is unique**

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

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

**OPSX-2484: FileAlreadyExistsException during timestamp filtering**

The timestamp filtering may result in FileAlreadyExistsException when there is a file with same name already existing in the tmp directory.

**OPSX-2772: For Account Administrator user, update roles functionality should be disabled**

An Account Administrator user holds the biggest set of privileges, and is not allowed to modify via current UI, even user try to modify permissions system doesn't support changing for account administrator.

**Known Issues for Management Console 1.3.x and lower****Recover fast in case of a Node failures with ECS HA**

When a node is deleted from cloud or made unavailable, it is observed that the it takes more than two minutes until the pods were rescheduled on another node.

It takes some time for the nodes to recover. Failure detection and pod-transitioning are not instantaneous.

**Cloudera Manager 7.6.1 is not compatible with CDP Private Cloud Data Servicesversion 1.3.4.**

You must use Cloudera Manager version 7.5.5 with this release.

**CDP Private Cloud Data Services ECS Installation: Failed to perform First Run of services.**

If an issue is encountered during the Install 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.

**Environment creation through the CDP CLI fails when the base cluster includes Ozone**

Problem: Attempt to create an environment using the CDP command-line interface fails in a CDP Private Cloud Data Services deployment when the Private Cloud Base cluster is in a degraded state and includes Ozone service.

Workaround: Stopping the Ozone service temporarily in the Private Cloud Base cluster during environment creation prevents the control plane from using Ozone as a logging destination, and avoids this issue.

**Filtering the diagnostic data by time range might result in a FileAlreadyExistsException**

Problem: Filtering the collected diagnostic data might result in a FileAlreadyExistsException if the /tmp directory already contains a file by that name.

There is currently no workaround for this issue.

**Full cluster name does not display in the Register Environment Wizard**

None

**Kerberos service does not always handle Cloudera Manager downtime**

Problem: The Cloudera Manager Server in the base cluster must be running to generate Kerberos principals for CDP Private Cloud. If there is downtime, you might observe Kerberos-related errors.

Resolve downtime issues on Cloudera Manager. If you encounter Kerberos errors, you can retry the concerned operation such as creating Virtual Warehouses.

**Management Console allows registration of two environments of the same name**

Problem: If two users attempt to register environments of the same name at the same time, this might result in an unusable environment.

Delete the environment and ensure that only one user attempts to register a new environment.

**Not all images are pushed during upgrade**

A retry of a failed upgrade intermittently fails at the Copy Images to Docker Registry step due to images not being found locally.

The failed images can be loaded manually (with a docker load), and the upgrade resumed. To identify which images need to be loaded take a look at the stderr file. The downloaded images are present in the Docker Data Directory.

**The Environments page on the Management Console UI for an environment in a deployment using ECS does not display the platform name**

Problem: When you view the details of an environment using the Management Console UI in a CDP Private Cloud Data Services deployment using ECS, the Platform field appears blank.

Use the relevant CDP CLI command from the environments module to view the required details.

**Updating user roles for the admin user does not update privileges**

In the Management Console, changing roles on the User Management page does not change privileges of the admin user.

None

**Upgrade applies values that cannot be patched**

If the size of a persistent volume claim in a Containerized Cluster is manually modified, subsequent upgrades of the cluster will fail.

**Incorrect warning about stale Kerberos client configurations**

If Cloudera Manager is configured to manage `krb5.conf`, ECS clusters may display a warning that they have stale Kerberos client configurations. Clicking on the warning may show an "Access denied" error.

No action is needed. ECS clusters do not require Kerberos client configurations to be deployed on those hosts.

**Vault becomes sealed**

If a host in an ECS cluster fails or restarts, the Vault may have become sealed. (You may see a Health Test alert in Cloudera Manager for the ECS service stating Vault instance is sealed.)

Unseal the Vault. In the Cloudera Manager Admin Console, go to the ECS service and click ActionsUnseal .

## Fixed Issues for the CDP Private Cloud Data Services Management Console

This section lists the issues that have been fixed since the last release of the CDP Private Cloud Management Console service.

**Fixed Issues in Management Console 1.5.1****OPSAPS-65551: Increase default fd limit for ECS**

The default maximum process FD limit for the ECS agent and server roles has been set to 1048576 to avoid a "too many open files" error. Changes have been made to `EcsParams`, `EcsAgentRoleHandler` and `EcsServerRoleHandler`.

**OPSAPS-65852: Any stop of an ECS role should include a drain**

Previously, stopping and starting an ECS Role only stopped and started the role respectively, which caused issues in Kubernetes and Longhorn volume health to turn bad. Now, when a user stops an ECS Role (Server or Agent), we perform a "cordon" followed by a "drain" on the node and then stop the ECS Role on the node. When starting an ECS Role, we first start the ECS Role, then we do an "uncordon" on the node to allow Kubernetes to reuse the node for its workload. A restart operation on ECS Service will perform a Rolling Restart, which does the same steps involved in stopping and starting roles, but one node at a time.

**OPSX-3716: Certificates updated against key "undefined" from control plane UI**

Previously users were able to upload certificates without choosing a certificate type. This caused certificates to be saved as undefined. This fix now enforces users to choose Certificate Type before they can save the certificate.

**OPSAPS-58019: krb5.conf had includedir DIRNAME that caused krb5.conf to not get copied into CML and CDW**

Fixed the issue where if the `/etc/krb5.conf` file on the Cloudera Manager host contained `include` or `includedir` directives, Kerberos-related failures sometimes occurred. Expanded the `include` and `includedir` contents as part of the `krb5.conf` content before return to the user so that the files referred by these two directives do not need to be de-referenced by the user.

**OPSX-3942, ENGESC-19665: CP logs occupies large amount of disk space**

Fixed the issue where control plane logs were taking up a large amount of disk space:

1. Clean up the files created under the /tmp directory after the bundle collection.
2. Include control plane logs while purging. CP logs will be present in the /data/cp directory.

**OPSX-3619: Installer exits even with pending pods in single node installation**

Fixed the issue with a single node ECS deployment where the installer exited prematurely while pods were still in a pending state.

**OPSX-4010: [UI Issue] Deletion Response sent immediately but deletion happens in 1 Min**

Fixed the issue where after a backup deletion request from the UI, and a confirmation pop-up, the actual deletion did not occur until approximately one minute later.

**ENGESC-20112: Unable to progress ECS Upgrade**

Fixed the issue where the Starting ECS Agent command failed during upgrade, but did not support retry, so the upgrade could not be resumed.

**OPSX-2062: Platform not shown on the Compute Cluster UI tab**

Fixed the issue on the CDP Private Console UI in ECS, where when listing the compute clusters, the Platform field is empty (null) instead of displaying RKE as the Platform.

**OPSAPS-66433: Support rolling upgrade for Docker service**

Added support for rolling restart of the embedded Docker service. Support rolling upgrade of the Docker service while upgrading Private Cloud.

**OPSAPS-66559: Create command to clear pending pods in the cluster**

The Refresh ECS command will restart the pods that are in pending state for 10 minutes. This value can be configured using the WAIT\_TIME\_FOR\_POD\_READINESS parameter.

**OPSAPS-65753: Upgrade CP before upgrading K8S**

When upgrading an ECS cluster, the control plane will be upgraded before Kubernetes is upgraded.

**PULSE-498: Alerts for Ozone health tests are not reported on the Control Plane dashboard**

The Cloudera Monitoring Grafana Control Plane dashboard now displays alerts for the Ozone service.

**PULSE-77: schemaVersion should be updated to 3**

The topology schema version has been upgraded to version 3, which has stopped the invalid schema version:2 error message appearing in the log files.

**PULSE-53: nil pointer reference on calling createSilence API via CDP CLI**

You can now silence your alerts from the CDP CLI. Alert silencing avoids repeated alert pings when troubleshooting issues.

**Fixed Issues in Management Console 1.5.0****COMPX-13184 YuniKorn-admission-controller not getting scheduled after restart due to lack of tolerations**

Fixed the issue where ecs-webhooks from the ECS platform failed to update the YuniKorn namespace. Due to this lack of toleration update from ECS, YuniKorn will insert this toleration from Liftie deployment as an interim update.

**Fixed Issues in Management Console 1.4.0****OPSX-2697 Not all images are pushed in upgrade**

Fixed the issue of a retry of an upgrade failing at the Copy Images to Docker Registry step due to images not being found locally.

**Fixed Issues in Management Console 1.3.x****CVE-2021-44228 (Apache Log4j 2 vulnerability) has been addressed in CDW on CDP Private Cloud Management Console version 1.4.0**

Log4j 2 has been upgraded to version 2.17.

**Fix copy-docker-template**

Fixed the issue of a retry of only the Push Images to Docker Registry failing due to the image not being available locally.

**NFS provisioner fails on cluster with more than ~10 nodes**

Fixed longhorn nfs\_provisioner failing to start on clusters with more than 10 nodes.

**Longhorn for Kubernetes is upgraded to version 1.2.x**

Longhorn has been upgraded from version 1.1.2 to 1.2.x

**ECS High Availability fails during installation**

Fixed an issue where selecting multiple ECS Server hosts during install would randomly result in a installation failure.