

Cloudera Data Services on premises Release Notes

Date published: 2023-12-16

Date modified: 2025-11-08

CLOUDERA

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 in Cloudera Data Services on premises 1.5.5.....4**
- Fixed Issues for the Cloudera Data Services on premises 1.5.5..... 4**
- Known issues for the Cloudera Data Services on premises 1.5.5.....5**
- Cumulative hotfixes..... 12**
 - Cloudera Data Services on premises 1.5.5-CHF1..... 12
- Service packs..... 13**
 - Cloudera Data Services on premises 1.5.5 SP1.....13
 - What's new in Cloudera Data Services on premises 1.5.5 SP1..... 13
 - Fixed Issues for the Cloudera Data Services on premises 1.5.5 SP1..... 14
 - Known issues for the Cloudera Data Services on premises 1.5.5 SP1..... 16

What's new in Cloudera Data Services on premises 1.5.5

Understand the functionalities and improvements to features of Cloudera Control Plane install and upgrade components in Cloudera Data Services on premises 1.5.5.

Certificate Management

Cert-manager is an open-source tool for Kubernetes that automates the provisioning, management, and renewal of TLS certificates. Its documentation at <https://cert-manager.io/docs/> provides comprehensive guidance on installing, configuring, and using cert-manager to secure workloads with trusted X.509 certificates. Cloudera provides out-of-the-box support for Venafi Trust Protection Platform (TPP) as part of the Cloudera Embedded Container Service installation. By integrating cert-manager, the Data services achieve secure communication, reduced manual overhead, and compliance with security standards, leveraging its robust automation and flexibility. For more information on setting Cert-manager using Venafi TPP, see [Setting up Certification Manager using Venafi TPP](#).

New upgrade prechecks

New pre-upgrade checks have been added to the list. The additional checks verify if the control plane and the docker registry is ready for upgrade. For more information, see [Pre-upgrade checklist](#).

Quota Management for multiple base cluster support

Quota management enables you to control how resources are allocated within your Cloudera Data Services on premises clusters. In order to prevent a single workload from consuming all available cluster resources, you can limit the number of CPUs, GPUs, and memory allocated by application, user, business units, or Data Service by defining resource pools that define resource limits. Pools are organized in a hierarchical manner by defining nodes in the hierarchy with resource limits, which can then be subdivided as needed to allocate resources for an organization and to allocate resources to cluster or environment wide services such as the monitoring service. For information, see [Quota Management](#).

Creating multiple environments with different base or Data Lake clusters

To register an environment with a data lake cluster managed by a Cloudera Manager that is different from your existing Cloudera Manager, you need to add the certificates of the new Cloudera Manager to the Cloudera Management Console UI. If the existing Cloudera Manager and the new Cloudera Manager share the same root CA, and the root CA is already uploaded as the data lake certificate, then no additional certificate needs to be added. For more information, see [Creating multiple environments with different base or Data Lake clusters](#).

Fixed Issues for the Cloudera Data Services on premises 1.5.5

You can review the list of reported issues and their fixes in Cloudera Data Services on premises 1.5.5. Fixed issues represent selected issues that were previously logged through Cloudera Support, but are now addressed in the current Cloudera Data Services on premises release. These issues may have been reported in previous versions of Cloudera Data Services on premises as a known issue; meaning they were reported by customers or identified by Cloudera Quality Engineering teams.

OPSX-4308 - Display error in UI if listEnvironments failed

Error is now displayed on the Environments Page of the Cloudera Management Console UI, if an API failure is encountered.

OPSX-6048 - Clean up delete backup Custom Resource (CR) after the job is run

DeleteBackup now removes the backup deletion CR from resource `deletebackuprequests.drs.cdp.cloudera.com`

OPSX-5944 - Issues while uncordoning nodes during restart

The uncordon step was added into Cloudera Manager and is removed from the Cloudera Embedded Container Service parcel.

OPSX-5852 - Remove warn logs for "Unexpected partition in crn" from Cloudera Data Services on premises

"Unexpected partition in crn" log entries are now removed from the logs.

OPSX-5403 - Typecasting fails when truststore password is integer

When truststore password is set to all numbers (integer or float), control plane installation was failing in both Cloudera Embedded Container Service and OpenShift Container Platform. Safe datatype conversion is done to treat even numbers as string password. Even if numbers are used for truststore passwords, control plane installation will be successful.

OPSX-5903 - Upgrade failed with rke2-ingress-nginx-controller" exceeded its progress deadline

Automated the manual workaround of scaling down and scaling up the deployment when the earlier rollout or its status check fails.

OPSAPS-72270- ECS Restart|Start ECS| Start ECS command fails on uncordon nodes step

To resolve this issue:

1. Ensure the kube-apiserver is up and running for at least 60 seconds before proceeding with the uncordon step.
2. Use the correct target node name, not just the name of the node where the uncordon command is executed.

Known issues for the Cloudera Data Services on premises 1.5.5

You must be aware of the known issues and limitations, the areas of impact, and workaround in Cloudera Data Services on premises 1.5.5 release.

Known Issues in Cloudera Data Services on premises 1.5.5

OBS-8038: When using the Grafana Dashboard URL shortener, the shortened URL defaults to localhost:3000. This behaviour happens because the URL shortener uses the local server address instead of the actual domain name of the Cloudera Observability instance. As a result, users cannot access the shortened URL.

You must not use the shortened URL. To ensure users can access the URL, update it to use the correct Cloudera Observability instance domain name, such as `cp_domain/{shorten_url}{}`.



Note: `cp_domain` refers to the Cloudera Control Plane domain.

DWX-20809: Cloudera Data Services on premises installations on RHEL 8.9 or lower versions may encounter issues

You may notice issues when installing Cloudera Data Services on premises on Cloudera Embedded Container Service clusters running on RHEL 8.9 or lower versions. Pod crashloops are noticed with the following error:

```
Warning FailedCreatePodSandBox          1s (x2 over 4s)  kubel
et Failed to create pod
sandbox: rpc error: code = Unknown desc = failed to create containe
nerd task: failed to create
shim task: OCI runtime create
```

```
failed: runc create failed: unable to start container process: u
nable to init seccomp: error
loading seccomp filter into kernel: error loading seccomp filt
er: errno 524: unknown
```

The issue is due to a memory leak with 'seccomp' (Secure Computing Mode) in the Linux kernel. If your kernel version is not on 6.2 or higher versions or if it is not part of the list of versions mentioned [here](#), you may face issues during installation.

To avoid this issue, increase the value of `net.core.bpf_jit_limit` by running the following command on all ECS hosts:

```
[root@host ~]# sysctl net.core.bpf_jit_limit=528482304
```

However, Cloudera recommends upgrading the Linux kernel to an appropriate version that contains a patch for the memory leak issue. For a list of versions that contain this patch, see this [link](#).

COMPX-20705: [153CHF-155] Post ECS upgrade pods are stuck in ApplicationRejected State

After upgrading the CDP installation pods on Kubernetes could be left in a failure state showing "ApplicationRejected". This is caused by a delay in settings being applied to Kubernetes as part of the post upgrade steps.

To resolve this issue, restart the scheduler to pick up the latest settings for Kubernetes. Also, restart YuniKorn using the following commands:

```
kubectl scale deployment yunikorn-scheduler --replicas=0 -n yun
ikorn
kubectl scale deployment yunikorn-scheduler --replicas=1 -n yu
nikorn
```

OPSX-6303 - ECS server went down - 'etcdserver: mvcc: database space exceeded'

ECS server may fail with error message - "etcdserver: mvcc: database space exceeded" in large clusters.

1. Add to the safety valve for server group:

```
etcd-arg:
- "quota-backend-bytes=4294967296"
```

2. Restart stale services (Select the option `re-deploy client configs`).
3. The default value for `quota-backend-bytes` is 2 GB. It can be increased up to 8 GB.

OPSX-6295 - Control Plane upgrade failing with cadence-matching and cadence-history

In case of extra cadence-matching and cadence-history pod stuck in `Init:CreateContainerError` state, Cloudera Embedded Container Service Upgrade to 1.5.5 will be stuck in retry loop because of all pods running validation failure.

You need to manually apply the workaround to proceed further upgrade and get it done successfully. Hence, delete the stuck cadence pods.

OPSX-4391 - External docker cert not base64 encoded

When using Cloudera Data Services on premises on ECS, in some rare situations, the CA certificate for the Docker registry in the `cdp` namespace is incorrectly encoded, resulting in TLS errors when connecting to the Docker registry.

Compare and edit the contents of the "cdp-private-installer-docker-cert" secret in the cdp namespace so that it matches the contents of the "cdp-private-installer-docker-cert" secret in other namespaces. The secrets and their corresponding namespaces can be identified using the command:

```
kubect1 get secret -A | grep cdp-private-installer-docker-cert
```

Inspect each secret using the command:

```
kubect1 get secret -n cdp cdp-private-installer-docker-cert -o y  
aml
```

Replace "cdp" with the different namespace names. If necessary, modify the secret in the cdp namespace using the command:

```
kubect1 edit secret -n cdp cdp-private-installer-docker-cert
```

OPSX-6245 - Airgap | Multiple pods are in pending state on rolling restart

Performing back-to-back rolling restarts on ECS clusters can intermittently fail during the Vault unseal step. During rapid consecutive rolling restarts, the kube-controller-manager pod may not return to a ready state promptly. This can cause a cascading effect where other critical pods, including Vault, fails to initialize properly. As a result, the unseal Vault step fails.

As a workaround, perform the following steps:

1. Stop the ECS role that failed.
2. Start the ECS role again.
3. If required, perform the rolling restart again.

OPSX-4684 - Start ECS command shows green(finished) even though start docker server failed on one of the hosts

Docker service starts with one or more docker roles failed to start because the corresponding host is unhealthy.

Make sure the host is healthy. Start the the docker role in the host.

OPSX-5986 - ECS fresh install failing with helm-install-rke2-ingress-nginx pod failing to come into Completed state

ECS fresh install fails at the "Execute command Reapply All Settings to Cluster on service ECS" step due to a timeout waiting for helm-install.

To confirm the issue, run the following kubect1 command on the ECS server host to check if the pod is stuck in a running state:

```
kubect1 get pods -n kube-system | grep helm-install-rke2-ingress-  
nginx
```

To resolve the issue, manually delete the pod by running:

```
kubect1 delete pod <helm-install-rke2-ingress-nginx-pod-name> -n  
kube-system
```

Then, click Resume to proceed with the fresh install process on the Cloudera Manager UI.

OPSX-6298 - Issue on service namespace cleanup

There might be cases in which uninstalling services from the Cloudera Data Services on premises UI will fail due to various reasons.

In case uninstallation of a Service fails, trigger again the service uninstall process, and mark "Force Delete" to ensure that all metadata of the service will be removed from Cloudera side. Then,

move to the OpenShift UI, and there search for that service namespace / project. On that project/namespace select the Action button on the top right of the screen and choose to Delete the Project.

If you move back to the main Project screen you could see that the project is moving to status “Terminating” after which it will be removed from the OCP platform. That action will ensure that all the entities linked to that project/namespace will also be removed by OpenShift.

OPSX-6265 - Setting inotify max_user_instances config

We cannot recommend an exact value for inotify max_user_instances config. It depends on all workloads that are run in a given node.

With newly introduced features like istio, cert manager, in Cloudera Control Plane, there is a need to set inotify max_user_instancesconfig to 256 instead of 128 to resolve this issue.

COMPX-20362 - Use API to create a pool that has a subset of resource types

The Resource Management UI supports displaying only three resource types: CPU, memory and GPU. The Resource Management UI will always set all three resource types it knows about: CPU, Memory and GPU (K8s resource nvidia.com/gpu) when creating a quota. If no value is chosen for a resource type a value of 0 will be set, blocking the use of that resource type.

To create a pool that has a subset of resource types the REST API must be used as follows:

POST /api/v1/compute/createResourcePool

Payload:

```
{
  "pool": {
    "path": "root.environment.service.mypool",
    "policy": {
      "allocation": "INELASTIC"
    },
    "quota": {
      "cpu": "100 m",
      "memory": "10 GB"
    }
  }
}
```



Note: Payload needs to be confirmed and checked.

Known issues from previous releases carried in Cloudera Data Services on premises 1.5.5

Known Issues identified in 1.5.4

DOCS-21833: Orphaned replicas/pods are not getting auto cleaned up leading to volume fill-up issues

By default, Longhorn will not automatically delete the orphaned replica directory. One can enable the automatic deletion by setting orphan-auto-deletion to true.

No workaround available.

OPSX-5310: Longhorn engine images were not deployed on ECS server nodes

Longhorn engine images were not deployed on ECS server nodes due to missing tolerations for Cloudera Control Plane taints. This caused the engine DaemonSet to schedule only on ECS agent nodes, preventing deployment on Cloudera Control Plane nodes.

1. Check the Engine DaemonSet Status. Run the following command to check if the Longhorn engine DaemonSet is missing on certain nodes:

```
kubectl get ds -n longhorn-system | grep engine
```

2. Identify Taints on Affected Nodes. Run the following command to check for taints on affected nodes:

```
kubectl describe node <node-name> | grep Taints
```



Note: If you see, `node-role.kubernetes.io/control-plane=true:NoSchedule`, this confirms the issue.

3. Manually Edit the DaemonSet to Add a Toleration. Edit the Longhorn engine DaemonSet YAML:

```
kubectl edit ds -n longhorn-system engine-image-ei-<your-engine-id>
```

4. Add the following under tolerations:

```
tolerations:
- effect: NoSchedule
  key: node-role.kubernetes.io/control-plane
  operator: Equal
  value: "true"
```

5. Apply the changes and verify deployment. Save and exit the editor. Then, check if the DaemonSet is now running on all necessary nodes:

```
kubectl get pods -n longhorn-system -o wide | grep engine
```

Verify that the engine pods are successfully scheduled on the affected ECS server nodes.

OPSX-5155: OS Upgrade | Pods are not starting after the OS upgrade from RHEL 8.6 to 8.8

After an OS upgrade and start of the Cloudera Embedded Container Service service, pods fail to come up due to stale state.

Restart the Cloudera Embedded Container Service cluster.

OPSX-5055: Cloudera Embedded Container Service upgrade failed at Unseal Vault step

During an Cloudera Embedded Container Service upgrade from 1.5.2 to 1.5.4 release, the vault pod fails to start due to an error caused by the Longhorn volume unable to attach to the host. The error is as below:

```
Warning FailedAttachVolume 3m16s (x166 over 5h26m) attachdetach-controller
AttachVolume.Attach failed for volume "pvc-0ba86385-9064-4ef9-9019-71976b4902a5" :
rpc error: code = Internal desc = volume pvc-0ba86385-9064-4ef9-9019-71976b4902a5
failed to attach to node host-1.cloudera.com with attachmentID
csi-7659ab0e6655d308d2316536269de47b4e66062539f135bf6012bfc8b41fc345: the volume is
currently attached to different node host-2.cloudera.com
```

Follow below steps provided by SUSE to ensure the Longhorn volume is correctly attached to the node where the vault pod is running.

```
# Find out the volume name that is failing to attach to the vault
pod.
For e.g. pvc-bc73e7d3-c7e7-468a-b8e0-afdb8033e40b from the pod
logs.
```

```
kubectl edit volumeattachments.longhorn.io -n longhorn-system
pvc-bc73e7d3-c7e7-468a-b8e0-afdb8033e40b

# Update the "spec:" section of the volumeattachment and replace
attachmentTickets section with {} as shown below and save.
spec:
  attachmentTickets: {}
  volume: pvc-bc73e7d3-c7e7-468a-b8e0-afdb8033e40b

# scale down the vault statefulset to 0 and scale it back up.
kubectl scale sts vault --replicas=0 -n vault-system
kubectl scale sts vault --replicas=1 -n vault-system
```

OPSX-4684: Start Cloudera Embedded Container Service command shows green(finished) even though start docker server failed on one of the hosts

The Docker service starts, but one or more Docker roles fail to start because the corresponding host is unhealthy.

Ensure the host is healthy. Start the the Docker role on the host.

OPX-735: Kerberos service should handle Cloudera Manager downtime

The Cloudera Manager Server in the base cluster operates to generate Kerberos principals for Cloudera on premises. If there is downtime, you may observe Kerberos-related errors.

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

Known Issues identified in 1.5.2

OPsx-4594: [ECS Restart Stability] Post rolling restart few volumes are in detached state (vault being one of them)

After rolling restart there may be some volumes in detached state.

1. Open the Longhorn UI to view the detached volumes.
2. Perform the following operations for each volume in a detached state:
 - a. Identify the workload name and type from the volume details.
 - b. Identify the workload and number of replicas using `kubectl` or the Kubernetes UI.
 - c. Scale the workload down to 0.
 - d. Wait for the pods associated with the workload to fully terminate.
 - e. Scale up the workload up to the number of replicas it had originally.

To prevent this issue, use the Longhorn UI to set the number of replicas for the volume to at least 3.

OPX-4392: Getting the real client IP address in the application

CML has a feature for adding the audit event for each user action ([Monitoring User Events](#)). In Private Cloud, instead of the client IP, we are getting the internal IP, which is logged into the internal DB.

In ECS, add the `enable-real-ip` configuration as true for the nginx ingress controller:

```
apiVersion: v1
data:
  allow-snippet-annotations: "true"
  enable-real-ip: "true"
kind: ConfigMap
metadata:
  annotations:
    meta.helm.sh/release-name: rke2-ingress-nginx
    meta.helm.sh/release-namespace: kube-system
  creationTimestamp: "2023-05-09T04:54:53Z"
```

```
labels:
  app.kubernetes.io/component: controller
  app.kubernetes.io/instance: rke2-ingress-nginx
  app.kubernetes.io/managed-by: Helm
  app.kubernetes.io/name: rke2-ingress-nginx
  app.kubernetes.io/part-of: rke2-ingress-nginx
  app.kubernetes.io/version: 1.6.4
  helm.sh/chart: rke2-ingress-nginx-4.5.201
name: rke2-ingress-nginx-controller
namespace: kube-system
resourceVersion: "162559439"
uid: cca67b0c-bc05-4e1f-8439-7d44323f4624
```

In OpenShift Container Platform, you may be able configure this using [HAproxy with X-forward-for pass to OpenShift 4](#).

CDPVC-1137, CDPAM-4388, COMPX-15083, and COMPX-15418: OpenShift Container Platform version upgrade from 4.10 to 4.11 fails due to a Pod Disruption Budget (PDB) issue

PDB can prevent a node from draining which makes the nodes to report the “Ready,SchedulingDisabled” state. As a result, the node is not updated to correct the Kubernetes version when you upgrade OpenShift Container Platform from 4.10 to 4.11.

To resolve this issue, confirm that the upgrade has failed due to the PDB issue, and then manually delete the PDBs from the Cloudera on premises namespace.

1. Run the following command to check whether the nodes are stuck in the “Ready,SchedulingDisabled” state:

```
oc get nodes
```

2. Get the machine config daemon details of the particular pod as follows:

```
oc get po -n openshift-machine-config-operator -l 'k8s-app=machine-config-daemon' -o wide
```

3. Check the logs of the machine config operator of that particular node as follows:

```
oc logs -f -n openshift-machine-config-operator [***MACHINE-CONFIG-DAEMON-NAME***] -c machine-config-daemon
```

Replace [***MACHINE-CONFIG-DAEMON-NAME***] with the actual machine config daemon name.

You may see one of the following errors in the node logs:

- error when evicting pods/cdp-release-cpx-liftie-****" -n "[***PRIVATE-CLOUD-NAMESPACE***] Cannot evict pod as it would violate the pod's disruption budget
- error when evicting pods/"cdp-release-cluster-proxy-[*****]" -n "[***PRIVATE-CLOUD-NAMESPACE***] Cannot evict pod as it would violate the pod's disruption budget

Delete the PDB from the Cloudera on premises namespace as follows:

- a. Obtain the PDB for the cdp-release-cluster-proxy namespace:

```
oc get pdb -n [***PRIVATE-CLOUD-NAMESPACE***] | grep cdp-release-cluster-proxy
```

- b. Back up the PDB:

```
oc get pdb [***PDB-NAME-OF-CLUSTER-PROXY***] -n [***PRIVATE-CLOUD-NAMESPACE***] -o yaml >> [***BACKUP-FILE-NAME***].yaml
```

- c. Delete the PDB:

```
oc delete pdb [***PDB-NAME-OF-CLUSTER-PROXY***] -n [***PRIVATE-CLOUD-NAMESPACE***]
```

Repeat the steps to delete the cdp-release-cpx-liftie PDB as well.

Cumulative hotfixes

Review the cumulative hotfixes for Cloudera Data Services on premises 1.5.5-CHF1.

Cloudera Data Services on premises 1.5.5-CHF1 supports:

- Cloudera Base on premises 7.1.9, 7.1.9 SP1 CHF5, 7.1.7 SP3 CHF10, and 7.3.1.400
- Cloudera Manager 7.13.1.402

Cloudera Data Services on premises 1.5.5-CHF1

There are no new features introduced in cumulative hotfix 1.5.5-CHF1.



Note: Before upgrading to Cloudera Data Services on premises 1.5.5 and above with Datalake 7.3.1, please follow the [Cloudera Data Engineering Prerequisites](#) to avoid breaking [Data Engineering's Spark support matrix](#).

Cloudera Data Services on premises 1.5.5 existing known issues are carried into Cloudera Data Services on premises 1.5.5-CHF1. For more details, see [Known Issues](#).

Known issue identified in 1.5.5-CHF1

The following are the known issues identified in 1.5.5 CHF1:

COMPX-20437 - Handle DB connection failures in running RPM and CAM

When a customer upgrades from version 1.5.5 to any 1.5.5 hotfix release, the cluster-access-manager and resource-pool-manager pods can enter a CrashLoopBackOff state if they are not restarted during the upgrade.

Restart the cluster-access-manager pod and then restart the resource-pool-manager pod (order of restart is very important). Commands to restart:

```
kubectl rollout restart deployment <cluster-access-manager-deployment-name> -n <namespace>  
kubectl rollout restart deployment <resource-pool-manager-deployment-name> -n <namespace>
```

Service packs

The service pack issues that have been shipped for Cloudera Data Services on premises 1.5.5-SP1.

Cloudera Data Services on premises 1.5.5 SP1

Review new features, behavioural changes, known issues, and fixed issues introduced in service pack 1.5.5 SP1.



Note: Before upgrading to Cloudera Data Services on premises 1.5.5 SP1 and above with Datalake 7.3.1, please follow the [Cloudera Data Engineering Prerequisites](#) to avoid breaking [Data Engineering's Spark support matrix](#).

What's new in Cloudera Data Services on premises 1.5.5 SP1

Understand the functionalities and improvements to features of Cloudera Data Services on premises 1.5.5 SP1.

Upgrade pre-checks

A new host inspection called Host Prerequisite Inspection is added to the upgrade pre-checks. It verifies that the minimum requirements for the system configurations are met on the hosts. For more information, see [Pre-upgrade checklist](#).

Post Upgrade Validation

This is a set of health checks added as the last step in the upgrade to verify that the upgrade was successful and the cluster is healthy. This includes verifying that the services are healthy, the hosts are upgraded to the new version, and that the Cloudera Control Plane is in a healthy state after upgrade. For the Cloudera Control Plane to be determined as good health, it will verify that RKE2, Longhorn, and Vault are healthy. It will also check that the critical pods in the Cloudera Control Plane namespace are up and running. For more information, see [Upgrading Cloudera Data Services on premises using Cloudera Embedded Container Service](#).

SAML enhancement in Cloudera Management Console

New settings such as Signing and encryption-decryption are added in Cloudera Management Console UI. For more information, see [Setting SAML as IdP](#).

Data Recovery Service (DRS) functionalities

Cloudera Data Services on premises 1.5.5 SP1 and higher versions provide the following DRS features:

- You can view the Backup and Restore Manager on the Cloudera Management Console Backup Manager Control Plane page.

- You can use the S3 locations in Ozone, in AWS buckets, or in both as an external backup location to back up and restore the Cloudera Control Plane.
- DRS automatically backs up Virtual Warehouse configurations from the database along with the namespace backup. You can restore the backup, when required.

For more information, see [Data Recovery Service overview](#).

Cloudera Embedded Container Service containerd And kubelet health tests

Implemented health test Runners for both the Ecs Server and Ecs Agent roles within the Cloudera Embedded Container Service. These runners executed by SMON on each Ecs Server and Ecs Agent to read metrics and determine the health status containerd and kubelet, which is then displayed on the Cloudera Embedded Container Service Roles page in the Cloudera Manager UI.

Metrics are collected (60 sec of interval) by process running alongside the Ecs Server and Ecs Agent process and sent to SMON.

For more details, see [ECS Agent Health test](#) and [ECS Server Health tests](#).

Fixed Issues for the Cloudera Data Services on premises 1.5.5 SP1

You can review the list of reported issues and their fixes in Cloudera Data Services on premises 1.5.5 SP1. Fixed issues represent selected issues that were previously logged through Cloudera Support, but are now addressed in the current Cloudera Data Services on premises release. These issues may have been reported in previous versions of Cloudera Data Services on premises as a known issue; meaning they were reported by customers or identified by Cloudera Quality Engineering teams.

OPSX-6589 - Cloudera Embedded Container Service 1.5.4 or 1.5.5, rke2_launch startup script fails to run because root is disallowed to use sudo

Startup failure in Cloudera Embedded Container Service 1.5.4 or 1.5.5 when sudo is disallowed for root. In 1.5.4 SP2, the rke2_launch startup script was updated to use sudo for setting the default NFSv4 configuration. This change caused Cloudera Embedded Container Service startup failures in environments where sudo is disabled for all users (including root).

The rke2_launch installer script has been updated so that inline sudo commands are no longer required and startup now proceeds without needing sudo.

OPSX-6049 - Diagnostics service did not purge logs from /data/env

Previously, the diagnostics service only purged logs under /data/bundles and /data/dp. Log-purging functionality has been extended to also include the /data/env directory.

OPSX-6602 - Break the dependency between unseal vault and restart validations steps

The restart validation step waits for pods in the cdp namespace to become Ready. However, some pods in this namespace depend on vault being unsealed first. Since the vault unsealing occurs after this step, the validation process waits until it times out unless vault is unsealed manually.

To prevent this issue, the validation of the cdp namespace is excluded from the restart validation to eliminate the manual unseal dependency.

OPSX-6349 - Upgrade logic missing for nvidia-device-plugin in Cloudera Embedded Container Service

After upgrading from Cloudera Embedded Container Service 1.5.2 to 1.5.4 (or later), Nvidia pods continued using legacy images because only install logic (not upgrade logic) was present.

The installer now includes a pre-upgrade step that upgrades the nvidia-device-plugin. After this change, upgrades of Cloudera Data Services on premises will ensure the plugin is updated too.

OPSX-6305 - Removal of SUID bit from /opt/cni/bin/install after Cloudera Embedded Container Service upgrade

The issue was identified during upgrades from older Calico versions, which left the binary with SUID permissions, posing a security risk. The fix involved removal of the SUID bit from the /opt/cni/bin/install binary after Cloudera Embedded Container Service upgrades to ensure compliance with security policies.



Note: The change applies specifically to Cloudera Embedded Container Service environments and does not impact OCP, with ongoing verification for Kubernetes 1.31 clusters.



Note: Latest updates will remove the binary entirely in newer Calico versions (v3.28.0+), but current upgrades from older versions require manual or scripted cleanup.

OPSX-6245 - Airgap | Multiple pods are in pending state on rolling restart

When performing consecutive rolling restarts on Cloudera Embedded Container Service clusters, the kube-controller-manager pod sometimes fails to become Ready promptly. This prevents other critical pods (including Vault) from initializing, causing the Vault-unseal step to fail.

The installer now implements pre-stop and post-start actions for Cloudera Embedded Container Service restarts in air-gapped environments to ensure pods return to Ready before proceeding.

OPSX-6241 - rke2-ingress-nginx-controller pod fails after Cloudera Embedded Container Service upgrade

After upgrading Cloudera Embedded Container Service, the ingress controller pod failed to start reliably. The restart sequence for the ingress component now includes forced termination of ingress pods, followed by scaling down and scaling up of the rke2-ingress-nginx-controller pod to ensure correct startup.

OPSX-5216 - Unsupported 'internal alias for registry' option

The internal alias for registry setting in Cloudera Embedded Container Service was enabled, but this configuration is not supported when running Cloudera AI workloads on fresh clusters.

Do not set the internal alias for registry flag if you intend to run Cloudera AI workloads.

OPSX-4763 - Environment state not reset after failed creation

In EnvironmentServiceAsyncUtils, when hive warehouse external directory creation failed, the isAvailable flag was set to false. While the subsequent API calls succeeded, the environment remained in a non-AVAILABLE state. The fix ensures that after failure and retry, the environment can be correctly marked as AVAILABLE when all tasks succeed.



Note: A pod#restart of the environment service may be required in earlier deployments.

OPSX-3323 - Custom log redaction does not work for JSON files in the diag bundle

Redaction rules (configured through the admin console) were not applied to JSON files in the diagnostic bundle, although they worked for .log and .txt files.

The diagnostic-bundle logic has been refactored so that JSON files are now included in log redaction. No behavioural change; only the file coverage is extended.

OPSAPS-74598 - Removal of Internal Alias Configuration from Cloudera Embedded Container Service Installation Wizard

ECS Internal Alias configuration is removed from Cloudera Embedded Container Service installation wizard, so that you do not enable this configuration.

OPSAPS-75358 - Configure the 'worker-shutdown-timeout' value for nginx ingress controller to 15 minutes

Based on the analysis for sample workloads, the worker-shutdown-timeout parameter for the nginx ingress controller has been lowered to 15 minutes to ensure graceful shutdown.

OPSX-6227 - Environment should be marked as FAILED if one of the async tasks fails

Previously, if one asynchronous task failed during environment creation, the environment remained in a CREATED state rather than FAILED, causing downstream workflows to proceed incorrectly.

The environment service now transitions the environment to FAILED status if any async task fails during creation.

OPSX-3953 - Upgrade postgresql image used in cdp-embedded-db

The embedded DB PostgreSQL component is upgraded from version 10.x to 17.4 as part of the Cloudera Control Plane upgrade. When Cloudera Control Plane 1.5.5 SP1 is installed or Cloudera Control Plane 1.5.x is upgraded to 1.5.5 SP1, cdp-embedded-db will participate in the Cloudera Control Plane upgrade. If the prior version is 1.5.x, then the data is automatically migrated to the new PostgreSQL 16 database.

Known issues for the Cloudera Data Services on premises 1.5.5 SP1

You must be aware of the known issues and limitations, the areas of impact, and workaround in Cloudera Data Services on premises 1.5.5 SP1 release.

Cloudera Data Services on premises 1.5.5 existing known issues are carried into Cloudera Data Services on premises 1.5.5 SP1. For more details, see [Known Issues](#).

Known issue identified in 1.5.5 SP1

The following are the known issues identified in 1.5.5 SP1:

OPSAPS-70612: Invalid URL error while installing Cloudera Data Services on premises 1.5.5 SP1 from Cloudera Manager 7.13.1.501

When attempting to install Cloudera Data Services on premises using a Cloudera Manager 7.13.1.501 hotfix, the installation may fail.

The installation fails to process the repository URL correctly in this specific scenario, resulting in an "Invalid URL" error, which blocks the installation.

Perform the following steps to mitigate this issue:

1. Login to Cloudera Manager and click **Hosts Add Hosts**.
2. In the **Add Hosts** page, select the **Add hosts to Cloudera Manager** option and then click **Continue**.
3. Follow the instructions on the **Setup Auto-TLS** page and then click **Continue**.
4. In the **Specify Hosts** page, enter a host name or pattern to search for new hosts to add to the cluster, and then click **Continue**. You can click the **Patterns** link for more information.

A list of matching hosts are displayed.

5. Select the hosts that you want to add and click **Continue**.
6. Select the repository location where Cloudera Manager can access the software required to install on the new hosts. Choose either the **Public Cloudera Repository** or a **Custom Repository**, and provide the URL of the custom repository available on your local network.
7. Click **Continue**.

After completing the above tasks, you can proceed with the Cloudera Data Services on premises installation.

OPSX-6794 - Upgrading from Cloudera Data Services on premises 1.5.4 to 1.5.5 Service Pack 1 fails with a Upgrade embedded DB error message midway through the upgrade process.

When attempting to upgrade from Cloudera Data Services on premises 1.5.4 to 1.5.5 Service Pack 1, the upgrade process could fail.

If you have performed a manual **Expand Volume** operation on the cdp-embedded-db-backend Persistent Volume Claim (PVC) of the cdp-embedded-db service in **cdp** namespace, it can cause the upgrade process to fail. The upgrade workflow attempts to apply the original value that was used during the initial cluster set up. As Longhorn does not support reducing the volume size, the upgrade process fails.

To avoid stalling the ongoing upgrade process, you must manually exclude the PVC for the cdp-embedded-db-backend in the upgrade process.



Important: You must abort the upgrade process before applying this workaround.

Perform the following steps:

```
#Login to ECS master node
ssh root@ecs_master_node

#Switch to embedded db directory in the ECS parcel
cd /opt/cloudera/parcels/ECS-1.5.5-h1000-b167-ecs-1.5.5-h1000-b167.p0.72522329/installer/helm/cdp-embedded-db

#Create a tmp working directory
mkdir temp-cdp-embedded-db-1.5.5-h1000-b167

#Extract the tgz contents into the temp working directory
tar -xzf cdp-embedded-db-1.5.5-h1000-b167.tgz -C ./temp-cdp-embedded-db-1.5.5-h1000-b167

#Manual changes
vi temp-cdp-embedded-db-1.5.5-h1000-b167/templates/persistentVolumeClaim.yaml

#This file will have two PVCs. cdp-embedded-db-backend and cdp-embedded-db-backend-v17. Remove cdp-embedded-db-backend PVC from the file. Save the file and exit

#Take a backup of the old tgz file
mv cdp-embedded-db-1.5.5-h1000-b167.tgz backup-cdp-embedded-db-1.5.5-h1000-b167.tgz

#Switch to temp working directory
cd temp-cdp-embedded-db-1.5.5-h1000-b167

#Prepare an updated tgz file
tar -czvf ../cdp-embedded-db-1.5.5-h1000-b167.tgz .

#Switch to previous directory and check the tgz file
cd ..
ls -lrth

#Move the backup tgz to a different directory
mv backup-cdp-embedded-db-1.5.5-h1000-b167.tgz /tmp/some_directory

#Cleanup the temp working directory
rm -rf temp-cdp-embedded-db-1.5.5-h1000-b167/

#Backup the following secret in cdp namespace to 'some_directory'. Make sure to replace 'some_directory' with an appropriate path
kubectl get secret cdp-thunderheaduserpreference-db-secret -n cdp -o yaml > /tmp/some_directory/backup-cdp-thunderheaduserpreference-db-secret

#Delete the following secret in cdp namespace
kubectl delete secret cdp-thunderheaduserpreference-db-secret -n cdp
```

Once you complete these steps, you can resume the upgrade process.

OPSX-6797 - A possible upgrade failure resulting from expanding the cdp-embedded-db volume from Longhorn UI after the initial installation of your Data Services on premises cluster

If you have expanded cdp-embedded-db volume from Longhorn UI after the initial installation of your Data Services on premises cluster, you must complete the workaround steps before planning your upgrade to Cloudera Data Services on premises 1.5.5 Service Pack 1 to avoid a potential upgrade failure.

Login to your ECS master node on your cluster and perform the following these steps:

```
# Find the size of cdp-embedded-db-backend PVC, look for the value in CAPACITY column
kubectl get pvc -n cdp | grep cdp-embedded-db-backend
# Find the latest/current copy of cdp-private-installer-values
# If the cluster is on 1.5.5-b212, then the secret to use will be cdp-private-installer-values-1.5.5-b212
kubectl get secret -n cdp | grep cdp-private-installer-values
# For this example, we will use 'cdp-private-installer-values-1.5.5-b212'
# Please use the correct one from your cluster

# Backup the existing values, in case you need to restore it later
kubectl get secret cdp-private-installer-values-1.5.5-b212 -n cdp -o yaml > backup-cdp-private-installer-values-1.5.5-b212.yaml

# Make a new copy before updating the values
cp backup-cdp-private-installer-values-1.5.5-b212.yaml updated-cdp-private-installer-values-1.5.5-b212.yaml
# Edit 'updated-cdp-private-installer-values-1.5.5-b212.yaml' using the following instructions:
# 1. Base64decode the value present in 'values.yaml.merged' attribute
# 2. Update Database.EmbeddedDbStorage to the latest value obtained in the first step
# 3. Base64encode the updated yaml string and update it in 'values.yaml.merged' attribute
# 4. Delete metadata.creationTimestamp, metadata.name, metadata.resourceVersion and metadata.uid attributes
# 5. Save these changes to updated-cdp-private-installer-values-1.5.5-b212.yaml

# Delete the existing secret
kubectl delete secret cdp-private-installer-values-1.5.5-b212 -n cdp

# Create a new secret using the updated file
kubectl apply -f updated-cdp-private-installer-values-1.5.5-b212.yaml

# Check if the secret got created
kubectl get secret -n cdp | grep cdp-private-installer-values

# Verify if the value got updated
kubectl get secret cdp-private-installer-values-1.5.5-b212 -n cdp -o jsonpath='{.data.values\.yaml\.merged}' | base64 --decode
```

COM PX-20437 - DB connection failures causing RPM and CAM pods to CrashLoopBackOff

During an upgrade from version 1.5.5 to any 1.5.5 hotfix release, the cluster-access-manager (CAM) and resource-pool-manager (RPM) pods can enter a CrashLoopBackOff state if they are not automatically restarted during the upgrade.

After upgrade, manually restart the CAM pod and then restart the RPM pod (order of restart is very important).

Commands to restart:

```
kubectl rollout restart deployment <cluster-access-manager-deployment-name> -n <namespace>
kubectl rollout restart deployment <resource-pool-manager-deployment-name> -n <namespace>
```

OBS-9491 - Prometheus configuration exceeds size limit in large environments

In environments with a large number of namespaces (approximately 300 or more per environment), the Prometheus configuration for Cloudera Monitoring might exceed the 1 MB Kubernetes Secret size limit. If the total size, which depends on factors such as the number of namespaces, the length of namespace names, their variability, and the size of the certificate store, exceeds 1 MB, the new Prometheus configuration will not be applied, and new namespaces will not be monitored. As a result, the telemetry data will not be collected from those namespaces and will not be reflected on the corresponding Grafana charts.

To resolve this issue, you must enable Prometheus configuration compression at the Cloudera Control Plane level.

1. Upgrade to Cloudera Data Services on premises 1.5.5 SP1 or a higher version.
2. Back up the secret from the monitoring namespace to allow for rollback and disabling compression later.
3. Set the environment variable `ENABLE_ENVIRONMENT_PROMETHEUS_CONFIG_COMPRESSION` to "true" on the `cdp-release-monitoring-pvc-service` deployment in the Cloudera Control Plane namespace.



Note: The upper bound for monitoring is based on the number of namespaces per environment, the length of the namespace name, the variability of the names, and the size of the cert store.

- For example, about 700 workload namespaces in a single environment with a short, consistent pattern ("aaaa-bbbb-XXXX") generate around 150 KB of configuration, which fits within the 1 MB limit.
- About 3000 namespaces in a single environment with a longer, consistent pattern ("aaaa-cccc-dddd-XXXXX") generate about 900 KB.

Rolling back the workaround:

1. Set the environment variable `ENABLE_ENVIRONMENT_PROMETHEUS_CONFIG_COMPRESSION` to "false" on the `cdp-release-monitoring-pvc-service` deployment in the Cloudera Control Plane namespace.
2. Reapply the secret from the backup.

OPSX-6618 - In Cloudera Embedded Container Service upgrade not all volumes are upgraded to the latest longhorn version

During restart of the Cloudera Embedded Container Service cluster from 1.5.5 to 1.5.5 SP1, the upgrade failed due to longhorn health issues. This is because one of the volumes was degraded.

Follow the steps to resolve the issue:

- Identify the problematic volumes (degraded state).
- Set the value of `spec.numberOfReplicas` of the volume to the number of active replicas. For example, set the value to 3 if two replicas are active.
- Apply the fix before or during the upgrade as per the workaround instructions (refer to Longhorn issue #11825). Longhorn Engineering is addressing in v1.11.0 and will back-port the fix. The workaround is included as part of the upgrade. However, if the issue is noticed even after

upgrade please follow the workaround steps documented here: <https://github.com/longhorn/longhorn/issues/11825>

COMPX-23842 / COMPX-24130/ COMPX-23319 - Pod status shown as OutOfcpu, OutOfmemory, or Pending after a cluster restart

During a cluster restart, due to an upgrade or normal maintenance, all nodes in the cluster are restarted. During this process, the cluster operates with reduced resource capacity. When this occurs, pod placement can be rejected, resulting in some pods entering OutOfcpu or OutOfmemory state.

Do not restart failed pods until the cluster restart process has fully completed. Once the cluster is fully operational:

- Restart any pods that are in a failed (OutOfcpu or OutOfmemory) state, so they can be rescheduled properly.
- If any pods remain in a pending state after the restart, restart the YuniKorn scheduler to refresh scheduling:

```
kubectl rollout restart deployment yunikorn-scheduler -n yunikorn
```

OPSX-6566 - Cloudera Embedded Container Service restart fails with etcd connectivity issues

Restart of the Cloudera Embedded Container Service server fails with etcd error: "error reading from server"

1. Identify the server role that failed.
2. Restart only the Cloudera Embedded Container Service server role which failed.
3. Once the server role is restarted and healthy, proceed with the remainder of the Cloudera Embedded Container Service server role restart sequence.

OPSX-6401 - Istio ingress-default-cert is not created in the upgrade scenario

After upgrading to 1.5.5 SP1, the Secret ingress-default-cert is not created in the istio-ingress namespace. Because this certificate is expected, failing to create it causes components like CAII & MR provisioning to fail.

To resolve this issue:

- After upgrading from an older version to 1.5.5 SP1, check for the secret called ingress-default-cert in the istio-ingress namespace (ns).

For example:

```
kubectl get secret ingress-default-cert -n istio-ingress
```

- If missing, copy the identical secret from kube-system namespace (or from the pre-upgrade environment) into istio-ingress namespace with the same name and contents.
- This will restore the expected certificate and allow CAII & MR provisioning to proceed.

OPSX-6645 - Cloudera Embedded Container Service upgrade failure at restart step

When the Cloudera Embedded Container Service role fails to start after a node reboot or service restart, the root cause can be that the etcd defragmentation process which runs on startup takes longer than the component timeout thresholds. As a result:

- The kubelet service may fail to start or time out.
- The kube-apiserver, kube-scheduler or kube-controller-manager roles may remain in a NotReady state.
- etcd may perform automatic defragmentation at startup.
- The API server may fail to connect to etcd (connection refused or timeout).

To fix this issue:

1. Resume the Cloudera Embedded Container Service start/restart from the Cloudera Manager UI once etcd has come up.
2. Ensure etcd meets production hardware and configuration requirements: For more information, see [etcd requirements](#) and [Knowledge Base](#).

OPSX-6767 - Cloudera Embedded Container Service cluster has stale configuration after Cloudera Manager upgrade to 7.13.1.501-b2 from 7.11.3.24

After upgrading Cloudera Manager to version 7.13.1.501, the Cloudera Embedded Container Service shows a staleness indicator. This occurs due to configuration changes applied by the upgrade:

- worker-shutdown-timeout: reduced from 24 hours (86,400 s) to 15 minutes (900 s).
- smon_host: a new monitoring configuration added.
- smon_port: a new monitoring port configuration (9997).

No action required — the staleness is expected following this upgrade and can be safely ignored. The indicator will automatically clear once Cloudera Embedded Container Service is upgraded to version 1.5.5 SP1 or later.

If you prefer to clear the staleness indicator right away, you may manually refresh the Cloudera Embedded Container Service service through the Cloudera Manager UI.



Note: Performing a refresh will restart the ingress controller and may cause a brief service interruption.

OPSX-6638 - Post rolling restart many pods are stuck in pending state

Pods remain in Pending state and fail to schedule with etcd performance warnings.

Restart the Cloudera Embedded Container Service master role which has etcd performance issues. For more information, see [etcd hardware recommendations](#) and [Knowledge Base](#).