

# Cloudera Management Console on premises Release Notes

Date published: 2023-12-16

Date modified: 2025-06-06

# 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.....4**

**Known issues.....5**

## 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
r 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:

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

Inspect each secret using the command:

```
kubectl 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:

```
kubectl 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 kubectl command on the ECS server host to check if the pod is stuck in a running state:

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

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

```
kubectl 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
```

### OPX-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.

## OPSX-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 `kubect`l 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.

## OPsx-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.

### **PULSE-944 and PULSE-941 Cloudera Observability namespace not created after platform upgrade from 151 to 152**

The Cloudera Observability namespace is not created after a platform upgrade from Cloudera Observability 1.5.1 to Cloudera Private Cloud Data Services 1.5.2.

During the creation of the resource pool the Cloudera Observability namespace is provided by the Cloudera on premises. If the provisioning flow is not completed, such as due to a timing difference between the start of the computeAPI pod and the call to the computeAPI pod by the service, the namespace is not created.

Trigger the Cloudera Observability namespace deployment by restarting the pvcservice pod.

### **PULSE-921 Cloudera Observability namespace has no pods**

The Cloudera Observability namespace should have the same number of pods and nodes. When the Cloudera Observability namespace has no pods the prometheus-node-exporter-1.6.0 helm release state becomes invalid and the Cloudera Data Services on premises is unable to uninstall and reinstall the namespace. Also, as the Node Exporter is not installed into the Cloudera Observability namespace its metrics are unavailable when querying Prometheus in the control plane, for example the node\_cpu\_seconds\_total metric.

Manually uninstall the invalid helm release with the --debug flag, verify that there are no helm releases listed by running `-n observability -a`, and then trigger the deployment process by restarting the pvcservice pod in the control plane.

### **PULSE-697 Add node-exporter to Cloudera Data Services on premises**

When expanding a cluster with new nodes and there is insufficient CPU and memory resources, the Node Exporter will encounter difficulties deploying new pods on the additional nodes.

To ensure sufficient resource allocation, such as when the Cloudera Observability namespace requires adjustment, delete the existing namespace and restart the pvcservice pod. This automatically initiates the creation of the Cloudera Observability namespace with the appropriate resource allocation.



**Note:** During the namespace recreation process the Node Exporter metrics are temporarily unavailable.

#### **PULSE-935 Longhorn volumes are over 90% of the capacity alerts on Prometheus volumes**

Cloudera Manager displays the following alert about your Prometheus volumes: Concerning: Firing alerts for Longhorn: The actual used space of Longhorn volume is over 90% of the capacity.

Longhorn stores historical data as snapshots that are calculated with the active data for the volume's actual size. This size is therefore greater than the volume's nominal data value.

When the alert is displayed on the Cloudera Manager UI and it is related to Longhorn volumes used by Prometheus, ignore. For more information, see the Longhorn space consumption guidelines in the Longhorn documentation.

#### **PULSE-937 Private-Key field change in Update Remote Write request does not reflect in enabling the metric flow**

When using the Cloudera Management Console UI for Remote Storage the Disable option does not deactivate the remote write configuration, even when the action returns a positive result message. Therefore, when disabling a remote storage configuration use the CLI client to disable the remote storage configuration directly from the API.

At this time when a remote storage configuration is incorrect, do not use the Edit or Disable option from the configuration's Actions menu (ellipsis icon) to change its configuration. Instead, delete the remote storage's configuration from the configuration's Actions menu with the Remove Configuration action and then re-create the remote write configuration with the Delete and Create operations of the API, using the CLI client.

#### **PULSE-841 Disabling the remote write configuration logs an error in both cp prometheus and env prometheus**

When a metric replication is set up between the cluster and Cloudera Observability and the connection is disabled or deleted, Prometheus writes an error message that states that it cannot replicate the metrics.

No workaround is required. After a few minutes the errors are no longer logged and Prometheus no longer tries to replicate the metrics.

#### **PULSE-895 Disabling the remote write config in the UI is broken in Cloudera Private Cloud Data Services**

The Remote Write Enable and Disable options in the Cloudera Management Console's User Interface do not work when a Remote Storage configuration is created with a requestSignerAuth type from either the HTTP API or using the CDP-CLI tool.

At this time, do not use the Enable or Disable options from the Remote Storage configuration's Actions menu in the Cloudera Management Console's UI. Instead, enable or disable the configuration from the HTTP API or using the CDP-CLI tool.

#### **PULSE-936 No Alert to prompt the metric flow being affected b/c of wrong private key configuration**

A remote write alert was not triggered when the wrong private key was used in a Remote Storage configuration.

No workaround. Incorrect configuration settings, such as in this case where a bad private key was used, may block the forwarding of metrics. When creating a Remote Storage configuration you must carefully verify each configuration setting.

### **Known Issues identified in 1.5.1**

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

As of Cloudera Private Cloud Data Services 1.5.1, external Control Plane metadata databases are no longer supported. New installs require the use of an embedded Cloudera Control Plane database. Upgrades from Cloudera 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 Cloudera Control Plane database to the embedded Cloudera Control Plane database. Upgrades from 1.4.0 or 1.5.0 with external Cloudera Control Plane metadata databases also require additional steps, which are described in the Cloudera Private Cloud Data Services 1.5.1 upgrade topics.

**DOCS-15855: Networking API is deprecated after upgrade to Cloudera 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 to fetch the correct HA Proxy image**

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.

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

**OPSX-4266: Cloudera Embedded Container Service upgrade from 1.5.0 to 1.5.1 is failing in Cadence schema update job**

When upgrading from Cloudera Embedded Container Service 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.

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
```

#### 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-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 Cloudera Data Services on premises 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.

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

#### 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.

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-3073 Cloudera Embedded Container ServiceFirst 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      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-1405: Able to create multiple Cloudera on premises 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 Cloudera 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.

**Known Issues identified in 1.5.0**

**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 Cloudera Embedded Container Service 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.

**Known Issues identified before 1.5.0**

**OPSX-5629: COE Insight from case 922848: Not able to connect to bit bucket**

After installing Cloudera AI on an Cloudera Embedded Container Service cluster, users were not able to connect the internal bitbucket repo.

Workaround:

In this case the MTU of the Cloudera Embedded Container Service virtual network interfaces were larger than that of host external interface, which may cause the network requests from Cloudera Embedded Container Service 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 Cloudera Embedded Container Service 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.

#### **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.

None

#### **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.

#### **Recover fast in case of a Node failures with Cloudera Embedded Container Service 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 Data Services on premises Cloudera Embedded Container Service: Failed to perform First Run of services.**

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

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 Cloudera Private Cloud Data Services deployment when the Cloudera Base on premises cluster is in a degraded state and includes Ozone service.

Workaround: Stopping the Ozone service temporarily in the Cloudera Base on premises 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.

**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 Cloudera on premises. 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.

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

In the Cloudera 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.

None