

CDP Private Cloud Data Services 1.5.3

CDP Private Cloud Data Services Management Console Release Notes

Date published: 2023-12-16

Date modified: 2024-03-05

CLUSTERA

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

Legal Notice

© Cloudera Inc. 2024. All rights reserved.

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

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

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

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

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

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

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

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

Contents

What's new.....	4
Known issues.....	4
Fixed issues.....	24

What's new

New features in the 1.5.3 release of the CDP Private Cloud Management Console service.

Certifications

OCP 4.12

K8s 1.26

Known issues for the CDP Private Cloud Data Services Management Console

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

Known Issues in Management Console 1.5.3

DOCS-20088/OPSX-4781: Vault pods may take long time to be ready during upgrades from 1.5.2 to 1.5.3

The 'vault-0' pod takes longer time to attach volume in some upgrade cases than usual. Due to the excess time taken the cluster upgrade may fail. But, usually in 15 minutes the volume can attach automatically and the pod would start running. In that case, the user can resume the upgrade.

No workaround available.

OPSX-4777: [151h2-153]post ECS upgrade Longhorn health test failing - helm-install-longhorn pod in crashloop state

After upgrading CDP Private Cloud Data Services on ECS, the Longhorn health test fails with the helm-install-longhorn pod in crashloop state.

To fix this issue:

```
helm history longhorn -n longhorn-system
      REVISION      UPDATED                               STATUS
  CHART              APP VERSION   DESCRIPTION
  longhorn-1.4.2    v1.4.2        Install complete
  longhorn-1.5.4    v1.5.4        Deletion in progress (or
silently failed)

      NAME                                     TYPE
  DATA  AGE      basic-auth                                     Opaque
  1      15h
      chart-values-longhorn                     Opaque
  0      10h
      longhorn-webhook-ca                       kubernetes.io/tls
  2      15h
      longhorn-webhook-tls                     kubernetes.io/tls
  2      15h
      sh.helm.release.v1.longhorn.v1          helm.sh/release.v1
  1      15h
      sh.helm.release.v1.longhorn.v2          helm.sh/release.v1
  1      21m
```

```

kubect1 get secrets sh.helm.release.v1.longhorn.v2
-n longhorn-system -o yaml > sh.helm.release.v1.longhorn.v2.yaml
helm get values --revision=2 longhorn -n longhorn-s
ystem > defaultSettings.yaml

kubect1 get jobs -n longhorn-system
NAME                                COMPLETIONS  DURATION
AGE
helm-install-longhorn              0/1           9h          9h
longhorn-post-upgrade             1/1           11m         10h
h
longhorn-uninstall                 0/1           10h
10h

kubect1 delete job helm-install-longhorn longhorn-
uninstall longhorn-post-upgrade -n longhorn-system
kubect1 delete secret sh.helm.release.v1.longhorn.v
2 -n longhorn-system

kubect1 patch HelmChart longhorn -n longhorn-syste
m --type=merge --patch-file /opt/cloudera/parcels/ECS/longhorn/l
onghorn.yaml

```

OPSX-4754 [ECS Restart Stability] DaemonSet rollout is stuck post rolling restart - DaemonSet kube-system/rke2-canal has not finished or progressed for at least 15 minutes

On RHEL 9.1, an ECS service DaemonSet rollout health alert appears after a rolling restart.

To fix the DaemonSet rollout issue:

1. Edit the DaemonSet file: to alter the value from auto to Legacy as shown below:

```
$KUBECTL -n kube-system edit ds/rke2-canal
```

Change the value of felixIptablesBackend to Legacy and save the DaemonSet file.

2. Reboot each node one-by-one.
3. Check to see if any of the nodes are cordoned off. If so, uncorordon them:

```

[root@host-1 ~]# $KUBECTL get nodes
NAME STATUS ROLES AGE VERSION
host-1.ecs-restart1.kcloud.cloudera.com Re
ady,SchedulingDisabled control-plane,etcd,mas
ter 17h v1.26.10+rke2r1
host-2.ecs-restart1.kcloud.cloudera.com R
eady <none> 17h v1.26.10+rke2r1
host-3.ecs-restart1.kcloud.cloudera.com Re
ady <none> 17h v1.26.10+rke2r1
host-4.ecs-restart1.kcloud.cloudera.com Rea
dy <none> 17h v1.26.10+rke2r1
[root@host-1 ~]# $KUBECTL uncordon host-1.ec
s-restart1.kcloud.cloudera.com
node/host-1.ecs-restart1.kcloud.cloudera.com
uncordoned

[root@host-1 ~]# $KUBECTL get nodes
NAME STATUS ROLES AGE VERSION
host-1.ecs-restart1.kcloud.cloudera.com Rea
dy control-plane,etcd,mas
ter 17h v1.26.10+rke2r1
host-2.ecs-restart1.kcloud.cloudera.com R
eady <none> 17h v1.26.10+rke2r1
host-3.ecs-restart1.kcloud.cloudera.com Re
ady <none> 17h v1.26.10+rke2r1
host-4.ecs-restart1.kcloud.cloudera.com Rea
dy <none> 17h v1.26.10+rke2r1

```

```
[root@host-1 ~]#
```

4. Ensure that the Vault is unsealed. If it is not, unseal the vault in Cloudera Manager by selecting Cluster > ECS > Actions > Unseal Vault.
5. Wait 5-6 minutes.

DOCS-19913: OCP upgrade – OCP namespace name must be 29 characters or less

Before upgrading on OCP, ensure that the OpenShift namespace name is 29 characters or less. Do not proceed with the upgrade if the namespace name is 30 or more characters in length.

None.

OPSAPS-69892: kube-proxy failure causing issues with cluster

After rebooting/restarting an ECS agent node, the kube-proxy Linux process may not start due to a race condition in the kubelet. When this happens, ECS cluster networking and other services – such as Vault, DNS, authentication, Longhorn storage, etc. – are affected. At the Kubernetes pod level, errors such as "connection refused", "connection timed out" and "i/o timeout" may be observed. If you suspect possible networking issues in your ECS cluster, checking kube-proxy is a good first step.

To fix this issue, perform the following steps on all of the affected nodes:

1. To identify which agent needs to be restarted, check the status of each kube-proxy pod to make sure it is in the "ready" state by running the following command on each host in the cluster.

```
kubectl describe pod host-1.test.cloudera.com -n kube-system
```

In the Conditions section of the describe pod output, confirm that the "ready" condition is "True".

```
Conditions:
  Type                Status
  Initialized         True
  Ready               True
  ContainersReady    True
  PodScheduled       True
```

Another option is to run the following command:

```
kubectl get pods -n kube-system -l component=kube-proxy -o go-template='{{range .items}}
{{.metadata.name}}{"\n"}{{"  "}}{{range .status.conditions}}
{{ if eq .type "Ready" }}
Ready:{{.status}}{"\n\n"}}{{end}}{{end}}{{end}}'
```

The output displays the status of all of the kube-proxy pods in the cluster:

```
// Some comments here
kube-proxy-host-1.cloudera.com
  Ready:True

kube-proxy-host-2.cloudera.com
  Ready:True

kube-proxy-host-3.cloudera.com
  Ready:True
```

- If the "ready" state is False, kube-proxy is not functioning properly, regardless of whether the kube-proxy process is running on that host or not. On each of the affected nodes, run the following command to delete the kube-proxy pod manifest:

```
rm /var/lib/rancher/rke2/agent/pod-manifests/kube-proxy.yaml
```

- Start the agent role.

After the agent role is started, you may not immediately see the kube-proxy process running, but a new kube-proxy process should start shortly. Check the pod status to make sure it is ready. After all of the problem agents have been restarted, the cluster may complain that the vault is sealed – if so, unseal it. At this point, the Control Plane should be functioning properly.

Additional details about this issue are available here: <https://www.suse.com/support/kb/doc/?id=000021284>

OPSX-4754 [ECS Restart Stability] DaemonSet rollout is stuck post rolling restart - DaemonSet kube-system/rke2-canal has not finished or progressed for at least 15 minutes

On RHEL 9.1, an ECS service DaemonSet rollout health alert appears after a rolling restart.

To fix the DaemonSet rollout issue:

- Edit the DaemonSet file: to alter the value from auto to Legacy as shown below:

```
$KUBECTL -n kube-system edit ds/rke2-canal
```

Change the value of felixIptablesBackend to Legacy and save the DaemonSet file.

- Reboot each node one-by-one.
- Check to see if any of the nodes are cordoned off. If so, uncorordon them:

```
[root@host-1 ~]# $KUBECTL get nodes
NAME STATUS ROLES AGE VERSION
host-1.ecs-restart1.kcloud.cloudera.com Ready,SchedulingDisabled control-plane,etcd,master 17h v1.26.10+rke2r1
host-2.ecs-restart1.kcloud.cloudera.com Ready <none> 17h v1.26.10+rke2r1
host-3.ecs-restart1.kcloud.cloudera.com Ready <none> 17h v1.26.10+rke2r1
host-4.ecs-restart1.kcloud.cloudera.com Ready <none> 17h v1.26.10+rke2r1
[root@host-1 ~]# $KUBECTL uncordon host-1.ecs-restart1.kcloud.cloudera.com
node/host-1.ecs-restart1.kcloud.cloudera.com uncordoned
[root@host-1 ~]# $KUBECTL get nodes
NAME STATUS ROLES AGE VERSION
host-1.ecs-restart1.kcloud.cloudera.com Ready control-plane,etcd,master 17h v1.26.10+rke2r1
host-2.ecs-restart1.kcloud.cloudera.com Ready <none> 17h v1.26.10+rke2r1
host-3.ecs-restart1.kcloud.cloudera.com Ready <none> 17h v1.26.10+rke2r1
host-4.ecs-restart1.kcloud.cloudera.com Ready <none> 17h v1.26.10+rke2r1
[root@host-1 ~]#
```

- Ensure that the Vault is unsealed. If it is not, unseal the vault in Cloudera Manager by selecting Cluster > ECS > Actions > Unseal Vault.
- Wait 5-6 minutes.

OPSX-4766: [ECS Restart] Host Reboot | start command failed with error - "Timed out waiting for kube-apiserver to be ready"

In an ECS cluster with HA enabled, ECS Start fails with an error after stopping the cluster and rebooting the hosts.

Steps to reproduce:

1. Stop ECS.
2. Reboot hosts.
3. Start ECS.

The start command fails with the following error message:

"Timed out waiting for kube-apiserver to be ready"

Option 1:

Start each master role instance individually without waiting each node to be up and running.

Option 2:

If Option 1 does not work, follow the steps from SUSE to recover the cluster:https://docs.rke2.io/backup_restore#cluster-reset

GPU Support: RHEL 8.8 only

GPU support is only offered with RHEL 8.8.

Known Issues in Management Console 1.5.2

OPSX-4650: CM - OCP pvc install Wizard - fails if route name is too long

If the cluster name given during Data Services installation on OpenShift is too long, you may encounter an installation failure, and the log will contain the following error: "The Route ... is invalid: spec.host: Invalid value: ...: must be no more than 63 characters"

Specify a shorter cluster name that is less than 63 characters.

OPSX-4369: FreeIPA generated krb5.conf must use a file-based cache

Private Cloud Data Services requires Kerberos to be enabled in the Private Cloud Base cluster. Furthermore, the /etc/krb5.conf file must be configured to use a file-based cache. Using a keyring-based cache or a KCM-based cache is not supported. When using FreeIPA, the krb5.conf file may be configured to use a KCM-based cache by default. This should be changed.

Ensure that the /etc/krb5.conf file on all hosts uses a file-based cache. This can be checked by looking for the default_ccache_name property and ensuring that its value begins with "FILE:". Any other type of cache, including those starting with "KEYRING:" or "KCM:" are not supported and must be changed.

OPSAPS-68923: CM - After CM upgrade from 7.9.5 to 7.11.3.x ECS cluster showing stale config

After Cloudera Manager upgrade from 7.9.5 to 7.11.3.x, an ECS 1.5.0 cluster may show a stale config to add ""limit_fds": 1048576"

This can be ignored – no restart of the ECS cluster is necessary. When the ECS 1.5.0 cluster is upgraded to 1.5.2, the stale config will be resolved.

COMPX-15475: [CM ECS UPG][150-152] post upgrade prometheus-node-exporter-1.6.0 pod stuck in pending state

As part of the upgrade all nodes in the cluster are restarted. If a set of pods remain in the pending state after the node restart, a YuniKorn restart is required.

You can use the following commands to restart YuniKorn:

```
kubectl scale deployment yunikorn-scheduler --replicas=0 -n yunikorn
kubectl scale deployment yunikorn-scheduler --replicas=1 -n yunikorn
```


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.

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

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

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

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

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

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

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

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

- PROCESSES
- PROCESSES_DETAIL
- COMMANDS_DETAIL

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

```
cm=> DELETE FROM COMMANDS where command_id=1546340765;  
ERROR: update or delete on table "commands" violates foreign  
key constraint "fk_commands_detail_command" on table "comm  
nds_detail"  
DETAIL: Key (command_id)=(1546340765) is still referenced f  
rom table "commands_detail".  
cm=>  
cm=> DELETE FROM commands_detail where command_id=1546340765;  
DELETE 1  
cm=> DELETE FROM COMMANDS where command_id=1546340765;  
DELETE 1
```

- Restart the Cloudera Manager server.

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" <<<<<<<<<< new config  
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 OCP, you may be able to configure this using [HAProxy with X-forward-for pass to OpenShift 4](#).

OPsx-4446: Duplicate Entries in cdp-pvc-truststore

Sometimes the cdp-pvc-truststore contains duplicate entries causing a 3M Request Entry to Large error.

Duplicate entries can be manually removed.

OPsx-4552: [ECS Restart] One of the docker servers failed to come up after starting the cluster post hosts reboot

At times the Docker server may fail to come up and return the following error message:

```
/var/run/docker.sock: Is a directory
```

On the Docker server role host, remove the /var/run/docker.sock directory, then restart the Docker server role.

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 OCP 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 Private Cloud 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 Private Cloud 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 Observability namespace not created after platform upgrade from 151 to 152

The Cloudera Observability namespace is not created after a platform upgrade from PvC DS 1.5.1 to PvC DS 1.5.2.

During the creation of the resource pool the Cloudera Observability namespace is provided by the CDP Private Cloud Service. 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 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 CDP Private Cloud Service 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 PvC DS

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 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 cdp-pvc

The Remote Write Enable and Disable options in the 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 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 in Management Console 1.5.1

External metadata databases are no longer supported on OCP

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

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

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

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

Also, the `df` command may hang on these hosts.

Run the following commands on each unavailable node:

1. Find the NFS mounts:

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

2. Force unmount:

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

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

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

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

CDPVC-1098: How to refresh the YuniKorn configuration

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

When upgrading from ECS 1.5.0 to 1.5.1, the CONTROL_PLANE_CANARY fails with the following error:

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

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

Use the following steps to manually execute the job:

1. Export the job manifest into a file:

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

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

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

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

```
resourceVersion
uid
selector
  matchLabels
```

```
controller-uid
labels
controller-uid
status section
```

4. Create the job:

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

If the job still fails, contact Cloudera Support.

OPXSX-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.

OPXSX-4295:

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

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

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

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

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

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

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

OPXSX-3666: mlx_crud_app DB connection fails with error "unable to create connection: x509: certificate relies on legacy Common Name field, use SANs instead"

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

4. Restart the Docker server role.

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

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

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

1. Scale down the workload. The Longhorn volume will be detached.
2. Wait for the Longhorn volume to be detached.
3. SSH into the node that has the replica.
4. cd into the replica folder (for example, /longhorn/replicas/pvc-126d40e2-7bff-4679-a310-e444e84df267-1a5dc941).

5. Change the "Rebuilding" field from true to false in the volume.meta file.
6. Scale up the workload to attach the volume.

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

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

Events:

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

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

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

Restart the canal pod running on that host:

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

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 CDP PVC Environments with the same name

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

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

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

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

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

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

Custom redaction rule for URLs does not work.

Cloudera Data Engineering service fails to start due to Ozone

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Events:

Type	Reason	Age	From
Message	-----	----	----

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

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

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

Restart the canal pod running on that host:

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Custom redaction rule for URLs does not work.

Cloudera Data Engineering service fails to start due to Ozone

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

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

Known Issues in Management Console 1.4.1

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

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

Workaround:

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

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

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

The default MTU of the virtual network interfaces is 1450.

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

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

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

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

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

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

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

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

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

1. Ensure that the problematic Docker server role has been stopped.
2. Log in to the failing Docker server host.

3. Run the following commands:

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

4. Restart the Docker server role.

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

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

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

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

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

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

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

Custom redaction rule for URLs does not work.

Cloudera Data Engineering service fails to start due to Ozone

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

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

Known Issues in Management Console 1.4.0**Cloudera Data Engineering service fails to start due to Ozone**

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

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

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

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

None.

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

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

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

OPSX-735: Kerberos service should handle CM downtime

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

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

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

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

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

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

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

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

OPX-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

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

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

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

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

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

Cloudera Manager 7.6.1 is not compatible with CDP Private Cloud Data Services version 1.3.4.

You must use Cloudera Manager version 7.5.5 with this release.

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

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

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

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

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

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

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

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

There is currently no workaround for this issue.

Full cluster name does not display in the Register Environment Wizard

None

Kerberos service does not always handle Cloudera Manager downtime

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

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

Management Console allows registration of two environments of the same name

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

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

Not all images are pushed during upgrade

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

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

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

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

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

Updating user roles for the admin user does not update privileges

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

None

Upgrade applies values that cannot be patched

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

None

Incorrect warning about stale Kerberos client configurations

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

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

Vault becomes sealed

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

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

Fixed Issues for the CDP Private Cloud Data Services Management Console

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

Fixed Issues in Management Console 1.5.3

PULSE-697: Add node-exporter to PvC DS

Fixed the issue that occurred when expanding a cluster with new nodes and there was insufficient CPU and memory resources, the Node Exporter encountered difficulties deploying new pods on the additional nodes.

OPSAPS-69556: 1.5.1->1.5.2 Upgrade[Public Registry+public bits] - fails with ImagePullErrors, registry modified to point to docker-private during upgrade

Previously, when upgrading using the Cloudera public registry with public bits, the Docker registry would incorrectly change to point to docker-private.infra.cloudera.com. This has now been fixed to point to the correct registry.

Fixed Issues in Management Console 1.5.2

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

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

OPSX-4397: The domain name parameter from environment service has an additional 'apps' subdomain

Previously, the domain name of any environment in an Openshift installation contained an additional "apps." subdomain. This has now been fixed.

OPSX-4407, OPSAPS-67046: Docker Server role fails to come up and returns a connection error during ECS upgrade

Fixed the ECS issue where when upgrading, a Docker server role would sometimes fail to come up and return the following error:

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

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

Fixed Issues in Management Console 1.5.1

OPSAPS-65551: Increase default fd limit for ECS

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

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

Previously, stopping and starting an ECS Role only stopped and started the role respectively, which caused issues in Kubernetes and Longhorn volume health to turn bad. Now, when a user stops an ECS Role (Server or Agent), we perform a "cordon" followed by a "drain" on the node and then

stop the ECS Role on the node. When starting an ECS Role, we first start the ECS Role, then we do an "uncordon" on the node to allow Kubernetes to reuse the node for its workload. A restart operation on ECS Service will perform a Rolling Restart, which does the same steps involved in stopping and starting roles, but one node at a time.

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

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

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

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

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

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

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

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

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

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

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

ENGESC-20112: Unable to progress ECS Upgrade

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

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

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

OPSAPS-66433: Support rolling upgrade for Docker service

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

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

The Refresh ECS command will restart the pods that are in pending state for 10 minutes. This value can be configured using the WAIT_TIME_FOR_POD_READINESS parameter.

OPSAPS-65753: Upgrade CP before upgrading K8S

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

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

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

PULSE-77: schemaVersion should be updated to 3

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

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

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

Fixed Issues in Management Console 1.5.0

COMPX-13184 Yunikorn-admission-controller not getting scheduled after restart due to lack of tolerations

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

Fixed Issues in Management Console 1.4.0

OPsx-2697 Not all images are pushed in upgrade

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

Fixed Issues in Management Console 1.3.x

CVE-2021-44228 (Apache Log4j 2 vulnerability) has been addressed in CDW on CDP Private Cloud Management Console version 1.4.0

Log4j 2 has been upgraded to version 2.17.

Fix copy-docker-template

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

NFS provisioner fails on cluster with more than ~10 nodes

Fixed longhorn nfs_provisioner failing to start on clusters with more than 10 nodes.

Longhorn for Kubernetes is upgraded to version 1.2.x

Longhorn has been upgraded from version 1.1.2 to 1.2.x

ECS High Availability fails during installation

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