

Troubleshooting CDP Private Cloud Data Services

Date published: 2023-12-16

Date modified: 2024-10-18



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

Refreshing the YuniKorn configuration.....	4
Deleting the CNI directory.....	4
Testing Longhorn health post ECS upgrade.....	5
Restarting the Docker Servers.....	6

Refreshing 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.

Procedure

1. Run the following commands to scale down the YuniKorn pods:

```
kubectl scale deployment yunikorn-admission-controller --replicas=0 -n yunikorn
kubectl scale deployment yunikorn-scheduler --replicas=0 -n yunikorn
```

The yunikorn-scheduler and yunikorn-admission-controller pods are managed by the yunikorn-scheduler and yunikorn-admission-controller deployments in the yunikorn namespace, so you can scale down these deployments to 0.

2. Run the following command to delete the yunikorn-configs ConfigMap:

```
kubectl delete cm yunikorn-configs -n yunikorn
```

3. Run the following commands to restart the resource-pool-manager pod:

```
kubectl scale deployment cdp-release-resource-pool-manager --replicas=0 -n <cdp-namespace>
kubectl scale deployment cdp-release-resource-pool-manager --replicas=1 -n <cdp-namespace>
```

The resource-pool-manager pod is managed by the cdp-release-resource-pool-manager deployment in your CDP control plane namespace, so you can scale that deployment down to 0 and then scale it back up to 1.

4. Run the following commands to scale up the YuniKorn pods:

```
kubectl scale deployment yunikorn-scheduler --replicas=1 -n yunikorn
kubectl scale deployment yunikorn-admission-controller --replicas=1 -n yunikorn
```

The yunikorn-scheduler and yunikorn-admission-controller pods are managed by the yunikorn-scheduler and yunikorn-admission-controller deployments in the yunikorn namespace, so you can scale up these deployments to 1.

Results

The preceding steps will refresh the YuniKorn configuration for the applicable control plane.

After the YuniKorn restart, Pending pods will be picked up and recovered automatically, but pods left in the ApplicationRejected state may need to be redeployed. If the pod is managed by a deployment, you can simply delete the pod. If the pod is unmanaged, you must delete and redeploy the pod.

Deleting the CNI directory

First run command fails at setup storage step with error "Timed out waiting for local path storage to come up". Pod is stuck in pending state on the host for a long time.

To fix this issue:

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:

```
kubect1 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-11kc9                                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>
kubect1 delete pod rke2-canal-11kc9 -n kube-system
```

Testing Longhorn health post ECS upgrade

Post ECS upgrade Longhorn health test fails and the helm-install-longhorn pod gets in crashloop state.

To fix this issue, run the following command:

```
#Get the history of longhorn helm chart so that we can identify the chart fo
r which installation is failing. #
    helm history longhorn -n longhorn-system
      REVISION    UPDATED                               STATUS          CHART
      APP VERSION  DESCRIPTION
orn-1.4.2        v1.4.2          Wed Feb 28 05:32:47 2024    deployed        longh
      2          Wed Feb 28 09:28:39 2024    uninstalling    longh
orn-1.5.4        v1.5.4          Deletion in progress (or silently failed)

    #The actual chart is saved as kubernetes secret. List the lon
ghorn helm chart saved as secrets.#
    kubect1 get secrets -n longhorn-system
      NAME                                TYPE          DATA
      AGE
15h    basic-auth                             Opaque        1
0h     chart-values-longhorn                  Opaque        0          1
15h    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          2
1m

    #We want to delete the latest chart i.e. sh.helm.release.v1.1
onghorn.v2. Save the back up of the secret as yaml before deleting. #
    kubect1 get secrets sh.helm.release.v1.longhorn.v2 -n longho
rn-system -o yaml > sh.helm.release.v1.longhorn.v2.yaml
    #Save the back up of the default values passed along with the
helm chart while installing.#
```

```

helm get values --revision=2 longhorn -n longhorn-system >
defaultSettings.yaml

#Find all jobs in longhorn-system and delete those. These jobs
will be re-triggered as part of the manual patch.#
kubectl get jobs -n longhorn-system
NAME                    COMPLETIONS   DURATION   AGE
helm-install-longhorn   0/1            9h          9h
longhorn-post-upgrade   1/1            11m         10h
longhorn-uninstall      0/1            10h         10h

#Delete the latest longhorn chart#
kubectl delete job helm-install-longhorn longhorn-uninstall 1
longhorn-post-upgrade -n longhorn-system
kubectl delete secret sh.helm.release.v1.longhorn.v2 -n longho
rn-system

#Apply the longhorn chart from the parcel directory.#
kubectl patch HelmChart longhorn -n longhorn-system --type=m
erge --patch-file /opt/cloudera/parcels/ECS/longhorn/longhorn.yaml

```

Restarting the Docker Servers

Docker servers fail 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
```

To fix this issue, remove the `/var/run/docker.sock` directory on the Docker server role host and then restart the Docker server role.