Red Hat OpenShift Container Storage 4.2

Troubleshooting OpenShift Container Storage

How to troubleshoot errors and issues in OpenShift Container Storage

Red Hat OpenShift Container Storage 4.2 Troubleshooting OpenShift Container Storage

How to troubleshoot errors and issues in OpenShift Container Storage

Legal Notice

Abstract

Read this document for instructions on troubleshooting Red Hat OpenShift Container Storage.

Table of Contents

CHAPTER 1. OVERVIEW

Troubleshooting OpenShift Container Storage is written to help administrators understand how to troubleshoot and fix their Red Hat OpenShift Container Storage cluster.

Most troubleshooting tasks focus on either a fix or a workaround. This document is divided into chapters based on the errors that an administrator may encounter:

- Chapter 2, *Downloading log files and diagnostic information using must-gather* shows you how to use the must-gather utility in OpenShift Container Storage.

- Chapter 3, *Commonly required logs for troubleshooting* shows you how to obrtain commonly required log files for OpenShift Container Storage.

- Chapter 4, *Replacing storage nodes for OpenShift Container Storage* shows you how to replace an operational or failed node on AWS UPI, AWS IPI, and VMware UPI for OpenShift Container Storage.

- Chapter 6, *Troubleshooting alerts and errors in OpenShift Container Storage* shows you how to identify the encountered error and perform required actions.

CHAPTER 2. DOWNLOADING LOG FILES AND DIAGNOSTIC INFORMATION USING MUST-GATHER

If Red Hat OpenShift Container Storage 4.2 is unable to automatically resolve a problem, use the must-gather tool to collect log files and diagnostic information so that you or Red Hat support can review the problem and determine a solution.

Procedure

1. Run the **must-gather** command from the client connected to the Openshift Container Storage cluster:

   ```
   $ oc adm must-gather --image=registry.redhat.io/ocs4/ocs-must-gather-rhel8 --dest-dir=
   <directory-name>
   ```

This collects the following information in the specified directory:

- Collects all OpenShift Container Storage cluster related Custom Resources (CRs) with their namespaces.

- Collects pod logs of all the OpenShift Container Storage related pods.

- Collects output of some standard Ceph commands like Status, Cluster health, and others.

CHAPTER 3. COMMONLY REQUIRED LOGS FOR TROUBLESHOOTING

Some of the commonly used logs for troubleshooting OpenShift Container Storage are listed, along with the commands to generate them.

- Generating logs for a specific pod:

  ```
  $ oc logs <pod-name> -n <namespace>
  ```

- Generating logs for Ceph or OpenShift Container Storage cluster:

  ```
  $ oc logs rook-ceph-operator-<ID> -n openshift-storage
  ```

- Generating logs for plugin pods like cephfs or rbd to detect any problem in the PVC mount of the app-pod:

  ```
  $ oc logs csi-cephfsplugin-<ID> -n openshift-storage
  ```

  ```
  $ oc logs csi-rbdplugin-<ID> -n openshift-storage
  ```

- Generating provisioner cephfs or rbd logs if PVC is not in **BOUND** state

  ```
  $ oc logs csi-cephfsplugin-provisioner-<ID> -n openshift-storage
  ```

  ```
  $ oc logs csi-rbdplugin-provisioner-<ID> -n openshift-storage
  ```

- Generating OpenShift Container Storage logs using cluster-info command:

  ```
  $ oc cluster-info dump -n openshift-storage --output-directory=<directory-name>
  ```

Additional resources

- Using must-gather

CHAPTER 4. REPLACING STORAGE NODES FOR OPENSHIFT CONTAINER STORAGE

For OpenShift Container Storage 4.2, node replacement can be performed proactively for an operational node and reactively for a failed node for the following deployments:

- For Amazon Web Services (AWS)
  - User-provisioned infrastructure
  - Installer-provisioned infrastructure

- For VMware
  - User-provisioned infrastructure

4.1. OPENSHIFT CONTAINER STORAGE DEPLOYED ON AWS

4.1.1. Replacing an operational AWS node on user-provisioned infrastructure

Perform this procedure to replace an operational node on AWS user-provisioned infrastructure.

Procedure

1. Identify the node that needs to be replaced.

2. Mark the node as unschedulable using the following command:

   ```
   $ oc adm cordon <node_name>
   ```

3. Drain the node using the following command:

   ```
   $ oc adm drain <node_name> --force --delete-local-data --ignore-daemonsets
   ```


   IMPORTANT

   This activity may take at least 5-10 minutes or more. Ceph errors generated during this period are temporary and are automatically resolved when the new node is labeled and functional.

4. Delete the node using the following command:

   ```
   $ oc delete nodes <node_name>
   ```

5. Create a new AWS machine instance with the required infrastructure. See Infrastructure requirements.

6. Create a new OpenShift Container Platform node using the new AWS machine instance.

7. Check for certificate signing requests (CSRs) related to OpenShift Container Platform that are in **Pending** state:

```
$ oc get csr
```

8. Approve all required OpenShift Container Platform CSRs for the new node:

```
$ oc adm certificate approve <Certificate_Name>
```

9. Click **Compute → Nodes**, confirm if the new node is in **Ready** state.

10. Apply the OpenShift Container Storage label to the new node using any one of the following:

    From User interface

    a. For the new node, click **Action Menu (⋮) → Edit Labels**.

    b. Add **cluster.ocs.openshift.io/openshift-storage** and click **Save**.

    From Command line interface

    - Execute the following command to apply the OpenShift Container Storage label to the new node:

    ```
    $ oc label node <new_node_name> cluster.ocs.openshift.io/openshift-storage=""
    ```

11. Restart the **mgr** pod to update the OpenShift Container Storage with the new hostname.

```
$ oc delete pod rook-ceph-mgr-xxxx
```

Verification steps

1. Execute the following command and verify that the new node is present in the output:

```
$ oc get nodes --show-labels | grep cluster.ocs.openshift.io/openshift-storage= |cut -d' ' -f1
```

2. Click **Workloads → Pods**, confirm that at least the following pods on the new node are in **Running** state:

   - **csi-cephfsplugin-***

   - **csi-rbdplugin-***

3. Verify that all other required OpenShift Container Storage pods are in **Running** state.

4. If verification steps fail, kindly contact Red Hat Support .

4.1.2. Replacing an operational AWS node on installer-provisioned infrastructure

Perform this procedure to replace an operational node on AWS installer-provisioned infrastructure (IPI).
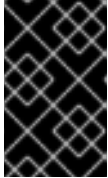
Procedure

1. Log in to OpenShift Web Console and click **Compute → Nodes**.

2. Identify the node that needs to be replaced. Take a note of its **Machine Name**.

3. Mark the node as unschedulable using the following command:

   ```
   $ oc adm cordon <node_name>
   ```

4. Drain the node using the following command:

   ```
   $ oc adm drain <node_name> --force --delete-local-data --ignore-daemonsets
   ```

   > **IMPORTANT**
   >
   > This activity may take at least 5-10 minutes or more. Ceph errors generated during this period are temporary and are automatically resolved when the new node is labeled and functional.

5. Click **Compute → Machines**. Search for the required machine.

6. Besides the required machine, click the **Action menu (⋮) → Delete Machine**.

7. Click **Delete** to confirm the machine deletion. A new machine is automatically created.

8. Wait for new machine to start and transition into **Running** state.
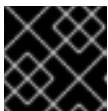
   > **IMPORTANT**
   >
   > This activity may take at least 5-10 minutes or more.

9. Click **Compute → Nodes**, confirm if the new node is in **Ready** state.

10. Apply the OpenShift Container Storage label to the new node using any one of the following:

    From User interface

    a. For the new node, click **Action Menu (⋮) → Edit Labels**

    b. Add **cluster.ocs.openshift.io/openshift-storage** and click **Save**.

    From Command line interface

    - Execute the following command to apply the OpenShift Container Storage label to the new node:

      ```
      $ oc label node <new_node_name> cluster.ocs.openshift.io/openshift-storage=""
      ```

11. Restart the **mgr** pod to update the OpenShift Container Storage with the new hostname.

    ```
    $ oc delete pod rook-ceph-mgr-xxxx
    ```

Verification steps

1. Execute the following command and verify that the new node is present in the output:

   ```
   $ oc get nodes --show-labels | grep cluster.ocs.openshift.io/openshift-storage= |cut -d' ' -f1
   ```

2. Click **Workloads → Pods**, confirm that at least the following pods on the new node are in **Running** state:

   - **csi-cephfsplugin-***

   - **csi-rbdplugin-***

3. Verify that all other required OpenShift Container Storage pods are in **Running** state.

4. If verification steps fail, kindly contact Red Hat Support .

4.1.3. Replacing a failed AWS node on user-provisioned infrastructure

Perform this procedure to replace a failed node which is not operational on AWS user-provisioned infrastructure (UPI) for OpenShift Container Storage 4.2.

Procedure

1. Identify the AWS machine instance of the node that needs to be replaced.

2. Log in to AWS and terminate the identified AWS machine instance.

3. Create a new AWS machine instance with the required infrastructure. See Infrastructure requirements.

4. Create a new OpenShift Container Platform node using the new AWS machine instance.

5. Check for certificate signing requests (CSRs) related to OpenShift Container Platform that are in **Pending** state:

   ```
   $ oc get csr
   ```

6. Approve all required OpenShift Container Platform CSRs for the new node:

   ```
   $ oc adm certificate approve <Certificate_Name>
   ```

7. Click **Compute → Nodes**, confirm if the new node is in **Ready** state.

8. Apply the OpenShift Container Storage label to the new node using any one of the following:

   From User interface

   a. For the new node, click **Action Menu (⋮) → Edit Labels**

   b. Add **cluster.ocs.openshift.io/openshift-storage** and click **Save**.

   From Command line interface

   - Execute the following command to apply the OpenShift Container Storage label to the new node:

     ```
     $ oc label node <new_node_name> cluster.ocs.openshift.io/openshift-storage=""
     ```

9. Restart the **mgr** pod to update the OpenShift Container Storage with the new hostname.

> $ oc delete pod rook-ceph-mgr-xxxx

Verification steps

1. Execute the following command and verify that the new node is present in the output:

   > $ oc get nodes --show-labels | grep cluster.ocs.openshift.io/openshift-storage= |cut -d' ' -f1

2. Click **Workloads → Pods**, confirm that at least the following pods on the new node are in **Running** state:

   - **csi-cephfsplugin-***

   - **csi-rbdplugin-***

3. Verify that all other required OpenShift Container Storage pods are in **Running** state.

4. If verification steps fail, kindly contact Red Hat Support .

4.1.4. Replacing a failed AWS node on installer-provisioned infrastructure

Perform this procedure to replace a failed node which is not operational on AWS installer-provisioned infrastructure (IPI) for OpenShift Container Storage 4.2.

Procedure

1. Log in to OpenShift Web Console and click **Compute → Nodes**.

2. Identify the faulty node and click on its **Machine Name**.

3. Click **Actions → Edit Annotations**, and click **Add More**.

4. Add **machine.openshift.io/exclude-node-draining** and click **Save**.

5. Click **Actions → Delete Machine**, and click **Delete**.

6. A new machine is automatically created, wait for new machine to start.
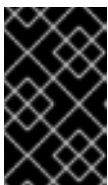
   IMPORTANT

   This activity may take at least 5-10 minutes or more. Ceph errors generated during this period are temporary and are automatically resolved when the new node is labeled and functional.

7. Click **Compute → Nodes**, confirm if the new node is in **Ready** state.

8. Apply the OpenShift Container Storage label to the new node using any one of the following:

   From User interface

   a. For the new node, click **Action Menu (⋮) → Edit Labels**

   b. Add **cluster.ocs.openshift.io/openshift-storage** and click **Save**.

   From Command line interface

- Execute the following command to apply the OpenShift Container Storage label to the new node:

  ```
  $ oc label node <new_node_name> cluster.ocs.openshift.io/openshift-storage=""
  ```

9. [Optional]: If the failed AWS instance is not removed automatically, terminate the instance from AWS console.

10. Restart the **mgr** pod to update the OpenShift Container Storage with the new hostname.

    ```
    $ oc delete pod rook-ceph-mgr-xxxx
    ```

Verification steps

1. Execute the following command and verify that the new node is present in the output:

   ```
   $ oc get nodes --show-labels | grep cluster.ocs.openshift.io/openshift-storage= |cut -d' ' -f1
   ```

2. Click **Workloads → Pods**, confirm that at least the following pods on the new node are in **Running** state:

   - **csi-cephfsplugin-***

   - **csi-rbdplugin-***

3. Verify that all other required OpenShift Container Storage pods are in **Running** state.

4. If verification steps fail, kindly contact Red Hat Support .

4.2. OPENSHIFT CONTAINER STORAGE DEPLOYED ON VMWARE

4.2.1. Replacing an operational VMware node on user-provisioned infrastructure

Perform this procedure to replace an operational node on VMware user-provisioned infrastructure (UPI).

Procedure

1. Identify the node and its VM that needs to be replaced.

2. Mark the node as unschedulable using the following command:

   ```
   $ oc adm cordon <node_name>
   ```

3. Drain the node using the following command:

   ```
   $ oc adm drain <node_name> --force --delete-local-data --ignore-daemonsets
   ```
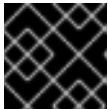
> **IMPORTANT**
>
> This activity may take at least 5-10 minutes or more. Ceph errors generated during this period are temporary and are automatically resolved when the new node is labeled and functional.

4. Delete the node using the following command:

   ```
   $ oc delete nodes <node_name>
   ```

5. Log in to vSphere and terminate the identified VM.

> **IMPORTANT**
>
> VM should be deleted only from the inventory and not from the disk.

6. Create a new VM on vSphere with the required infrastructure. See Infrastructure requirements.

7. Create a new OpenShift Container Platform worker node using the new VM.

8. Check for certificate signing requests (CSRs) related to OpenShift Container Platform that are in **Pending** state:

   ```
   $ oc get csr
   ```

9. Approve all required OpenShift Container Platform CSRs for the new node:

   ```
   $ oc adm certificate approve <Certificate_Name>
   ```

10. Click **Compute → Nodes**, confirm if the new node is in **Ready** state.

11. Apply the OpenShift Container Storage label to the new node using any one of the following:

    From User interface

    a. For the new node, click **Action Menu (⋮) → Edit Labels**

    b. Add **cluster.ocs.openshift.io/openshift-storage** and click **Save**.

    From Command line interface

    - Execute the following command to apply the OpenShift Container Storage label to the new node:

      ```
      $ oc label node <new_node_name> cluster.ocs.openshift.io/openshift-storage=""
      ```

12. Restart the **mgr** pod to update the OpenShift Container Storage with the new hostname.

    ```
    $ oc delete pod rook-ceph-mgr-xxxx
    ```

Verification steps

1. Execute the following command and verify that the new node is present in the output:

   ```
   $ oc get nodes --show-labels | grep cluster.ocs.openshift.io/openshift-storage= |cut -d' ' -f1
   ```

2. Click **Workloads → Pods**, confirm that at least the following pods on the new node are in **Running** state:

   - **csi-cephfsplugin-***

   - **csi-rbdplugin-***

3. Verify that all other required OpenShift Container Storage pods are in **Running** state.

4. If verification steps fail, kindly contact Red Hat Support .

4.2.2. Replacing a failed VMware node on user-provisioned infrastructure

Perform this procedure to replace a failed node on VMware user-provisioned infrastructure (UPI).

Procedure

1. Identify the node and its VM that needs to be replaced.

2. Delete the node using the following command:

   ```
   $ oc delete nodes <node_name>
   ```

3. Log in to vSphere and terminate the identified VM.
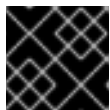
   > **IMPORTANT**
   >
   > VM should be deleted only from the inventory and not from the disk.

4. Create a new VM on vSphere with the required infrastructure. See Infrastructure requirements.

5. Create a new OpenShift Container Platform worker node using the new VM.

6. Check for certificate signing requests (CSRs) related to OpenShift Container Platform that are in **Pending** state:

   ```
   $ oc get csr
   ```

7. Approve all required OpenShift Container Platform CSRs for the new node:

   ```
   $ oc adm certificate approve <Certificate_Name>
   ```

8. Click **Compute → Nodes**, confirm if the new node is in **Ready** state.

9. Apply the OpenShift Container Storage label to the new node using any one of the following:

   From User interface

   a. For the new node, click **Action Menu (⋮) → Edit Labels**

      b. Add **cluster.ocs.openshift.io/openshift-storage** and click **Save**.

From Command line interface

- Execute the following command to apply the OpenShift Container Storage label to the new node:

  ```
  $ oc label node <new_node_name> cluster.ocs.openshift.io/openshift-storage=""
  ```

10. Restart the **mgr** pod to update the OpenShift Container Storage with the new hostname.

    ```
    $ oc delete pod rook-ceph-mgr-xxxx
    ```

Verification steps

1. Execute the following command and verify that the new node is present in the output:

   ```
   $ oc get nodes --show-labels | grep cluster.ocs.openshift.io/openshift-storage= |cut -d' ' -f1
   ```

2. Click **Workloads → Pods**, confirm that at least the following pods on the new node are in **Running** state:

   - **csi-cephfsplugin-***

   - **csi-rbdplugin-***

3. Verify that all other required OpenShift Container Storage pods are in **Running** state.

4. If verification steps fail, kindly contact Red Hat Support .

CHAPTER 5. RECOVERING PODS

When a first node (say **NODE1**) goes to NotReady state because of some issue, the hosted pods that are using PVC with ReadWriteOnce (RWO) access mode try to move to the second node (say **NODE2**) but get stuck due to multi-attach error. In such a case, you can recover MON, OSD, and application pods by using the following steps.

Procedure

1. Power off **NODE1** (from AWS or vSphere side) and ensure that **NODE1** is completely down.

2. Force delete the pods on **NODE1** by using the following command:

   ```
   $ oc delete pod <pod-name> --grace-period=0 --force
   ```

CHAPTER 6. TROUBLESHOOTING ALERTS AND ERRORS IN OPENSHIFT CONTAINER STORAGE

6.1. RESOLVING ALERTS AND ERRORS

Red Hat OpenShift Container Storage can detect and automatically resolve a number of common failure scenarios. However, some problems require administrator intervention.

To know the errors currently firing, check one of the following locations:

- **Monitoring → Alerting → Firing** option

- **Home → Dashboards → Overview** tab

- **Home → Dashboards → Persistent Storage** tab

- **Home → Dashboards → Object Service** tab

Copy the error displayed and search it in the following section to know its severity and resolution:

Name: **CephMonVersionMismatch**

Message: **There are multiple versions of storage services running.**

Description: **There are {{ $value }} different versions of Ceph Mon components running.**

Severity: Warning

Resolution: Fix

Procedure: Inspect the user interface and log, and verify if an update is in progress.

- If an update in progress, this alert is temporary.

- If an update is not in progress, restart the upgrade process.

Name: **CephOSDVersionMismatch**

Message: **There are multiple versions of storage services running.**

Description: **There are {{ $value }} different versions of Ceph OSD components running.**

Severity: Warning

Resolution: Fix

Procedure: Inspect the user interface and log, and verify if an update is in progress.

- If an update in progress, this alert is temporary.

- If an update is not in progress, restart the upgrade process.

Name: **CephClusterCriticallyFull**

Message: **Storage cluster is critically full and needs immediate expansion**

Description: **Storage cluster utilization has crossed 85%.**

Severity: Crtical

Resolution: Fix

Procedure: Remove unnecessary data or expand the cluster.

Name: **CephClusterNearFull**

Fixed: **Storage cluster is nearing full. Expansion is required.**

Description: **Storage cluster utilization has crossed 75%.**

Severity: Warning

Resolution: Fix

Procedure: Remove unnecessary data or expand the cluster.

Name: **NooBaaBucketErrorState**

Message: **A NooBaa Bucket Is In Error State**

Description: **A NooBaa bucket {{ $labels.bucket_name }} is in error state for more than 6m**

Severity: Warning

Resolution: Workaround

Procedure: Resolving NooBaa Bucket Error State

Name: **NooBaaBucketExceedingQuotaState**

Message: **A NooBaa Bucket Is In Exceeding Quota State**

Description: **A NooBaa bucket {{ $labels.bucket_name }} is exceeding its quota - {{ printf "%0.0f" $value }}% used message: A NooBaa Bucket Is In Exceeding Quota State**

Severity: Warning

Resolution: Fix

Procedure: Resolving NooBaa Bucket Exceeding Quota State

Name: **NooBaaBucketLowCapacityState**

Message: **A NooBaa Bucket Is In Low Capacity State**

Description: **A NooBaa bucket {{ $labels.bucket_name }} is using {{ printf "%0.0f" $value }}% of its capacity**

Severity: Warning

Resolution: Fix

Procedure: Resolving NooBaa Bucket Capacity or Quota State

Name: **NooBaaBucketNoCapacityState**

Message: **A NooBaa Bucket Is In No Capacity State**

Description: **A NooBaa bucket {{ $labels.bucket_name }} is using all of its capacity**

Severity: Warning

Resolution: Fix

Procedure: Resolving NooBaa Bucket Capacity or Quota State

Name: **NooBaaBucketReachingQuotaState**

Message: **A NooBaa Bucket Is In Reaching Quota State**

Description: **A NooBaa bucket {{ $labels.bucket_name }} is using {{ printf "%0.0f" $value }}% of its quota**

Severity: Warning

Resolution: Fix

Procedure: Resolving NooBaa Bucket Capacity or Quota State

Name: **NooBaaResourceErrorState**

Message: **A NooBaa Resource Is In Error State**

Description: **A NooBaa resource {{ $labels.resource_name }} is in error state for more than 6m**

Severity: Warning

Resolution: Workaround

Procedure: Resolving NooBaa Bucket Error State

Name: **NooBaaSystemCapacityWarning100**

Message: **A NooBaa System Approached Its Capacity**

Description: **A NooBaa system approached its capacity, usage is at 100%**

Severity: Warning

Resolution: Fix

Procedure: Resolving NooBaa Bucket Capacity or Quota State

Name: **NooBaaSystemCapacityWarning85**

Message: **A NooBaa System Is Approaching Its Capacity**

Description: **A NooBaa system is approaching its capacity, usage is more than 85%**

Severity: Warning

Resolution: Fix

Procedure: Resolving NooBaa Bucket Capacity or Quota State

Name: **NooBaaSystemCapacityWarning95**

Message: **A NooBaa System Is Approaching Its Capacity**

Description: **A NooBaa system is approaching its capacity, usage is more than 95%**

Severity: Warning

Resolution: Fix

Procedure: Resolving NooBaa Bucket Capacity or Quota State

Name: **CephMdsMissingReplicas**

Message: **Insufficient replicas for storage metadata service.**

Description: `Minimum required replicas for storage metadata service not available.

Might affect the working of storage cluster.`

Severity: Warning

Resolution: Contact Red Hat support

Procedure:

1. Check for alerts and operator status.

2. If the issue cannot be identified, contact Red Hat support.

Name: **CephMgrIsAbsent**

Message: **Storage metrics collector service not available anymore.**

Description: **Ceph Manager has disappeared from Prometheus target discovery.**

Severity: Critical

Resolution: Contact Red Hat support

Procedure:

1. Inspect the user interface and log, and verify if an update is in progress.

   - If an update in progress, this alert is temporary.

   - If an update is not in progress, restart the upgrade process.

2. Once the upgrade is complete, check for alerts and operator status.

3. If the issue persistents or cannot be identified, contact Red Hat support.

Name: **CephNodeDown**

Message: **Storage node {{ $labels.node }} went down**

Description: **Storage node {{ $labels.node }} went down. Please check the node immediately.**

Severity: Critical

Resolution: Contact Red Hat support

Procedure:

1. Check which node stopped functioning and its cause.

2. Take appropriate actions to recover the node. If node cannot be recovered:

   - See Replacing storage nodes for OpenShift Container Storage

   - Contact Red Hat support

Name: **CephClusterErrorState**

Message: **Storage cluster is in error state**

Description: **Storage cluster is in error state for more than 10m.**

Severity: Critical

Resolution: Contact Red Hat support

Procedure:

1. Check for alerts and operator status.

2. If the issue cannot be identified, download log files and diagnostic information using must-gather.

3. Open a Support Ticket with Red Hat Support with an attachment of the output of must-gather.

Name: **CephClusterWarningState**

Message: **Storage cluster is in degraded state**

Description: **Storage cluster is in warning state for more than 10m.**

Severity: Warning

Resolution: Contact Red Hat support

Procedure:

1. Check for alerts and operator status.

2. If the issue cannot be identified, download log files and diagnostic information using must-gather.

3. Open a Support Ticket with Red Hat Support with an attachment of the output of must-gather.

Name: **CephDataRecoveryTakingTooLong**

Message: **Data recovery is slow**

Description: **Data recovery has been active for too long.**

Severity: Warning

Resolution: Contact Red Hat support

Name: **CephOSDDiskNotResponding**

Message: **Disk not responding**

Description: **Disk device {{ $labels.device }} not responding, on host {{ $labels.host }}.**

Severity: Critical

Resolution: Contact Red Hat support

Name: **CephOSDDiskUnavailable**

Message: **Disk not accessible**

Description: **Disk device {{ $labels.device }} not accessible on host {{ $labels.host }}.**

Severity: Critical

Resolution: Contact Red Hat support

Name: **CephPGRepairTakingTooLong**

Message: **Self heal problems detected**

Description: **Self heal operations taking too long.**

Severity: Warning

Resolution: Contact Red Hat support

Name: **CephMonHighNumberOfLeaderChanges**

Message: **Storage Cluster has seen many leader changes recently.**

Description: **'Ceph Monitor "{{ $labels.job }}": instance {{ $labels.instance }} has seen {{ $value printf "%.2f" }} leader changes per minute recently.'**

Severity: Warning

Resolution: Contact Red Hat support

Name: **CephMonQuorumAtRisk**

Message: **Storage quorum at risk**

Description: **Storage cluster quorum is low.**

Severity: Critical

Resolution: Contact Red Hat support

6.2. RESOLVING NOOBAA BUCKET ERROR STATE

Procedure

1. Log in to OpenShift Web Console and click **Object Service**.

2. In the **Details** card, click the link under **System Name** field.

3. In the left pane, click **Buckets** option and search for the bucket in error state.

4. Click on it's **Bucket Name**. Error encountered in bucket is displayed.

5. Depending on the specific error of the bucket, perform one or both of the following:

a. For space related errors:

   i. In the left pane, click **Resources** option.

   ii. Click on the resource in error state.

   iii. Scale the resource by adding more agents.

b. For resource health errors:

   i. In the left pane, click **Resources** option.

   ii. Click on the resource in error state.

   iii. Connectivity error means the backing service is not available and needs to be restored.

   iv. For access/permissions errors, update the connection's **Access Key** and **Secret Key**.

6.3. RESOLVING NOOBAA BUCKET EXCEEDING QUOTA STATE

To resolve **A NooBaa Bucket Is In Exceeding Quota State** error perform one of the following:

- Cleanup some of the data on the bucket.

- Increase the bucket quota by performing the following steps:

  1. Log in to OpenShift Web Console and click **Object Service**.

  2. In the **Details** card, click the link under **System Name** field.

  3. In the left pane, click **Buckets** option and search for the bucket in error state.

  4. Click on it's **Bucket Name**. Error encountered in bucket is displayed.

  5. Click **Bucket Policies → Edit Quota** and increase the quota.

6.4. RESOLVING NOOBAA BUCKET CAPACITY OR QUOTA STATE

Procedure

  1. Log in to OpenShift Web Console and click **Object Service**.

  2. In the **Details** card, click the link under **System Name** field.

  3. In the left pane, click **Resources** option and search for the PV pool resource.

  4. For the PV pool resource with low capacity status, click on it's **Resource Name**.

  5. Edit the pool configuration and increase the number of agents.