



Red Hat OpenShift Container Storage 4.4

Troubleshooting OpenShift Container Storage

How to troubleshoot errors and issues in OpenShift Container Storage

Red Hat OpenShift Container Storage 4.4 Troubleshooting OpenShift Container Storage

How to troubleshoot errors and issues in OpenShift Container Storage

Legal Notice

Copyright © 2020 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, the Red Hat logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux[®] is the registered trademark of Linus Torvalds in the United States and other countries.

Java[®] is a registered trademark of Oracle and/or its affiliates.

XFS[®] is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL[®] is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js[®] is an official trademark of Joyent. Red Hat is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack[®] Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

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

Table of Contents

CHAPTER 1. OVERVIEW	3
CHAPTER 2. DOWNLOADING LOG FILES AND DIAGNOSTIC INFORMATION USING MUST-GATHER	4
CHAPTER 3. COMMONLY REQUIRED LOGS FOR TROUBLESHOOTING	5
CHAPTER 4. OVERRIDING THE CLUSTER-WIDE DEFAULT NODE SELECTOR FOR OPENSIFT CONTAINER STORAGE POST DEPLOYMENT	6
CHAPTER 5. TROUBLESHOOTING ALERTS AND ERRORS IN OPENSIFT CONTAINER STORAGE	7
5.1. RESOLVING ALERTS AND ERRORS	7
5.2. RESOLVING NOOBAA BUCKET ERROR STATE	13
5.3. RESOLVING NOOBAA BUCKET EXCEEDING QUOTA STATE	14
5.4. RESOLVING NOOBAA BUCKET CAPACITY OR QUOTA STATE	14
5.5. RECOVERING PODS	14
5.6. RECOVERING FROM EBS VOLUME DETACH	15
CHAPTER 6. CHECKING FOR LOCAL STORAGE OPERATOR DEPLOYMENTS	16

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 obtain commonly required log files for OpenShift Container Storage.
- [Chapter 5, *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 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

- 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> --node-name=<node-name>
```

where `<node-name>` is any master node in Ready state.



NOTE

--node-name is optional and needs to be used specifically for situations when either one or more worker nodes are not in Ready state.

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. OVERRIDING THE CLUSTER-WIDE DEFAULT NODE SELECTOR FOR OPENSIFT CONTAINER STORAGE POST DEPLOYMENT

When a cluster-wide default node selector is used for OpenShift Container Storage, the pods generated by CSI daemonsets are able to start only on the nodes that match the selector. To be able to use OpenShift Container Storage from nodes which do not match the selector, override the **cluster-wide default node selector** by performing the following steps in the command line interface :

Procedure

1. Specify a blank node selector for the openshift-storage namespace.

```
$ oc annotate namespace openshift-storage openshift.io/node-selector=
```

2. Delete the original pods generated by the DaemonSets.

```
oc delete pod -l app=csi-cephfsplugin -n openshift-storage
oc delete pod -l app=csi-rbdplugin -n openshift-storage
```

CHAPTER 5. TROUBLESHOOTING ALERTS AND ERRORS IN OPENSIFT CONTAINER STORAGE

5.1. 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** → **Overview** → **Cluster** tab
- **Home** → **Overview** → **Persistent Storage** tab
- **Home** → **Overview** → **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: Critical

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 persists 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](#)

<p>Name: CephOSDDiskUnavailable</p> <p>Message: Disk not accessible</p> <p>Description: Disk device {{ \$labels.device }} not accessible on host {{ \$labels.host }}.</p> <p>Severity: Critical</p> <p>Resolution: Contact Red Hat support</p>
<p>Name: CephPGRepairTakingTooLong</p> <p>Message: Self heal problems detected</p> <p>Description: Self heal operations taking too long.</p> <p>Severity: Warning</p> <p>Resolution: Contact Red Hat support</p>
<p>Name: CephMonHighNumberOfLeaderChanges</p> <p>Message: Storage Cluster has seen many leader changes recently.</p> <p>Description: 'Ceph Monitor "{{ \$labels.job }}" instance {{ \$labels.instance }} has seen {{ \$value printf "%.2f" }} leader changes per minute recently.'</p> <p>Severity: Warning</p> <p>Resolution: Contact Red Hat support</p>
<p>Name: CephMonQuorumAtRisk</p> <p>Message: Storage quorum at risk</p> <p>Description: Storage cluster quorum is low.</p> <p>Severity: Critical</p> <p>Resolution: Contact Red Hat support</p>

5.2. 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 its **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**.

5.3. 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.

5.4. 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.

5.5. 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
```

5.6. RECOVERING FROM EBS VOLUME DETACH

When an OSD or MON elastic block storage (EBS) volume where the OSD disk resides is detached from the worker Amazon EC2 instance, the volume gets reattached automatically within one or two minutes. However, the OSD pod gets into a **CrashLoopBackOff** state. To recover and bring back the pod to **Running** state, you must restart the EC2 instance.

CHAPTER 6. CHECKING FOR LOCAL STORAGE OPERATOR DEPLOYMENTS

OpenShift Container Storage clusters with Local Storage Operator are deployed using local storage devices. To find out if your existing cluster with OpenShift Container Storage was deployed using local storage devices, use the following procedure:

Prerequisites

- OpenShift Container Storage is installed and running in the **openshift-storage** namespace.

Procedure

By checking the storage class associated with your OpenShift Container Storage cluster's persistent volume claims (PVCs), you can tell if your cluster was deployed using local storage devices.

- Check the storage class associated with OpenShift Container Storage cluster's PVCs with the following command:

```
$ oc get pvc -n openshift-storage
```

- Check the output. For clusters with Local Storage Operator, the PVCs associated with **ocs-deviceset** use the storage class **localblock**. The output looks similar to the following:

NAME	STATUS	VOLUME	CAPACITY	ACCESS
db-noobaa-db-0	Bound	pvc-d96c747b-2ab5-47e2-b07e-1079623748d8	50Gi	
RWO	ocs-storagecluster-ceph-rbd	114s		
ocs-deviceset-0-0-lzfrd	Bound	local-pv-7e70c77c	1769Gi	RWO
localblock	2m10s			
ocs-deviceset-1-0-7rggl	Bound	local-pv-b19b3d48	1769Gi	RWO
localblock	2m10s			
ocs-deviceset-2-0-znhk8	Bound	local-pv-e9f22cdc	1769Gi	RWO
localblock	2m10s			

Additional Resources

- [Deploying OpenShift Container Storage using local storage devices](#)