Red Hat OpenShift Container Storage 4.6

4.6 Release Notes

Release notes for feature and enhancements, known issues, and other important release information

Last Updated: 2021-10-27

Red Hat OpenShift Container Storage 4.6 4.6 Release Notes

Release notes for feature and enhancements, known issues, and other important release information

Legal Notice

Abstract

The release notes for Red Hat OpenShift Container Storage 4.6 summarize all new features and enhancements, notable technical changes, and any known bugs upon general availability.

Table of Contents

CHAPTER 1. INTRODUCTION

Red Hat OpenShift Container Storage is software-defined storage that is optimized for container environments. It runs as an operator on OpenShift Container Platform to provide highly integrated and simplified persistent storage management for containers.

Red Hat OpenShift Container Storage is integrated into the latest Red Hat OpenShift Container Platform to address platform services, application portability, and persistence challenges. It provides a highly scalable backend for the next generation of cloud-native applications, built on a new technology stack that includes Red Hat Ceph Storage, the Rook.io Operator, and NooBaa's Multicloud Object Gateway technology.

Red Hat OpenShift Container Storage provides a trusted, enterprise-grade application development environment that simplifies and enhances the user experience across the application lifecycle in a number of ways:

- Provides block storage for databases.
- Shared file storage for continuous integration, messaging, and data aggregation.
- Object storage for cloud-first development, archival, backup, and media storage.
- Scale applications and data exponentially.
- Attach and detach persistent data volumes at an accelerated rate.
- Stretch clusters across multiple data-centers or availability zones.
- Establish a comprehensive application container registry.
- Support the next generation of OpenShift workloads such as Data Analytics, Artificial Intelligence, Machine Learning, Deep Learning, and Internet of Things (IoT).
- Dynamically provision not only application containers, but data service volumes and containers, as well as additional OpenShift Container Platform nodes, Elastic Block Store (EBS) volumes and other infrastructure services.

1.1. ABOUT THIS RELEASE

Red Hat OpenShift Container Storage 4.6 (RHBA-2020:5606 and RHBA-2020:5605) is now available. New enhancements, features, and known issues that pertain to OpenShift Container Storage 4.6 are included in this topic.

Red Hat OpenShift Container Storage 4.6 is supported on the latest Red Hat OpenShift Container Platform version. For more information, see Red Hat OpenShift Container Storage Supportability and Interoperability Guide.

OpenShift Container Storage 4.6 is an Extended Update Support (EUS) release. More information on Red Hat OpenShift Container Storage EUS is available in OpenShift Life Cycle, under layered products, and OpenShift EUS Overview.

With the release of OpenShift Container Storage 4.6, version 4.3 is now end of life. For more information, see the Red Hat OpenShift Container Platform Life Cycle Policy .

CHAPTER 2. NEW FEATURES AND ENHANCEMENTS

This section describes new features and major enhancements introduced in Red Hat OpenShift Container Storage 4.6.

Encrypted storage data

Administrators can now choose to encrypt all data in the OpenShift Container Storage 4.6 cluster as part of the deployment process. See Data encryption options for more information, and follow the OpenShift Container Storage documentation for deploying on your cloud or bare metal environment.

Flexible environment for applications

Users can now create multiple storage pools which map to the custom storage classes.

These multiple pools:

- Enable applications with their own high availability to use persistent volumes with two replicas to make them more efficient
- Save space for persistent volume claims using storage classes with compression enabled

Expanding Persistent Volume Claims

OpenShift Container Storage 4.5 introduced the ability to expand Persistent Volume Claims as a Technology Preview feature providing more flexibility in the management of persistent storage resources.This feature is fully supported for expanding new or pre-existing persistent volume claims as of OpenShift Container Storage 4.6.

For more information, see Expanding Persistent Volume Claims .

Unified view of multicloud object buckets

Associate your object buckets with a namespace bucket to see all objects across the associated buckets, while only writing to your current preferred storage provider. See Configuring namespace buckets for more information.

Automated scaling of Multicloud Object Gateway endpoint pods

Improved performance and serviceability to manage resources for S3 load with provision for automatic scaling when experiencing load increases or decreases.

Automated device discovery with filter

With this release, it is possible to view (or filter) all the available storage devices in the UI while deploying and adding devices for OpenShift Container Storage using the local storage devices in VMware and bare metal infrastructures.

For more information, see Deploying using local storage devices in VMware infrastructure and Deploying using local storage devices in bare metal infrastructure .

User-friendly method to replace a failed local storage device.

With this release, users can replace their failed local storage devices in the VMware and bare metal infrastructures using UI. Users are provided with an inventory list of the disks where they can find the status of the disk and initiate the replacement of a failed device. Also, failed disk replacement can be initiated from the dashboard and alert notifications.

For more information, see Replacing failed storage devices on VMware and bare metal infrastructures using user interface.

Volume snapshot

Developers and administrators can now take a volume snapshot which is the state of the storage volume in a cluster at a particular point in time. These snapshots help to use storage more efficiently by not having to make a full copy each time and can be used as building blocks for developing an application. The volume snapshots can be restored as a new persistent volume claim (PVC).

For more information, see Volume snapshots.

Volume Clone

Clone an OpenShift Container Storage 4.6 persistent volume, as well as use the cloned persistent volume for read/write operations in another application (on RHOCP 4.6) with no performance or application impact.

For more information, see Volume cloning.

RADOS Gateway status and health

The Object Service dashboard now provides status and health information for the RADOS Gateway. For more information, see Viewing metrics in the Object Service dashboard.

Improved PV creation time

The Ceph CSI driver, **ceph-csi**, has moved to use the native Go binding against Red Hat Ceph Storage instead of invoking the **ceph** command line. This leads to improved PV creation time.

MDS and RGW configurations can be provided after external cluster installation

Previously, MDS and RGW configurations were accepted only at the initial time of external cluster creation. With this update, MDS and RGW configurations can be provided at a later point by updating the external secret. The OpenShift Container Storage Operator reconciles the changes and updates the resources accordingly. For more information, see Adding file and object storage to an existing external OpenShift Container Storage cluster.

Telemetry

New usage information regarding RADOS Block Device (RBD) and Ceph File System (CephFS) is now collected via Telemetry. For more information, see Information collected by Telemetry.

Improved uninstallation methods

User-friendly method to uninstall OpenShift Container Storage 4.6 so that there is no data, etc., remaining.

Consistent and comprehensive documentation

Re-organized the managing guide into smaller books so that relevant information can be easily found.

Seamless upgrade process from OpenShift Container Storage version 4.x to version 4.6

OpenShift Container Storage customers now have an easy to use method for upgrading their existing clusters to OpenShift Container Storage 4.6. For more information, see Updating OpenShift Container Storage.

CHAPTER 3. BUG FIXES

This section describes notable bug fixes introduced in Red Hat OpenShift Container Storage 4.6.

MGR pod restarts even if the MONs are down

Previously, when the nodes restarted the MGR pod might get stuck in a pod initialisation state which resulted in the inability to create new persistent volumes (PVs). With this update, the MGR pod restarts even if the MONs are down.

(BZ#1990031)

Old OSD pods no longer stay in Terminating state after disk replacement

Old OSD pod sometimes stayed behind in **Terminating** state after the disk replacement procedure. With this update, if the **rook-ceph-osd pod** is in **Terminating** state, you can now use the **force** option to delete the pod.

MAX HPA value exceeding 1 no longer triggers an alert

In previous versions of Red Hat OpenShift Container Storage, the autoscaling feature for pods was not available. Therefore, the **MAX HPA** value could not be greater than **1**, or an alert was triggered. With this update, this feature is enabled and the alert is no longer triggered.

(BZ#1836299)

ceph-mgr no longer causes errors during requests

Previously, certain **ceph-mgr** modules (**fs**) always connected to the MONs that were passed in as part of the initial **ceph-mgr** pod creation. Therefore, when the MON endpoints were changed, these modules failed to connect to the Red Hat Ceph Storage cluster for various requests, such as provisioning and staging CephFS volumes, causing errors. With this update, **ceph-mgr** has been fixed to keep its MON endpoints updated as they change, and to not rely only on the initial MON addresses that are passed during pod creation, and **ceph-mgr** operations continue and work as expected.

(BZ#1858195)

RGW endpoints no longer run in http mode

Previously, external python scripts always tried to see the given RGW endpoint, reachable or not, and TLS enabled **https** URLs were not supported. This compromised security because the user was forced to run RGW endpoints in **http** mode. The external script has been fixed to include **https** URLs reachability check.

(BZ#1878853)

Broken links removed from UI

Previously, there were broken links in the UI caused by outdated Multicloud Object Gateway documentation. These links have been removed, and all information is now covered in the OpenShift Container Storage documentation.

(BZ#1881398)

OpenShift Container Storage pods are no longer being scheduled on nodes not labeled for OpenShift Container Storage

Previously, some OpenShift Container Storage pods were being scheduled on nodes not labeled for OpenShift Container Storage that belonged on the labeled nodes. This fix adds the proper **NodeAffinity** to those pods, resolving that issue.

(BZ#1883828)

CSI driver and other resources no longer disappear unexpectedly

Previously, there was an invalid owner reference in Rook on the CSI driver. Because of this, OpenShift Container Platform would periodically incorrectly garbage collect the CSI driver and other resources in the **openshift-storage** namespace causing resources to disappear unexpectedly. The invalid owner reference in Rook to the CSI driver has been removed, and the CSI driver and other resources no longer disappear.

(BZ#1884318)

MON PDBs are now reconciled allowing node drains

Previously, the reconciler for MON **PodDisruptionBudget** was static. It would create the PDB only once based on the MON count, but would not update it if the MON count changed. In OpenShift Container Storage versions 4.3 and 4.4, the default MON count was increased to 5 when the cluster had 5 nodes. In OpenShift Container Storage versions 4.5 and later, the MON count was kept to 3, despite a different number of nodes. On upgrading from OpenShift Container Storage versions 4.3 and 4.4 to 4.5 or later, the **ALLOWED DISRUPTIONS** would become 0. This would not allow nodes to drain. With this update, the MON PDB reconciler now creates a new PDB for MONs if the MON count changes. Therefor, **ALLOWED DISRUPTIONS** will always be 1 and allow node drains.

(BZ#1888713)

No issues when setting public access policy to a bucket

Previously, there was a translation issue when setting the public access policy to a bucket and the desired policy would not be set correctly. This translation issue has been fixed, and the desired policy is set correctly so public access can be set.

(BZ#1889683)

CHAPTER 4. TECHNOLOGY PREVIEWS

This section describes technology preview features introduced in Red Hat OpenShift Container Storage 4.6.

Red Hat Virtualization technology preview support

OpenShift Container Storage can now be installed and managed using Red Hat Virtualization. For more information see, the Deploying and managing OpenShift Container Storage using Red Hat Virtualization guide.

Red Hat OpenStack Platform technology preview support

OpenShift Container Storage can now be installed and managed using Red Hat OpenStack Platform. For more information see, the Deploying and managing OpenShift Container Storage using Red Hat OpenStack Platform guide.

IBM Power Systems technology preview support

OpenShift Container Storage can now be installed and managed using IBM Power Systems. For more information see, the Deploying and managing OpenShift Container Storage using IBM Power Systems Guide.

IBM Z and LinuxONE technology preview support

OpenShift Container Storage can now be installed and managed using IBM Z and LinuxONE. For more information see, the Deploying and managing OpenShift Container Storage using IBM Z and LinuxONE .

Minimum deployment technology preview support

OpenShift Container Storage can now be deployed with minimum configuration when the standard deployment resource requirement is not met. For more information, see minimum deployment resource requirements in the Planning Guide.

Compact deployment technology preview support

OpenShift Container Storage can now be installed on a three-node OpenShift compact bare metal cluster, where all the workloads run on three strong master nodes. There are no worker or storage nodes.

For information on how to configure OpenShift Container Platform on a compact bare metal cluster, see Configuring a three-node cluster and Delivering a Three-node Architecture for Edge Deployments.

Expand storage capacity using additional device

Administrators can now scale up storage capacity using storage classes other than the one defined during deployment. First define a new storage class based on an existing storage provider, and select that storage class when you expand your OpenShift Container Storage capacity. See Scaling storage for more information.

CHAPTER 5. KNOWN ISSUES

This section describes known issues in Red Hat OpenShift Container Storage 4.6.

Issue with nooba-db

noobaa-core-0 does not migrate to other nodes when a node goes down. NooBaa will not work when a node is down as migration of **noobaa-core** pod is blocked.

(BZ#1783961)

PodDisruptionBudget alert continuously shown

The **PodDisruptionBudget** alert, which is an OpenShift Container Platform alert, is continuously shown for object storage devices (OSDs). This alert can be ignored. You can choose to silence this alert by following the instructions in the Managing cluster alerts section of the Red Hat Openshift Container Platform documentation.

For more information, refer to this Red Hat Knowledgebase article .

(BZ#1788126)

Restore Snapshot/Clone operations with greater size than parent PVC results in endless loop

Ceph CSI doesn't support restoring a snapshot or creating clones with a size greater than the parent PVC. Therefore, **Restore Snapshot/Clone** operations with a greater size results in an endless loop. To workaround this issue, delete the pending PVC. In order to get a larger PVC, complete one of the following based on the operation you are using:

- If using Snapshots, restore the existing snapshot to create a volume of the same size as the parent PVC, then attach it to a pod and expand the PVC to the required size. For more information, see Volume snapshots.
- If using Clone, clone the parent PVC to create a volume of the same size as the parent PVC, then attach it to a pod and expand the PVC to the required size. For more information, see Volume cloning.

(BZ#1870334)

Prometheus listens only on port 9283

The Prometheus service on **ceph-mgr** in an external cluster is expected to listen on port 9283. Other ports are not supported. Red Hat Ceph Storage administrators must use only port 9283 for the Prometheus exporter.

(BZ#1890971)

Ceph status is HEALTH_WARN after disk replacement

After disk replacement, a warning **1 daemons have recently crashed** is seen even if all OSD pods are up and running. This warning causes a change in Ceph's status. The Ceph status should be **HEALTH_OK** instead of **HEALTH_WARN**. To workaround this issue, **rsh** to the **ceph-tools** pod and silence the warning, the Ceph health will then be back to **HEALTH_OK**.

(BZ#1896810)

You cannot create a PVC from a volume snapshot in the absence of volume snapshotclass

Deleting **volume snapshotclass** changes **volume snapshot status** to **Error**. Therefore, you cannot restore a PVC from the **volume snapshot**. To workaround this issue, recreate the **volume snapshotclass**.

(BZ#1902711)

Tainted nodes cannot be discovered by the device discovery wizard for Local Storage based deployments

Local Storage based deployments can now be deployed using the user interface on Openshift Container Platform v4.6. During the storage cluster creation, nodes cannot be discovered if Red Hat OpenShift Container Storage nodes have the taint **node.ocs.openshift.io/storage="true":NoSchedule** as **localvolumeset** and **localvolumediscovery** custom resources do not have the required toleration. For a workaround see Red Hat Knowledgebase Solution .

(BZ#1905373)

Device replacement action cannot be performed via UI for an encrypted OpenShift Container Storage cluster

On an encrypted OpenShift Container Storage cluster, the discovery result CR discovers the device backed by a Ceph OSD (Object Storage Daemon) differently from the one reported in the Ceph alerts. When clicking the alert, the user is presented with **Disk not found** message. Due to the mismatch, console UI cannot enable the disk replacement option for an OCS user. To workaround this issue, use the CLI procedure for failed device replacement.

(BZ#1906002)

CHAPTER 6. UNSUPPORTED FEATURES FOR IBM POWER SYSTEMS AND Z

The following OpenShift Container Storage 4.6 features are not supported on clusters deployed using IBM Power Systems and IBM Z. Click each feature for detailed information about the feature.

- Object storage
- FIPS
- Proxy environment
- External mode
- Disconnected environment