



Container-Native Storage 3.9

3.9 Release Notes

Release Notes for Container-Native Storage on OpenShift Container Platform 3.9
Edition 1

Container-Native Storage 3.9 3.9 Release Notes

Release Notes for Container-Native Storage on OpenShift Container Platform 3.9
Edition 1

Bhavana Mohan
Customer Content Services Red Hat
bmohanra@redhat.com

Legal Notice

Copyright © 2018 Red Hat, Inc.

This document is licensed by Red Hat under the [Creative Commons Attribution-ShareAlike 3.0 Unported License](https://creativecommons.org/licenses/by-sa/3.0/). If you distribute this document, or a modified version of it, you must provide attribution to Red Hat, Inc. and provide a link to the original. If the document is modified, all Red Hat trademarks must be removed.

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, 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 Software Collections 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

These release notes provide high-level coverage of the improvements and additions that have been implemented in Container-Native Storage for OpenShift Container Platform 3.9.

Table of Contents

CHAPTER 1. WHAT'S NEW IN THIS RELEASE?	3
CHAPTER 2. NOTABLE BUG FIXES	4
CHAPTER 3. KNOWN ISSUES	5
APPENDIX A. REVISION HISTORY	6
INDEX	7

CHAPTER 1. WHAT'S NEW IN THIS RELEASE?

This section describes the key features and enhancements in the Container-Native Storage 3.9 release.

- **Expanding the Persistent Volume Size:** With this release, the persistent volume size for dynamically provisioned volumes on file storage can be increased by increasing the persistent volume claim. The parameter `allowVolumeExpansion` has to be set to "true" in the storage class file to enable this feature. For more information, refer https://access.redhat.com/documentation/en-us/container-native_storage/3.9/html/container-native_storage_for_openshift_container_platform/chap-Documentation-Red_Hat_Gluster_Storage_Container_Native_with_OpenShift_Platform-OpenShift_Creating_Persistent_Volumes#sect_expanding_pv.
- **Providing a Custom Volume Name Prefix for a Persistent Volume :** You can provide a custom volume name prefix to the persistent volume that is created in both file and block storage. By providing a custom volume name prefix, users can now easily search/filter the volumes based on:
 - Any string that was provided as the field value of "volnameprefix" in the storageclass file.
 - Persistent volume claim name.
 - Project/Namespace name.

The parameter `volumenameprefix` has to be included in the storage class file to enable this feature. For more information to enable this for block storage, refer https://access.redhat.com/documentation/en-us/container-native_storage/3.9/html/container-native_storage_for_openshift_container_platform/Block_Storage#sect_block-custom-volname-prefix. For more information to enable this for file storage, refer, https://access.redhat.com/documentation/en-us/container-native_storage/3.9/html/container-native_storage_for_openshift_container_platform/chap-Documentation-Red_Hat_Gluster_Storage_Container_Native_with_OpenShift_Platform-OpenShift_Creating_Persistent_Volumes#sect_file-custom-volname-prefix.

- **Viewing Volume Metrics:** With this release you can view the details and various metrics that show the persistent volume consumption in addition to the allocated persistent volume size for dynamically provisioned volumes on file storage. The different metrics that can be viewed on Prometheus are:
 - `kubelet_volume_stats_available_bytes`: Number of available bytes in the volume.
 - `kubelet_volume_stats_capacity_bytes`: Capacity in bytes of the volume.
 - `kubelet_volume_stats_inodes`: Maximum number of inodes in the volume.
 - `kubelet_volume_stats_inodes_free`: Number of free inodes in the volume.
 - `kubelet_volume_stats_inodes_used`: Number of used inodes in the volume.
 - `kubelet_volume_stats_used_bytes`: Number of used bytes in the volume.

For more information about Volume metrics, refer https://access.redhat.com/documentation/en-us/container-native_storage/3.9/html/container-native_storage_for_openshift_container_platform/chap-documentation-red_hat_gluster_storage_container_native_with_openshift_platform-openshift_creating_persistent_volumes#enable_vol_metrics.

CHAPTER 2. NOTABLE BUG FIXES

This chapter describes bugs fixed in this release of Container-Native Storage for OpenShift Container Platform that have significant impact on users.

This release addresses many customers reported issues related to Heketi and Gluster database going out of synchronization leaving the system with stale metadata about a number of PV's, old deleted PV's, and failed transaction metadata. Container-Native Storage 3.9 handles Heketi and underlying Gluster subsystem more robustly.

This release also addresses some of the issues reported by customers on Block backed PV's specifically under failure scenario and after recovery from such failures.

heketi

BZ#1415750

Previously, deleting an heketi pod while some heketi operation was in progress would result in incomplete entries in the database. With this fix, such entries are marked "pending" until the operation is completed, thus leading to a consistent database view.

BZ#1434668

Earlier, the 'device info' output displayed the state of the device as 'failed' after a device remove operation was completed. With this fix, the state of the device is changed to 'removed' which matches with the operation performed.

BZ#1437798

Earlier, it was possible to run multiple device remove operations in parallel on the same device. This led to race conditions and database inconsistencies. With this fix, an error is returned while another device remove operation on the same device is already in progress.

kubernetes

BZ#1505290

Previously, the gluster-block provisioner did not identify the storage units correctly in the PVC. For example, it would identify 1 as 1GiB by default and the provisioner would fail on 1Gi. With this enhancement, gluster-block provisioner identifies the storage units correctly, ie, 1 will be treated as 1 byte, 1Gi will be treated as 1 GibiByte, and 1Ki will be treated as 1KibiByte.

CHAPTER 3. KNOWN ISSUES

This chapter provides a list of known issues at the time of release.

- [BZ#1560840](#)

In a gluster cluster with more than three nodes, if one or more nodes are down, but at least three nodes are up, heketi intermittently fails to create new replica 3 volumes, even if the healthy nodes have sufficient space available.

To workaround this issue, execute the following command on the unavailable nodes, which makes heketi volume creation 100% reliable:

```
# heketi-cli node disable <node-id>
```

After the nodes are available again, execute the following command to let heketi take this node into account again when creating volumes:

```
heketi-cli node enable <node-id>
```

- [BZ#1461131](#)

Volumes that were created using Container-Native Storage 3.5 or previous do not have the GID stored in heketi database. Hence, when a volume expansion is performed, new bricks do not get the group ID set on them which might lead to I/O errors.

- [BZ#1409848](#)

The following two lines might be repeatedly logged in the rhgs-server-docker container/gluster container logs.

```
[MSGID: 106006] [glusterd-svc-  
mgmt.c:323:glusterd_svc_common_rpc_notify] 0-management: nfs has  
disconnected from glusterd.  
[socket.c:701:__socket_rwv] 0-nfs: readv on  
/var/run/gluster/1ab7d02f7e575c09b793c68ec2a478a5.socket failed  
(Invalid argument)
```

These logs are added as glusterd is unable to start the NFS service. There is no functional impact as NFS export is not supported in Containerized Red Hat Gluster Storage.

APPENDIX A. REVISION HISTORY

Revision 1.0-4 Publishing for CNS 3.9 release.	Thu Apr 05 2018	Bhavana Mohan
Revision 1.0-3 Included a known issue and incorporated review comments.	Wed Apr 04 2018	Bhavana Mohan
Revision 1.0-2 Updated the What's New chapter and the Notable Bug Fixes chapter.	Tue Mar 27 2018	Bhavana Mohan
Revision 1.0-1 Initial creation by publican	Wed Mar 14 2018	Bhavana Mohan

INDEX