



Red Hat OpenShift Container Storage 3.11

3.11 Release Notes

Release Notes for Red Hat OpenShift Container Storage

Edition 1

Last Updated: 2018-10-25

Red Hat Openshift Container Storage3.11 3.11 Release Notes

Release Notes for Red Hat Openshift Container Storage
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](#). 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 Red Hat OpenShift Container Storage 3.11.

Table of Contents

CHAPTER 1. NOTABLE BUG FIXES	3
CHAPTER 2. KNOWN ISSUES	5
APPENDIX A. REVISION HISTORY	6
INDEX	7

CHAPTER 1. NOTABLE BUG FIXES

This chapter describes bugs fixed in this release of Red Hat Openshift Container Storage that have significant impact on users.

rhgs-server-container

BZ#1623433

Previously, bricks were not properly mounted in the gluster pod and lvmetad processes (one in the container, one on the host) were competing for resources. Hence, this prevented certain logical volumes from being detected and/or available which resulted in failed mounted bricks. With this fix, do not rely on the lvmetad service, and ensure it is not started in the container. Hence, resulting in a single metadata daemon process running and managing the devices and logical volumes.

gluster-block

BZ#1638988

Block volumes created on CNS version 3.9 or below use default Asymmetric Logical Unit Access (alua) group setting at target side configuration. These are meant to be used with 'prio const' at the initiator side multipath.conf. After upgrading OCS to 3.10 or above, the new recommendation at multipath is 'prio alua'. Old block volume having target side alua group set to default are not compatible with newly recommended multipath.conf. Hence, to maintain compatibility of old block volumes with the new multipath.conf recommendation, the default alua group should be changed to glfs alua group. With this update, the block volume target configuration is generated automatically as part of the upgrade to fix the compatibility issue. Hence, the compatibility of old block volume with the new multipath.conf recommendation is maintained.

BZ#1575908

The CLI audit logs have been enhanced to capture every CLI command and its response. This makes logs more detailed and easier to understand so that the problems are easier to debug.

BZ#1589666

Updates to the gluster-block configuration file are now applied automatically, removing the need to restart the gluster-block service when configuration changes.

BZ#1594331

Previously, using gluster-block on a non-recommended kernel version resulted in unexpected user space processes entering an uninterruptible sleep state. With this update, gluster-block now checks for minimum kernel version. If the minimum recommendation is not met, then gluster-block service stops. This results in a more stable gluster-block user experience.

tcmu-runner

BZ#1561323

Updates to the tcmu-runner configuration file are now applied automatically, removing the need to restart the tcmu-runner service when configuration changes.

BZ#1476730

Gluster-block operations (create/delete/modify) or gluster-block-target service restart, performed when tcmu-runner is in offline state, can trigger netlink hung issue, with targetcli process entering

uninterruptible sleep (D state) state forever. To recover from this state, restart the tcmu-runner daemon.

heketi

BZ#160042

When many volume create requests were made simultaneously, heketi spawned a very large number of goroutines, which greatly increased memory consumption by heketi. Heketi is now restricted to a certain number of concurrent operations in order to prevent this issue.

BZ#1629889

Previously, heketi ignored pvremove and vgrename errors when a device was removed while the device remove commands were used. Attempting to add the same disk again failed because it had not been properly removed in the first place. Heketi no longer ignores pvremove and vgrename errors, ensuring that devices are removed correctly, and can be re-added to Heketi after removal. Alternatively, you can also use the "--force-forget" flag with the device remove command to ignore any errors to ensure the same device can be added back to Heketi.

CHAPTER 2. KNOWN ISSUES

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

- [BZ#1597320](#)

Few paths are missing for iscsi mpath device. This is because, either CHAP security values mismatch or iscsi CSG: stage 0 has been skipped. Both will result in iSCSI login negotiation to fail.

To workaround this issue, delete the app pod which will trigger the restart of the pod. The pod start will then update the credentials and relogins.

- [BZ#1596021](#)

Sometimes in OCP, instead of using the multipath mapper device, such as `/dev/mapper/mpatha`, to mount the device, individual paths, such as `/dev/sdb`, is used. This is because OCP is currently not waiting for the mpath checker default timeout of 30 seconds, but waits for only 10 second and then picks the individual path if the mapper device is not ready. Due to this, we are not benefited by all the advantages of high availability and multipathing for block device.

To workaround this issue, delete the app pod which will then trigger the restart of the pod. The pod then relogins and undergoes a multipath check.

- [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 1.0-1
Publishing for 3.11 release

Wed Oct 24 2018

Bhavana Mohan

INDEX