Red Hat Hyperconverged Infrastructure for Virtualization 1.8

1.8 Release Notes

Release notes and known issues

Red Hat Hyperconverged Infrastructure for Virtualization 1.8 1.8 Release Notes

Release notes and known issues

Laura Bailey

lbailey@redhat.com

Legal Notice

Abstract

This document outlines known issues in Red Hat Hyperconverged Infrastructure for Virtualization 1.8 at release time, and highlights the major changes since the previous release.

Table of Contents

CHAPTER 1. WHAT CHANGED IN THIS RELEASE?

1.1. MAJOR CHANGES IN VERSION 1.8

Be aware of the following differences between Red Hat Hyperconverged Infrastructure for Virtualization 1.8 and previous versions:

Changed behavior

- Red Hat Hyperconverged Infrastructure for Virtualization 1.8 and Red Hat Virtualization 4.4 are based on Red Hat Enterprise Linux 8. Read about the key differences in Red Hat Enterprise Linux 8 in Considerations in adopting RHEL 8 .

- Cluster upgrades now require at least 10 percent free space on gluster disks in order to reduce the risk of running out of space mid-upgrade. This is available post the upgrade to Red Hat Hyperconverged Infrastructure for Virtualization 1.8. (BZ#1783750)

- "Hosts" and "Additional Hosts" tabs in the Web Console have been combined into a new "Hosts" tab that shows information previously shown on both. (BZ#1762804)

- Readcache and readcachesize options have been removed from VDO volumes, as they are not supported on Red Hat Enterprise Linux 8 based operating systems. (BZ#1808081)

- The Quartz scheduler is replaced with the standard Java scheduler to match support with Red Hat Virtualization. (BZ#1797487)

Enhancements

- The Administrator Portal can now upgrade all hosts in a cluster with one click. This is available post the upgrade to Red Hat Hyperconverged Infrastructure for Virtualization 1.8. (BZ#1721366)

- At-rest encryption using Network-Bound Disk Encryption (NBDE) is now supported on new Red Hat Hyperconverged Infrastructure for Virtualization deployments. (BZ#1821248, BZ#1781184)

- Added support for IPv6 networking. Environments with both IPv4 and IPv6 addresses are not supported. (BZ#1721383)

- New roles, playbooks, and inventory examples are available to simplify and automate the following tasks:

    - Upgrading (BZ#1500728, BZ#1832654)

    - Backing up and restoring configuration (BZ#1850488)

    - Replacing hosts (BZ#1840123)

    - Blacklisting multipath devices (BZ#1807808)

    - Creating the gluster logical network (BZ#1832966)

    - Deploying on IPv6 networks. (BZ#1688217)

- Added an option to select IPv4 or IPv6 based deployment in the web console. (BZ#1688798)

- **fsync** in the replication module now uses **eager-lock** functionality which improves the performance of small-block of size approximately equal to 4k write-heavy workloads by more than 50 percent on Red Hat Hyperconverged Infrastructure for Virtualization 1.8. (BZ#1836164)

- The web console now supports blacklisting multipath devices. (BZ#1814120)

- New fencing policies **skip_fencing_if_gluster_bricks_up** and **skip_fencing_if_gluster_quorum_not_met** are now added and enabled by default. (BZ#1775552)

- Red Hat Hyperconverged Infrastructure for Virtualization now ensures that the "performance.strict-o-direct" option in Red Hat Gluster Storage is enabled before creating a storage domain. (BZ#1807400)

- Red Hat Gluster Storage volume options can now be set for all volumes in the Administrator Portal by using "all" as the volume name. (BZ#1775586)

- Read-only fields are no longer included in the web console user interface, making the interface simpler and easier to read. (BZ#1814553)

CHAPTER 2. BUG FIXES

This section documents important bugs that affected previous versions of Red Hat Hyperconverged Infrastructure for Virtualization that have been corrected in version 1.8.

BZ#1688239 - Geo-replication with IPv6 networks

Previously, geo-replication could not be used with IPv6 addresses. With this release, all the helper scripts used for gluster geo-replication are made compatible with IPV6 hostnames (FQDN).

BZ#1719140 - Virtual machine availability

Virtualization group options are updated so that the virtual machines are available when one of the hosts is powered down.

BZ#1792821 - Split-brain after heal

Previously, healing of entries in directories could be triggered when only the heal source (and not the heal target) was available. This led to replication extended attributes being reset and resulted in a GFID split-brain condition when the heal target became available again. Entry healing is now triggered only when all bricks in a replicated set are available, to avoid this issue.

BZ#1721097 - VDO statistics

Previously, Virtual Disk Optimization (VDO) statistics were not available for VDO volumes. With this release, Virtual Disk Optimization (VDO) handles different outputs correctly from the VDO statistics tool to avoid traceback caused by Virtual Desktop and Servers Manager.

BZ#1836164 - Replication of write-heavy workloads

fsync in the replication module uses **eager-lock** functionality which improves the performance of small-block of size approximately equal to 4k write-heavy workloads with an improvement of more than 50 percent on Red Hat Hyperconverged Infrastructure for Virtualization 1.8.

BZ#1821763 - VDO volumes and maxDiscardSize parameter

The virtual disk optimization (VDO) Ansible module now supports the maxDiscardSize parameter and sets this parameter by default. As a result, VDO volumes created with this parameter no longer fail.

BZ#1808081 - Readcache and readcachesize on VDO volumes

The *readcache* and *readcachesize* options for virtual disk optimization (VDO) are not supported on VDO volumes based on Red Hat Enterprise Linux 8, which includes Red Hat Hyperconverged Infrastructure for Virtualization 1.8. These options are now removed so that VDO volume creation succeeds on version 1.8.

BZ#1793398 - Deployment using Ansible

Previously, running the deployment playbook from the command line interface failed because of incorrect values for variables (**he_ansible_host_name** and **he_mem_size_MB**). Variable values have been updated and the deployment playbook now runs correctly.

BZ#1809413 - Activating glusterd service caused quorum loss

Previously, activating the host from the Administrator Portal restarted the glusterd service which led to quorum loss when the glusterd process ID changed. With this release, the glusterd service does not restart if it is already up and running during the activation of the host, so the glusterd process ID does not change and there is no quorum loss.

BZ#1796035 - Additional hosts in Administrator Portal

Previously, additional hosts were not added to the Administrator Portal automatically post deployment. With this release, gluster ansible roles have been updated to ensure that any additional hosts are automatically added to the Administrator Portal

BZ#1774900 - Disconnected host detection

Previously, the detection of disconnected hosts took a long time leading to sanlock timeout. With

this release, the socket and rpc timeouts in gluster have been improved so that disconnected hosts are detected before sanlock timeout occurs and reboot of a single host does not impact virtual machines running on other hosts.

BZ#1795928 - Erroneous deployment failure message

When deployment playbooks were run on the command line interface, a web hook was successfully added to gluster-eventsapi, but reported a failure instead of a success, causing deployment to fail the first time it was attempted. This has been corrected and deployment now works correctly.

BZ#1715428 - Storage domain creation

Previously, storage domains were automatically created only when additional hosts were specified. The two operations have been separated, since they are logically unrelated, and storage domains are now created regardless of whether additional hosts are specified.

BZ#1733413 - Incorrect volume type displayed

Previously, the web console contained an unnecessary drop-down menu for volume type selection and showed the wrong volume type (replicated) for single node deployments. The menu is removed and the correct volume type (distributed) is now shown.

BZ#1754743 - Cache volume failure on VDO volumes

Previously, configuring volumes that used both virtual disk optimization (VDO) and a cache volume caused deployment in the web console to fail. This occurred because the underlying volume path was specified in the form "/dev/sdx" instead of the form "/dev/mapper/vdo_sdx". VDO volumes are now specified using the correct form and deployment no longer fails.

BZ#1432326 - Network out of sync after associating a network with a host

When the Gluster network was associated with a new node's network interface, the Gluster network entered an out of sync state. This no longer occurs in Red Hat Hyperconverged Infrastructure for Virtualization 1.8.

BZ#1567908 - Multipath entries for devices visible after rebooting

The vdsm service makes various configuration changes after Red Hat Virtualization is first installed. One such change made multipath entries for devices visible in Red Hat Hyperconverged Infrastructure for Virtualization, including local devices. This caused issues on hosts that were updated or rebooted before the Red Hat Hyperconverged Infrastructure for Virtualization deployment process was complete. Red Hat Hyperconverged Infrastructure for Virtualization now provides the option to blacklist multipath devices, which prevents any entries from being used by RHHI for Virtualization.

BZ#1590264 - Storage network down after Hosted Engine deployment

During Red Hat Hyperconverged Infrastructure for Virtualization setup, two separate network interfaces are required to set up Red Hat Gluster Storage. After storage configuration is complete, the hosted engine is deployed and the host is added to the engine as a managed host. Previously, during deployment, the storage network was altered to remove **BOOTPROTO=dhcp**. This meant that the storage network did not have an IP addresses assigned automatically, and was not available after the hosted engine was deployed. This line is no longer removed during deployment, and the storage network is available as expected.

BZ#1609451 - Volume status reported incorrectly after reboot

When a node rebooted, including as part of upgrades or updates, subsequent runs of **gluster volume status** sometimes incorrectly reported that bricks were not running, even when the relevant **glusterfsd** processes were running as expected. State is now reported correctly in these circumstances.

BZ#1856629 - Warning **use device with format** /**dev/mapper/<WWID>** seen with gluster devices enabled

When expanding the volume from the web console as part of day2 operation with device name /**dev/sdx**, a warning **use device name with format /dev/mapper/<WWID>** was seen even when the blacklist gluster device checkbox was enabled. In version 1.8, Red Hat recommends to ignore this

warning note and proceed with the next step of deployment, as the warning was not valid. In version 1.8 Batch Update 1, this issue has been corrected and the spurious warning no longer appears.

CHAPTER 3. KNOWN ISSUES

This section documents unexpected behavior known to affect Red Hat Hyperconverged Infrastructure for Virtualization (RHHI for Virtualization).

BZ#1851114 – **Error message** **device path is not a valid name for this device** is shown

When logical volume (LV) name exceeds 55 characters which is a limitation with python–blivet, error message like **ValueError: gluster_thinpool_gluster_vg_<WWID> is not a valid name for this device** is seen in **vdsm.log** and **supervdsm.log** files.
To work around this issue, follow these steps:

1. Rename the volume group (VG):

   ```
   # vgrename gluster_vg_<WWID> gluster_vg_<last-4-digit-WWID>
   ```

2. Rename thinpool:

   ```
   # lvrename gluster_vg_<last-4-digit-WWID> gluster_thinpool_gluster_vg_<WWID>
   gluster_thinpool_gluster_vg_<last-4-digit-WWID>
   ```

BZ#1853995 – **Updating storage domain gives error dialog box**

While replacing the primary volfile during host replacement to update the storage domain via the Administrator Portal, the portal gives an **operation canceled dialog box**. However, in the backend the values get updated.

BZ#1855945 – **RHHI for Virtualization depolyment fails using multipath configuration and lvm cache**

During the deployment of RHHI for Virtualization with multipath device names, volume groups (VG) and logical volumes (LV) are created with the suffix of WWID leading to LV names longer than 128 characters. This results in failure of LV cache creation.
To work around this issue, follow these steps:

When initiating RHHI for Virtualization deployment with multipath device names as ``/dev/mapper/<WWID>``, replace VG and thinpool suffix with last 4 digits of WWID as:

1. During deployment from the web console, provide a multipath device name as **/dev/mapper/<WWID>** for bricks.

2. Click **Next** to generate an inventory file.

3. Login in to the deployment node via SSH.

4. Find the *<WWID>* with LVM components:

   ```
   # grep vg /etc/ansible/hc_wizard_inventory.yml
   ```

5. For all WWIDs, replace WWID with the last 4 digits of WWID.

   ```
   # sed -i 's/<WWID>/<last-4-digit-WWID>/g' /etc/ansible/hc_wizard_inventory.yml
   ```

6. Continue deployment from web console.

BZ#1856577 – **Shared storage volume fails to mount in IPv6 environment**

When gluster volumes are created with the **gluster_shared_storage** option during the deployment of RHHI for Virtualization using IPv6 addresses, the mount option is not added in the **fstab** file. As a result, the shared storage fails to mount.

To workaround this issue, add mount option as **xlator-option=transport.address-family=inet6** in the **fstab** file.

BZ#1856594 – Fails to create VDO enabled gluster volume with day2 operation from web console

Virtual Disk Optimization (VDO) enabled gluster volume with day2 operation fails from the web console.

To work around this issue, modify the playbook **vdo_create.yml** at **/etc/ansible/roles/gluster.infra/roles/backend_setup/tasks/vdo_create.yml** and change the ansible tasks as:

```
- name: Restart all VDO volumes
  shell: "vdo stop -n {{item.name}} && vdo start -n {{item.name}}"
  with_items: "{{ gluster_infra_vdo }}"
```

BZ#1858197 – Pending self-heal on the volume, post the bricks are online

In dual network configurations (one for gluster and other for ovirt management), Automatic File Replication (AFR) healer threads are not spawned with the restart of self-heal daemon resulting in pending self heal entries in the volume.

To work around this issue, follow these steps:

1. change the hostname to the other network FQDN using the command

   ```
   # hostnamectl set-hostname <other-network-FQDN>
   ```

2. Start the heal using the command:

   ```
   # gluster volume heal <volname>
   ```

BZ#1554241 – Cache volumes must be manually attached to asymmetric brick configurations

When bricks are configured asymmetrically, and a logical cache volume is configured, the cache volume is attached to only one brick. This is because the current implementation of asymmetric brick configuration creates a separate volume group and thin pool for each device, so asymmetric brick configurations would require a cache volume per device. However, this would use a large number of cache devices, and is not currently possible to configure using the Web Console.

To work around this issue, first remove any cache volumes that have been applied to an asymmetric brick set.

```
# lvconvert --uncache volume_group/logical_cache_volume
```

Then, follow the instructions in Configuring a logical cache volume to create a logical cache volume manually.

BZ#1690820 – Create volume populates host field with IP address not FQDN

When you create a new volume using the Web Console using the **Create Volume** button, the value for hosts is populated from gluster peer list, and the first host is an IP address instead of an FQDN. As part of volume creation, this value is passed to an FQDN validation process, which fails with an IP address.

To work around this issue, edit the generated variable file and manually insert the FQDN instead of the IP address.

BZ#1506680 – Disk properties not cleaned correctly on reinstall

The installer cannot clean some kinds of metadata from existing logical volumes. This means that reinstalling a hyperconverged host fails unless the disks have been manually cleared beforehand. To work around this issue, run the following commands to remove all data and metadata from disks attached to the physical machine.

> **WARNING**
>
> Back up any data that you want to keep before running these commands, as these commands completely remove all data and metadata on all disks.

```
# pvremove /dev/* --force -y
# for disk in $(ls /dev/{sd*,nv*}); do wipefs -a -f $disk; done
```

BZ#1506680 – Disk properties not cleaned correctly on reinstall