



Red Hat Hyperconverged Infrastructure for Virtualization 1.6

1.6 Release Notes

Release notes and known issues

Red Hat Hyperconverged Infrastructure for Virtualization 1.6 1.6 Release Notes

Release notes and known issues

Laura Bailey
lbailey@redhat.com

Legal Notice

Copyright © 2019 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

This document outlines known issues in Red Hat Hyperconverged Infrastructure for Virtualization 1.6 at release time, and highlights the major changes since the previous release.

Table of Contents

CHAPTER 1. WHAT CHANGED IN THIS RELEASE?	3
1.1. MAJOR CHANGES IN VERSION 1.6	3
1.2. TECHNOLOGY PREVIEWS	3
1.3. PROOF OF CONCEPT TECHNOLOGY	3
1.4. DEPRECATIONS	4
CHAPTER 2. BUG FIXES	5
CHAPTER 3. KNOWN ISSUES	7

CHAPTER 1. WHAT CHANGED IN THIS RELEASE?

1.1. MAJOR CHANGES IN VERSION 1.6

Be aware of the following differences between Red Hat Hyperconverged Infrastructure for Virtualization 1.6 and previous versions:

Ansible-based deployment and management

RHHI for Virtualization now uses Ansible playbooks for all deployment and management tasks. Documentation has been updated accordingly.

Expand volumes across more than 3 nodes

Volumes can now span across 3, 6, 9, or 12 nodes. See [Volume Types](#) for support details. See [Expanding an existing volume across more hyperconverged nodes](#) for instructions on expanding existing volumes across more nodes.

RHHI for Virtualization sizing tool

Visit the [RHHI for Virtualization Sizing Tool](#), enter your deployment requirements and click **Solve** to receive an example configuration with suggested nodes, memory, and resource commitments for your deployment.

1.2. TECHNOLOGY PREVIEWS

Technology Preview features are provided with a limited support scope, as detailed on the Customer Portal: [Technology Preview Features Support Scope](#).

The features listed in this section are provided under Technology Preview support limitations.

IPv6 networking support

Technology Preview support is now available for IPv6 addresses in IPv6-only environments (including DNS and gateway addresses). There are several important limitations to IPv6 support during this Technology Preview:

- Disaster recovery using geo-replication is not supported ([BZ#1688269](#))
- An additional mount option is required when mounting storage provided by IPv6 hosts: **xlator-option=transport.address-family=inet6** ([BZ#1688243](#))
- Hosts cannot be automatically added to the RHV Administration Portal ([BZ#1688239](#))

1.3. PROOF OF CONCEPT TECHNOLOGY

Proof of Concept technology is available for use in testing environments.

The technology listed in this section is not supported and not intended for production environments.

OpenShift Container Storage on RHHI for Virtualization

OpenShift Container Storage can now be deployed on top of RHHI for Virtualization to let you create and manage containers and their underlying storage on a small cluster of hyperconverged hosts. The basic deployment structure uses four hyperconverged hosts that provide storage and virtualization capabilities. These hyperconverged nodes host six virtual machines installed with Red Hat OpenShift Container Platform. Contact your Red Hat support representative to access this technology for Proof of Concept testing purposes.

1.4. DEPRECATIONS

The following items are considered deprecated or not supported as of Red Hat Hyperconverged Infrastructure for Virtualization 1.6:

Gdeploy is no longer supported

Ansible playbooks are now used for all deployment and management tasks. Gdeploy is therefore no longer supported after upgrading to Red Hat Hyperconverged Infrastructure for Virtualization 1.6.

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 1.6.

[BZ#1638639](#) - Mount options in single-node deployment UI

When a single node deployment was created using the Web Console, the **backup-volfile-servers** option was automatically added to the **Mount options** field on the **Storage** tab of the deployment wizard, even though the deployment used a single node. This has now been corrected.

[BZ#1643730](#) - Tuning during host upgrade

When a hyperconverged host was upgraded to any Red Hat Hyperconverged infrastructure for Virtualization 1.5 release, the performance profile for the virt group was not updated with parameters added during development. This meant that upgraded deployments were not optimally tuned for most use cases. This has now been corrected, and the following options exist in the **/var/lib/glusterd/groups/virt** file on each server after upgrade.

[BZ#1666714](#) - Edits during deployment

Previously, when the generated configuration file was edited by the user editing values in the tabs of the UI during the *Review* step of deployment, the original contents of the generated configuration file were not updated. This has now been corrected and the file is updated with new values as expected.

[BZ#1667209](#) - Mount options for VDO and non-VDO volumes

Previously, when a deployment used virtual disk optimization (VDO), options related to VDO devices were added to non-VDO disks in the **/etc/fstab** file. This meant that disks failed to mount after a reboot. This has been corrected so that devices are configured correctly and all disks mount correctly after reboot.

[BZ#1610699](#) - Thin provisioned volumes and VDO

Parts of the Web Console implied that thin provisioning and VDO could be used together, even though this combination is not supported. The user interface has now been updated so that thin provisioning is disabled when VDO volumes are present.

[BZ#1629542](#) - VDO logical size

Previously, the logical size of a VDO volume was 10 times the size of the brick by default. This meant that using the default settings wasted storage space. The default size of the brick and the default logical size of the VDO volume are now the same.

[BZ#1641592](#) - Virtual disk allocation policy

By default, virtual disks created during deployment used the *sparse* allocation policy, which reduced performance for Red Hat Hyperconverged Infrastructure for Virtualization. The default allocation policy is now *preallocated* to improve performance.

[BZ#1674713](#) - Cleaning up failed deployments

If Red Hat Hyperconverged Infrastructure for Virtualization deployment fails, you can now start the deployment process from scratch by running the `gluster_cleanup.yml` playbook. See documentation for details.

[BZ#1518164](#) - Fetching brick profile statistics

Previously, attempts to fetch brick profile statistics failed because of a missing check for NFS access, even though the statistics were fetched correctly. This meant that attempting to view the profile of a volume failed with the following error message: "Could not fetch brick profile stats". This is now checked correctly, and works as expected.

[BZ#1641483](#) - Single node volume type

When a single node deployment was created using the Web Console, the UI indicated that the volume created was a replicated volume type, when the volume was actually a distributed volume

type. This has been corrected and the UI now correctly shows a distributed volume for this deployment type.

CHAPTER 3. KNOWN ISSUES

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

BZ#1754743 - Enabling LV cache along with VDO volumes fails during Deployment

If LV cache is attached to a Virtual Data Optimizer (VDO) enabled thinpool device, then the path for the cache device should be `/dev/mapper/vdo_<device name>` instead of `/dev/<device name>`

To work around this issue, edit the generated Ansible inventory file.

If VDO is created from **sdb** and cache disk is **sdC** then the configuration should be as follows:

```
gluster_infra_cache_vars:
  - vgname: gluster_vg_sdb
    cachedisk: `/dev/mapper/vdo_sdb,/dev/sdC`
    cachelvname: cachelv_gluster_thinpool_gluster_vg_sdb
    cachethinpoolname: gluster_thinpool_gluster_vg_sdb
    cachelvsize: 1G
    cachemode: writethrough
```

BZ#1432326 - Associating a network with a host makes network out of sync

When the Gluster network is associated with a Red Hat Gluster Storage node's network interface, the Gluster network enters an out of sync state. To work around this issue, click the **Management** tab that corresponds to the node and click **Refresh Capabilities**.

BZ#1547674 - Default slab size limits VDO volume size

Virtual Data Optimizer (VDO) volumes are comprised of up to 8096 slabs. The default slab size is 2GB, which limits VDO volumes to a maximum physical size of 16TB, and a maximum supported logical size of 160TB. This means that creating a VDO volume on a physical device that is larger than 16TB fails.

To work around this issue, specify a larger slab size at VDO volume creation.

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#1590264 - Storage network down after Hosted Engine deployment

During RHHI 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 Red Hat Virtualization as a managed host. During host deployment, the storage network is altered, and **BOOTPROTO=dhcp** is removed. This means that the storage network does not have IP addresses assigned automatically, and is not available after the hosted engine is deployed.

To work around this issue, add the line **BOOTPROTO=dhcp** to the network interface configuration file for your storage network after deployment of the hosted engine is complete.

BZ#1567908 - Multipath entries for devices are visible after rebooting

The vdsmd service makes various configuration changes after Red Hat Virtualization is first installed. One such change makes multipath entries for devices visible on RHHI for Virtualization, including local devices. This causes issues on hosts that are updated or rebooted before the RHHI for Virtualization deployment process is complete.

To work around this issue, add all devices to the multipath blacklist:

- a. Add the following to **/etc/multipath.conf**:

```
blacklist {
    devnode "*"
}
```

- b. Reboot all hyperconverged hosts.

BZ#1589041 - NFS statistics queried on bricks without NFS access

When the profiling window is opened, NFS profile statistics are queried even when a volume does not have any NFS clients or NFS access enabled. This results in the brick profile UI reporting the following error: **Could not fetch brick profile stats**.

To work around this issue, select a brick from the profile information window and refresh details to see profiling details for the brick.

BZ#1637898 - Hosted Engine VM pauses on network attach

During deployment using DHCP, the hosted engine virtual machine pauses after attaching the storage logical network to the network interface. There is currently no workaround for this issue.

BZ#1608268 - the Web Console UI does not allow logical volume cache on thickly provisioned volumes

When you attempt to configure a logical volume cache for a thickly provisioned volume using the Web Console, the deployment fails. You can manually configure a logical volume cache after deployment by adding a faster disk to your volume group using the following procedure. Note that device names are examples.

- a. Add the new SSD (**/dev/sdc**) to the volume group (**gluster_vg_sdb**).

```
# vgextend gluster_vg_sdb /dev/sdc
```

- b. Create a logical volume (**cachelv**) from the SSD to use as a cache.

```
# lvcreate -n cachelv -L 220G gluster_vg_sdb /dev/sdc
```

- c. Create a cache pool from the new logical volume.

```
# lvconvert --type cache-pool gluster_vg_sdb/cachelv
```

- d. Attach the cache pool to the thickly provisioned logical volume as a cache volume.

```
# lvconvert --type cache gluster_vg_sdb/cachelv gluster_vg_sdb/gluster_thick_lv1
```

BZ#1609451 - Volume status reported incorrectly after reboot

When a node reboots, including as part of upgrades or updates, subsequent runs of **gluster volume status** sometimes incorrectly report that bricks are not running, even when the relevant **glusterfsd** processes are running as expected.

To work around this issue, restart **glusterd** on the hyperconverged node after the node is upgraded, updated, or rebooted.

BZ#1688269 - IPv6 hosts not added to Red Hat Virtualization Manager

When IPv6 addresses are used to deploy Red Hat Hyperconverged Infrastructure for Virtualization, the second and third hyperconverged hosts are not automatically configured to be managed by the hosted engine virtual machine during the deployment process.

To work around this issue, manually add the second and third hosts to Red Hat Virtualization Manager when deployment is complete: [Adding a host](#).

BZ#1688239 - Geo-replication session creation fails with IPv6

The **gverify.sh** script is used during geo-replication to verify that secondary volumes can be mounted before data is synchronized. When IPv6 is used, the script fails, because it does not use the **--xlator-option=transport.address-family=inet6** mount option as part of its checks. As a result, geo-replication cannot currently be used with IPv6 addresses.

BZ#1692671 - Deployment fails for static host names

The ansible-based deployment process introduced in Red Hat Hyperconverged Infrastructure for Virtualization uses the **dig** tool to verify FQDNs. This means that when static IP addresses are used and defined in the **/etc/hosts** file instead of a DHCP server being used, **dig** cannot detect them.

To work around this issue, add **gluster_features_fqdn_check: false** to each host definition in the generated playbook file, like so:

```
hosts:
  rhh1host1.example.com:
    gluster_features_fqdn_check: false
    gluster_infra_volume_groups:
      - vgname: gluster_vg_sdb
        pvname: /dev/sdb
      - vgname: gluster_vg_sdc
        pvname: /dev/mapper/vdo_sdc
```

BZ#1688243 - IPv6 deployments require additional mount option

Deployments that use IPv6 require an additional mount option to work correctly. Work around this issue by adding **xlator-option="transport.address-family=inet6"** to the **Mount Options** field on the **Storage** tab during Hosted Engine deployment.

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#1649502 - Vague host upgrade error

If pending heals exist or gluster quorum cannot be maintained when you attempt to upgrade an online host, the host upgrade fails with a message that indicates the host could not be moved to maintenance, but does not specify the exact cause of error.

To work around this issue, move the host into maintenance mode before attempting to upgrade.

BZ#1693144 - Expand cluster instead creates new cluster

When adding a new volume using the **Expand Cluster** feature in the Web Console, the new volume is added as a new cluster instead of adding a new volume to an existing cluster.

To work around this issue, peer probe the new hyperconverged hosts from hosts in the existing cluster before using **Expand Cluster**.

Log in to a host that is part of the cluster you wanted to expand, and run the following command for each host to be added:

```
# gluster peer probe <new-host-fqdn>
```

Then use **Expand Cluster** in the Web Console, using the same FQDN used with the peer probe command.

BZ#1694604 - FUSE mount process crash when LRU-cache fills during file operation

When sharding is enabled on a volume, a single file allocation operation creates all shards in a single batch. If the number of shards involved in an operation is greater than the number of entries allowed in the lru-cache, the inode associated with the file operation is freed while it is still in use. This leads to a crash in the mount process when deleting large files from volumes with sharding enabled, which causes all virtual machines that had mounted that storage to pause. To work around this issue, forcefully shut down the virtual machines and start them again.

BZ#1697704 - Null pointer exception while configuring remote data sync

When you attempt to use *Remote data sync setup* to configure geo-replication for a storage domain that was created automatically during RHHI for Virtualization deployment, *GetGeoRepSessionsForStorageDomainQuery* throws a null pointer exception. This issue does not occur for manually created storage domains.

To work around this issue, edit the storage domain and check the **Use managed gluster volume** checkbox.

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
```

