



Red Hat Hyperconverged Infrastructure for Virtualization 1.5

1.5 Release Notes

Release notes and known issues

Red Hat Hyperconverged Infrastructure for Virtualization 1.5 1.5 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, 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

This document outlines known issues in Red Hat Hyperconverged Infrastructure for Virtualization 1.5 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 1.5	3
CHAPTER 2. BUG FIXES	4
CHAPTER 3. KNOWN ISSUES	5
CHAPTER 4. TECHNOLOGY PREVIEWS	8

CHAPTER 1. WHAT CHANGED IN THIS RELEASE?

1.1. MAJOR CHANGES IN VERSION 1.5

Be aware of the following differences between Red Hat Hyperconverged Infrastructure for Virtualization 1.5 and previous versions.

Deduplication and compression support with Virtual Data Optimizer (VDO)

Configuring VDO at deployment time lets you reduce the amount of storage space required for data. You can configure Gluster bricks to use deduplication and compression, and monitor and configure notifications for VDO capacity usage so that you know when your storage is running out of space. The space saved by using Virtual Disk Optimization is displayed on the Brick and Volume detail pages of the Cockpit UI. See [Understanding VDO](#) and [Monitoring VDO](#) for more information.

Configure disaster recovery with failover and failback

Red Hat Hyperconverged Infrastructure for Virtualization now supports backup, failover, and failback to a remote secondary site. See [Configuring backup and recovery options](#) for an overview and information on configuring disaster recovery. See [Recovering from disaster](#) for the recovery process.

Scale using the user interface

New nodes can now be prepared and configured in Cockpit. See [Expanding the hyperconverged cluster by adding a new volume on new nodes using Cockpit](#) for details.

Upgrade using the user interface

Upgrade your deployment using the Administration Portal. See [Upgrading Red Hat Hyperconverged Infrastructure](#) for details.

Manage your storage and virtual machines in Cockpit

You can now view and manage your storage and your virtual machines from Cockpit. The Red Hat Virtualization Administration Console is still required for some more complex tasks, such as geo-replication. See [Managing Red Hat Gluster Storage using Cockpit](#) for more information.

Configure different devices and device types

Previous versions of RHHI for Virtualization expected each virtualization host to be set up the same way, with the same device types and device names on each. As of version 1.5, you can specify different device and sizes as appropriate for each host and size arbiter bricks appropriately.

Updated user interfaces

Cockpit and the Administration Portal have seen a number of updates to their user interfaces. Operations are now better organised and easier to find, and a number of new options are available.

- Specify additional hosts during Cockpit setup instead of adding them manually after Hosted Engine deployment.
- Reset brick configuration after reinstalling a virtualization host.

Deploy on a single node

Single node deployments of Red Hat Hyperconverged Infrastructure for Virtualization are now supported. See [Deploying RHHI for Virtualization in the data center](#) for support limitations and deployment details.

Convert virtualization hosts

Red Hat Virtualization hosts can now be converted into hyperconverged hosts. See [Converting a virtualization cluster to a hyperconverged cluster](#) for details.

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 1.5.

BZ#1489346 - Storage details needed to be entered twice

After configuring Red Hat Gluster Storage using Cockpit, users can either continue to Hosted Engine configuration right away, or continue deployment at another time. Previously, continuing deployment later resulted in users being prompted for information they had entered during the Red Hat Gluster Storage configuration process in Cockpit, as this information was not retained to be used during Hosted Engine deployment. Red Hat Gluster Storage configuration information is now retained and accessed by the Hosted Engine deployment wizard so that users do not need to enter this information twice.

BZ#1554487 - All geo-replication sessions failed when one failed

When a storage domain was configured to synchronize data to two secondary sites, and one of the sessions moved to a faulty state, the second session also failed to synchronize. This occurred because the locks for the faulty session were not released correctly on session failure, so the second session waited for locks to be released and could not synchronize. Locks are now released correctly when geo-replication sessions move to a faulty state, and this issue no longer occurs.

CHAPTER 3. KNOWN ISSUES

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

BZ#1425767 - Sanity check script does not fail

The sanity check script sometimes returns zero (success) even when disks do not exist, or are not empty. Since the sanity check appears to succeed, gdeploy attempts to create physical volumes and fails. To work around this issue, ensure that the **disk** value in the gdeploy configuration file is correct, and that the disk has no partitions or labels, and retry deployment.

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 Cockpit.

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 the engine 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 virtualization 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#1641483 - Inaccurate display in single-node deployment UI

When Cockpit is used to create a single-node deployment, the *Volumes* tab shows a volume type of **replicated**. However, the volume created in this deployment type is a single-brick distributed volume.

BZ#1640465 - Hosted Engine virtual machine migration fails

Migration of the Hosted Engine virtual machine from one host to another fails because libvirt incorrectly identifies FUSE mounted gluster volumes as being part of a shared file system, and as such does not identify the Hosted Engine virtual machine storage as being in need of migration. There is currently no workaround for this issue.

BZ1643730 - Upgraded hosts are not completely tuned

When a virtualization host is upgraded to any Red Hat Hyperconverged infrastructure for Virtualization 1.5 release, the performance profile for the virt group is not updated with parameters added during development. This means that upgraded deployments are not optimally tuned for most use cases. To work around this issue, add the following options to the `/var/lib/glusterd/groups/virt` file on each server, or manually set these options on your upgraded volumes.

```
cluster.choose-local=off
client.event-threads=4
server.event-threads=4
performance.client-io-threads=on
```

Run the following command for all volumes to apply the updated `virt` profile.

```
# gluster volume set <volname> group virt
```

BZ#1638639 - Incorrect mount option in single-node deployment UI

When you create a single node deployment using the Cockpit UI, the **backup-volfile-servers** option is automatically added to the **Mount options** field on the **Storage** tab of the deployment wizard. To work around this issue, delete **backup-volfile-servers=:** from the Mount options field before continuing the deployment.

BZ#1638630 - Generated configuration file does not display

On the final step of the deployment process, the **Review** page of the deployment wizard does not display the generated configuration file. Click **Reload** to force the configuration file to display.

BZ#1637898 - Hosted Engine VM pauses on network attach

During deployment, 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 - Cockpit 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 Cockpit UI, 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 (**cache1v**) from the SSD to use as a cache.

```
# lvcreate -n cache1v -L 220G gluster_vg_sdb /dev/sdc
```

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

```
# lvconvert --type cache-pool gluster_vg_sdb/cache1v
```

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

```
# lvconvert --type cache gluster_vg_sdb/cache1v  
gluster_vg_sdb/gluster_thick_lv1
```

CHAPTER 4. 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.

Automating deployments using Ansible

You can use Ansible to deploy Red Hat Hyperconverged Infrastructure for Virtualization without needing to watch and tend to the deployment process. Read [Automating RHHI for Virtualization deployment](#) for more information about this process.