



Red Hat Hyperconverged Infrastructure for Virtualization 1.6

Maintaining Red Hat Hyperconverged Infrastructure for Virtualization

Common maintenance tasks for Red Hat Hyperconverged Infrastructure for
Virtualization

Red Hat Hyperconverged Infrastructure for Virtualization 1.6 Maintaining Red Hat Hyperconverged Infrastructure for Virtualization

Common maintenance tasks for Red Hat Hyperconverged Infrastructure for Virtualization

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

Red Hat Hyperconverged Infrastructure for Virtualization (RHVI for Virtualization) combines compute, storage, networking, and management capabilities into a single solution, simplifying deployment and reducing the cost of acquisition and maintenance. This document explains how to perform maintenance tasks specific to Red Hat Hyperconverged Infrastructure for Virtualization.

Table of Contents

PART I. CONFIGURATION TASKS	4
CHAPTER 1. ADD COMPUTE AND STORAGE RESOURCES	5
1.1. CREATING NEW BRICKS USING ANSIBLE	5
1.2. CREATING NEW BRICKS ABOVE VDO LAYER USING ANSIBLE	8
1.3. EXPANDING VOLUME FROM RED HAT VIRTUALIZATION MANAGER	10
1.4. EXPANDING THE HYPERCONVERGED CLUSTER BY ADDING A NEW VOLUME ON NEW NODES USING THE WEB CONSOLE	12
CHAPTER 2. CONFIGURE HIGH AVAILABILITY USING FENCING POLICIES	16
2.1. CONFIGURING FENCING POLICIES IN THE CLUSTER	16
2.2. CONFIGURING FENCING PARAMETERS ON THE HOSTS	16
CHAPTER 3. CONFIGURING BACKUP AND RECOVERY OPTIONS	20
3.1. PREREQUISITES	20
3.1.1. Prerequisites for geo-replication	20
3.1.2. Prerequisites for failover and failback configuration	20
3.2. SUPPORTED BACKUP AND RECOVERY CONFIGURATIONS	20
3.3. CONFIGURING BACKUP TO A SECONDARY VOLUME	21
3.3.1. Prerequisites	21
3.3.1.1. Enable shared storage on the source volume	21
3.3.1.2. Match encryption on source and target volumes	21
3.3.2. Create a suitable target volume for geo-replication	22
3.3.3. Configuring geo-replication for backing up volumes	22
3.3.3.1. Creating a geo-replication session	22
3.3.3.2. Verifying creation of a geo-replication session	23
3.3.3.3. Synchronizing volume state using the Administration Portal	23
3.3.4. Scheduling regular backups using geo-replication	23
3.4. CONFIGURING FAILOVER TO AND FAILBACK FROM A SECONDARY CLUSTER	24
3.4.1. Creating a secondary cluster for failover	24
3.4.2. Creating a mapping file between source and target clusters	24
3.4.3. Creating a failover playbook between source and target clusters	25
3.4.4. Creating a failover cleanup playbook for your primary cluster	26
3.4.5. Create a failback playbook between source and target clusters	26
CHAPTER 4. CONFIGURE ENCRYPTION WITH TRANSPORT LAYER SECURITY (TLS/SSL)	27
4.1. CONFIGURING TLS/SSL USING SELF-SIGNED CERTIFICATES	27
4.2. CONFIGURING TLS/SSL USING CERTIFICATE AUTHORITY SIGNED CERTIFICATES	29
CHAPTER 5. CONFIGURE PERFORMANCE IMPROVEMENTS	33
5.1. IMPROVING VOLUME PERFORMANCE BY CHANGING SHARD SIZE	33
5.1.1. Changing shard size on replicated volumes	33
5.1.2. Changing shard size on arbitrated volumes	36
5.2. CONFIGURING A LOGICAL VOLUME CACHE (LVMCACHE) FOR AN EXISTING VOLUME	39
CHAPTER 6. CONFIGURE MONITORING	42
6.1. CONFIGURING EVENT NOTIFICATIONS	42
PART II. MAINTENANCE TASKS	43
CHAPTER 7. BASIC OPERATIONS	44
7.1. CREATING A SHUTDOWN PLAYBOOK	44
7.2. SHUTTING DOWN RHHI FOR VIRTUALIZATION	44
7.3. STARTING UP A HYPERCONVERGED CLUSTER	45

CHAPTER 8. UPGRADING TO RED HAT HYPERCONVERGED INFRASTRUCTURE FOR VIRTUALIZATION 1.6	47
8.1. MAJOR CHANGES IN VERSION 1.6	47
8.2. UPGRADE WORKFLOW	47
8.3. PREPARING TO UPGRADE	47
8.3.1. Update to the latest version of the previous release	47
8.3.2. Update subscriptions	48
8.3.3. Verify that data is not currently being synchronized using geo-replication	48
8.4. UPGRADING RED HAT HYPERCONVERGED INFRASTRUCTURE FOR VIRTUALIZATION	48
8.4.1. Upgrading the Hosted Engine virtual machine	48
8.4.2. Upgrading the hyperconverged hosts	49
CHAPTER 9. MONITORING RED HAT HYPERCONVERGED INFRASTRUCTURE FOR VIRTUALIZATION	52
9.1. MONITORING VIRTUAL DATA OPTIMIZER (VDO)	52
9.1.1. Monitoring VDO using the command line interface	52
9.1.2. Monitoring VDO using the Web Console	52
CHAPTER 10. FREEING SPACE ON THINLY-PROVISIONED LOGICAL VOLUMES USING FSTRIM	53
CHAPTER 11. ADD HYPERCONVERGED HOSTS TO RED HAT VIRTUALIZATION MANAGER	54
CHAPTER 12. REINSTALLING A HYPERCONVERGED HOST	55
CHAPTER 13. REPLACING HOSTS	56
13.1. REPLACING THE PRIMARY HYPERCONVERGED HOST USING ANSIBLE	56
13.2. REPLACING OTHER HYPERCONVERGED HOSTS USING ANSIBLE	59
13.2.1. Replacing a hyperconverged host to use a different FQDN	59
13.2.2. Replacing a hyperconverged host to use the same FQDN	62
13.3. PREPARING A REPLACEMENT HYPERCONVERGED HOST USING ANSIBLE	64
CHAPTER 14. RECOVERING FROM DISASTER	68
14.1. MANUALLY RESTORING DATA FROM A BACKUP VOLUME	68
14.1.1. Restoring a volume from a geo-replicated backup	68
14.2. FAILING OVER TO A SECONDARY CLUSTER	70
14.3. FAILING BACK TO A PRIMARY CLUSTER	70
14.4. STOPPING A GEO-REPLICATION SESSION USING RHV MANAGER	70
14.5. TURNING OFF SCHEDULED BACKUPS BY DELETING THE GEO-REPLICATION SCHEDULE	71
PART III. REFERENCE MATERIAL	72
APPENDIX A. FENCING POLICIES FOR RED HAT GLUSTER STORAGE	73
APPENDIX B. GLOSSARY OF TERMS	74
B.1. VIRTUALIZATION TERMS	74
B.2. STORAGE TERMS	74
B.3. HYPERCONVERGED INFRASTRUCTURE TERMS	75

PART I. CONFIGURATION TASKS

CHAPTER 1. ADD COMPUTE AND STORAGE RESOURCES

Red Hat Hyperconverged Infrastructure for Virtualization (RHVI for Virtualization) can be scaled to 6, 9, or 12 nodes.

You can add compute and storage resources in several ways:

- [Add new volumes on new nodes](#)
- [Expand an existing volume across new nodes](#)

You can also increase the space available on your existing nodes to expand storage without expanding compute resources.

- [Growing a thin pool using the Web Console](#)
- [Growing a logical volume using the Web Console](#)
- [Growing the logical size of a VDO device using cockpit](#)

1.1. CREATING NEW BRICKS USING ANSIBLE

If you want to create bricks on a lot of hosts at once, you can automate the process by creating an ansible playbook. Follow this process to create and run a playbook that creates, formats, and mounts bricks for use in a hyperconverged environment.

Prerequisites

- Install the physical machines to host your new bricks. Follow the instructions in [Install Physical Host Machines](#).
- Configure key-based SSH authentication without a password between all nodes. Configure this from the node that is running the Web Console to all new nodes, and from the first new node to all other new nodes.



IMPORTANT

RHVI for Virtualization expects key-based SSH authentication without a password between these nodes for both IP addresses and FQDNs. Ensure that you configure key-based SSH authentication between these machines for the IP address and FQDN of all storage and management network interfaces.

Follow the instructions in [Using key-based authentication](#) to configure key-based SSH authentication without a password.

- Verify that your hosts do not use a Virtual Disk Optimization (VDO) layer. If you have a VDO layer, use [Section 1.2, "Creating new bricks above VDO layer using ansible"](#) instead.

Procedure

1. Create an inventory file

Create a new **inventory** file in the `/etc/ansible/roles/gluster.infra/playbooks` directory using the following example.

This file lists the hosts on which to create new bricks.

Example inventory file

```
[hosts]
server4.example.com
server5.example.com
server6.example.com
```

2. Create a `bricks.yml` variables file

Create a new `bricks.yml` file in the `/etc/ansible/roles/gluster.infra/playbooks` directory using the following example.

This file defines the underlying storage infrastructure and settings to be created or used on each host.

Example `bricks.yml` variable file

```
# gluster_infra_disktype
# Set a disk type. Options: JBOD, RAID6, RAID10 - Default: JBOD
gluster_infra_disktype: RAID10

# gluster_infra_dalign
# Dataalignment, for JBOD default is 256K if not provided.
# For RAID{6,10} dataalignment is computed by multiplying
# gluster_infra_diskcount and gluster_infra_stripe_unit_size.
gluster_infra_dalign: 256K

# gluster_infra_diskcount
# Required only for RAID6 and RAID10.
gluster_infra_diskcount: 10

# gluster_infra_stripe_unit_size
# Required only in case of RAID6 and RAID10. Stripe unit size always in KiB, do
# not provide the trailing `K' in the value.
gluster_infra_stripe_unit_size: 128

# gluster_infra_volume_groups
# Variables for creating volume group
gluster_infra_volume_groups:
  - { vgname: 'vg_vdb', pvname: '/dev/vdb' }
  - { vgname: 'vg_vdc', pvname: '/dev/vdc' }

# gluster_infra_thick_lvs
# Variable for thick lv creation
gluster_infra_thick_lvs:
  - { vgname: 'vg_vdb', lvname: 'vg_vdb_thicklv1', size: '10G' }


# gluster_infra_thinpools
# thinpoolname is optional, if not provided `vgname' followed by _thinpool is
# used for name. poolmetadatasize is optional, default 16G is used
gluster_infra_thinpools:
  - {vgname: 'vg_vdb', thinpoolname: 'foo_thinpool', thinpoolsizesize: '10G', poolmetadatasize:
'1G' }
  - {vgname: 'vg_vdc', thinpoolname: 'bar_thinpool', thinpoolsizesize: '20G', poolmetadatasize:
```

```
'1G' }

# gluster_infra_lv_logicalvols
# Thin volumes for the brick. `thinpoolname' is optional, if omitted `vgname'
# followed by _thinpool is used
gluster_infra_lv_logicalvols:
  - { vgname: 'vg_vdb', thinpool: 'foo_thinpool', lvname: 'vg_vdb_thinlv', lvsize: '500G' }
  - { vgname: 'vg_vdc', thinpool: 'bar_thinpool', lvname: 'vg_vdc_thinlv', lvsize: '500G' }

# Setting up cache using SSD disks
gluster_infra_cache_vars:
  - { vgname: 'vg_vdb', cachedisk: '/dev/vdd',
      cachethinpoolname: 'foo_thinpool', cachelvname: 'cachelv',
      cachelvsize: '20G', cachemetalvname: 'cachemeta',
      cachemetalvsize: '100M', cachemode: 'writethrough' }

# gluster_infra_mount_devices
gluster_infra_mount_devices:
  - { path: '/rhgs/thicklv', vgname: 'vg_vdb', lvname: 'vg_vdb_thicklv1' }
  - { path: '/rhgs/thinlv1', vgname: 'vg_vdb', lvname: 'vg_vdb_thinlv' }
  - { path: '/rhgs/thinlv2', vgname: 'vg_vdc', lvname: 'vg_vdc_thinlv' }
```



IMPORTANT

If the **path:** defined does not begin with **/rhgs** the bricks are not detected automatically by the Administration Portal. Synchronize the host storage after running the **create_brick.yml** playbook to add the new bricks to the Administration Portal.

3. Create a **create_brick.yml** playbook file

Create a new **create_brick.yml** file in the **/etc/ansible/roles/gluster.infra/playbooks** directory using the following example.

This file defines the work involved in creating a brick using the **gluster.infra** role and the variable file you created above.

Example **create_brick.yml** playbook file

```
---
- name: Create a GlusterFS brick on the servers
  remote_user: root
  hosts: all
  gather_facts: false
  vars_files:
    - bricks.yml

  roles:
    - gluster.infra
```

4. Execute the playbook

Run the following command from the **/etc/ansible/roles/gluster.infra/playbooks** directory to run the playbook you created using the inventory and the variables files you defined above.

```
# ansible-playbook -i inventory create_brick.yml
```

5. Verify that your bricks are available

- a. Click **Compute** → **Hosts** and select the host.
- b. Click **Storage Devices** and check the list of storage devices for your new bricks. If you cannot see your new bricks, click **Sync** and wait for them to appear in the list of storage devices.

1.2. CREATING NEW BRICKS ABOVE VDO LAYER USING ANSIBLE

If you want to create bricks on a lot of hosts at once, you can automate the process by creating an ansible playbook.

Prerequisites

- Install the physical machines to host your new bricks. Follow the instructions in [Install Physical Host Machines](#).
- Configure key-based SSH authentication without a password between all nodes. Configure this from the node that is running the Web Console to all new nodes, and from the first new node to all other new nodes.



IMPORTANT

RHHI for Virtualization expects key-based SSH authentication without a password between these nodes for both IP addresses and FQDNs. Ensure that you configure key-based SSH authentication between these machines for the IP address and FQDN of all storage and management network interfaces.

Follow the instructions in [Using key-based authentication](#) to configure key-based SSH authentication without a password.

- Verify that your hosts use a Virtual Disk Optimization (VDO) layer. If you do not have a VDO layer, use [Section 1.1, "Creating new bricks using ansible"](#) instead.

Procedure

1. Create an inventory file

Create a new **inventory** file in the `/etc/ansible/roles/gluster.infra/playbooks` directory using the following example.

This file lists the hosts on which to create new bricks.

Example inventory file

```
[hosts]
server4.example.com
server5.example.com
server6.example.com
```

2. Create a bricks.yml variables file

Create a new **bricks.yml** file in the `/etc/ansible/roles/gluster.infra/playbooks` directory using the following example.

This file defines the underlying storage infrastructure and settings to be created or used on each host.

Example `vdo_bricks.yml` variable file

```
# gluster_infra_disktype
# Set a disk type. Options: JBOD, RAID6, RAID10 - Default: JBOD
gluster_infra_disktype: RAID10

# gluster_infra_dalign
# Dataalignment, for JBOD default is 256K if not provided.
# For RAID{6,10} dataalignment is computed by multiplying
# gluster_infra_diskcount and gluster_infra_stripe_unit_size.
gluster_infra_dalign: 256K

# gluster_infra_diskcount
# Required only for RAID6 and RAID10.
gluster_infra_diskcount: 10

# gluster_infra_stripe_unit_size
# Required only in case of RAID6 and RAID10. Stripe unit size always in KiB, do
# not provide the trailing `K' in the value.
gluster_infra_stripe_unit_size: 128

# VDO creation
gluster_infra_vdo:
  - { name: 'hc_vdo_1', device: '/dev/vdb' }
  - { name: 'hc_vdo_2', device: '/dev/vdc' }


# gluster_infra_volume_groups
# Variables for creating volume group
gluster_infra_volume_groups:
  - { vgname: 'vg_vdb', pvname: '/dev/mapper/hc_vdo_1' }
  - { vgname: 'vg_vdc', pvname: '/dev/mapper/hc_vdo_2' }

# gluster_infra_thick_lvs
# Variable for thick lv creation
gluster_infra_thick_lvs:
  - { vgname: 'vg_vdb', lvname: 'vg_vdb_thicklv1', size: '10G' }

# gluster_infra_thinpools
# thinpoolname is optional, if not provided `vgname' followed by _thinpool is
# used for name. poolmetadatasize is optional, default 16G is used
gluster_infra_thinpools:
  - {vgname: 'vg_vdb', thinpoolname: 'foo_thinpool', thinpoolsizesize: '10G', poolmetadatasize:
'1G' }
  - {vgname: 'vg_vdc', thinpoolname: 'bar_thinpool', thinpoolsizesize: '20G', poolmetadatasize:
'1G' }

# gluster_infra_lv_logicalvols
# Thinvolumes for the brick. `thinpoolname' is optional, if omitted `vgname'
# followed by _thinpool is used
gluster_infra_lv_logicalvols:
  - { vgname: 'vg_vdb', thinpool: 'foo_thinpool', lvname: 'vg_vdb_thinlv', lvsize: '500G' }
  - { vgname: 'vg_vdc', thinpool: 'bar_thinpool', lvname: 'vg_vdc_thinlv', lvsize: '500G' }
```

```
# gluster_infra_mount_devices
gluster_infra_mount_devices:
  - { path: '/rhgs/thickl', vgname: 'vg_vdb', lvname: 'vg_vdb_thickl' }
  - { path: '/rhgs/thinlv1', vgname: 'vg_vdb', lvname: 'vg_vdb_thinlv1' }
  - { path: '/rhgs/thinlv2', vgname: 'vg_vdc', lvname: 'vg_vdc_thinlv2' }
```



IMPORTANT

If the **path:** defined does not begin with **/rhgs** the bricks are not detected automatically by the Administration Portal. Synchronize the host storage after running the **create_brick.yml** playbook to add the new bricks to the Administration Portal.

3. Create a **create_brick.yml** playbook file

Create a new **create_brick.yml** file in the **/etc/ansible/roles/gluster.infra/playbooks** directory using the following example.

This file defines the work involved in creating a brick using the **gluster.infra** role and the variable file you created above.

Example **create_brick.yml** playbook file

```
---
- name: Create a GlusterFS brick on the servers
  remote_user: root
  hosts: all
  gather_facts: false
  vars_files:
    - vdo_bricks.yml

  roles:
    - gluster.infra
```

4. Execute the playbook

Run the following command from the **/etc/ansible/roles/gluster.infra/playbooks** directory to run the playbook you created using the inventory and the variables files you defined above.

```
# ansible-playbook -i inventory create_brick.yml
```

5. Verify that your bricks are available

- a. Click **Compute** → **Hosts** and select the host.
- b. Click **Storage Devices** and check the list of storage devices for your new bricks. If you cannot see your new bricks, click **Sync** and wait for them to appear in the list of storage devices.

1.3. EXPANDING VOLUME FROM RED HAT VIRTUALIZATION MANAGER

Follow this section to expand an existing volume across new bricks on new hyperconverged nodes.

Prerequisites

- Verify that your scaling plans are supported: [Requirements for scaling](#).
- If your existing deployment uses certificates signed by a Certificate Authority for encryption, prepare the certificates required for the new nodes.
- Install three physical machines to serve as the new hyperconverged nodes. Follow the instructions in [Install Physical Host Machines](#).
- Configure key-based SSH authentication without a password. Configure this from the node that is running the Web Console to all new nodes, and from the first new node to all other new nodes.



IMPORTANT

RHHI for Virtualization expects key-based SSH authentication without a password between these nodes for both IP addresses and FQDNs. Ensure that you configure key-based SSH authentication between these machines for the IP address and FQDN of all storage and management network interfaces.

Follow the instructions in [Using key-based authentication](#) to configure key-based SSH authentication without a password.

Procedure

1. Create new bricks

Create the bricks on the servers you want to expand your volume across by following the instructions in [Creating bricks using ansible](#) or [Creating bricks above a VDO layer using ansible](#) depending on your requirements.



IMPORTANT

If the **path:** defined does not begin with **/rhgs** the bricks are not detected automatically by the Administration Portal. Synchronize the host storage after running the **create_brick.yml** playbook to synchronize the new bricks to the Administration Portal.

1. Click **Compute** → **Hosts** and select the host.
2. Click **Storage Devices**
3. Click **Sync**.

Repeat for each host that has new bricks.

2. Add new bricks to the volume

- a. Log in to RHV Administration Console.
- b. Click **Storage** → **Volumes** and select the volume to expand.
- c. Click the **Bricks** tab.
- d. Click **Add**. The *Add Bricks* window opens.
- e. Add new bricks.

- i. Select the brick host from the **Host** dropdown menu.
 - ii. Select the brick to add from the **Brick Directory** dropdown menu and click **Add**.
- f. When all bricks are listed, click **OK** to add bricks to the volume.


The volume automatically syncs the new bricks.

1.4. EXPANDING THE HYPERCONVERGED CLUSTER BY ADDING A NEW VOLUME ON NEW NODES USING THE WEB CONSOLE

Follow these instructions to use the Web Console to expand your hyperconverged cluster with a new volume on new nodes.

Prerequisites

- Verify that your scaling plans are supported: [Requirements for scaling](#).
- If your existing deployment uses certificates signed by a Certificate Authority for encryption, prepare the certificates that will be required for the new nodes.
- Install three physical machines to serve as the new hyperconverged nodes. Follow the instructions in [Deploying Red Hat Hyperconverged Infrastructure for Virtualization](#).
- Configure key-based SSH authentication without a password. Configure this from the node that is running the Web Console to all new nodes, and from the first new node to all other new nodes.



IMPORTANT

RHHI for Virtualization expects key-based SSH authentication without a password between these nodes for both IP addresses and FQDNs. Ensure that you configure key-based SSH authentication between these machines for the IP address and FQDN of all storage and management network interfaces.

Follow the instructions in [Using key-based authentication](#) to configure key-based SSH authentication without a password.

Procedure

1. Log in to the Web Console.
2. Click **Virtualization** → **Hosted Engine** and then click **Manage Gluster**.
3. Click **Expand Cluster**. The *Gluster Deployment* window opens.
 - a. On the *Hosts* tab, enter the FQDN or IP address of the new hyperconverged nodes and click **Next**.

Expand Cluster ✕

Hosts Volumes Bricks Review

1 ————— 2 ————— 3 ————— 4

Host1	<input type="text" value="newhost1.example.com"/>
Host2	<input type="text" value="newhost2.example.com"/>
Host3 ⓘ	<input type="text" value="newhost3.example.com"/>

- b. On the *Volumes* tab, specify the details of the volume you want to create.

Expand Cluster ✕

Hosts Volumes Bricks Review

1 ————— 2 ————— 3 ————— 4

Name	Volume Type	Arbiter	Brick Dirs
<input type="text" value="new_volume"/>	<input style="border: 1px solid #ccc; background-color: #f0f0f0; border-radius: 3px; padding: 2px 5px; font-size: 0.9em; font-family: sans-serif; margin: 0;" type="text" value="Replicate"/> ▾	<input type="checkbox"/>	<input type="text" value="/gluster_bricks/new_volume/nc"/> <input type="button" value="🗑"/>

[⊕ Add Volume](#)

- c. On the *Bricks* tab, specify the details of the disks to be used to create the Gluster volume.

Expand Cluster
✕

Hosts Volumes Bricks Review

① ————— ② ————— ③ ————— ④

Raid Information ⓘ

Raid Type: RAID 6 ▾

Stripe Size(KB): 256 ▾

Data Disk Count: 12 ▾

Brick Configuration

Select Host: newhost1.example.com ▾

LV Name	Device Name	Size(GB)	Thinp	Mount Point	Enable Dedupe & Compression
new_volume	sdb	500	<input checked="" type="checkbox"/>	/gluster_bricks/new_volume	<input type="checkbox"/> 🗑

[+ Add Bricks](#)

Configure LV Cache

ⓘ **Arbiter bricks will be created on the third host in the host list.**

Cancel < Back Next >

- d. On the *Review* tab, check the generated file for any problems. When you are satisfied, click **Deploy**.

Expand Cluster
✕

Hosts Volumes Bricks Review

① ————— ② ————— ③ ————— ④

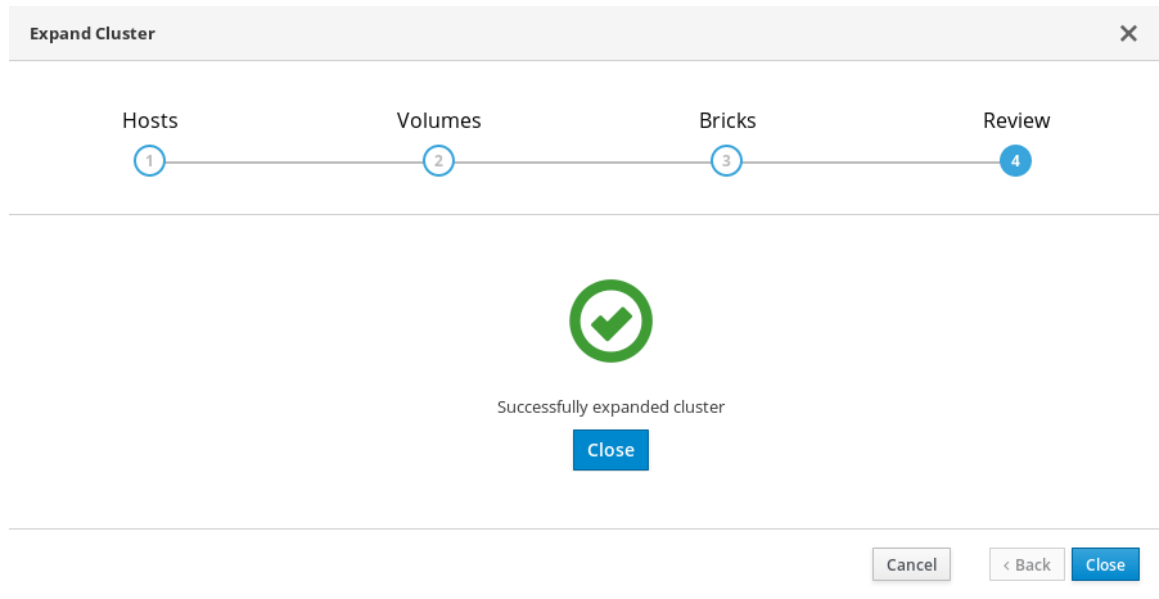
Generated Ansible inventory : /etc/ansible/hc_wizard_inventory.yml
✎ Edit 🔄 Reload

```

hc_nodes:
hosts:
newhost1.example.com:
gluster_infra_volume_groups:
- vgname: gluster_vg_sdb
  pvname: /dev/sdb
gluster_infra_mount_devices:
- path: /gluster_bricks/new_volume
  lvname: gluster_lv_new_volume
  vgname: gluster_vg_sdb
gluster_infra_thinpools:
- vgname: gluster_vg_sdb
  thinpoolname: gluster_thinpool_gluster_vg_sdb
                    
```

Cancel < Back Deploy

Deployment takes some time to complete. The following screen appears when the cluster has been successfully expanded.



CHAPTER 2. CONFIGURE HIGH AVAILABILITY USING FENCING POLICIES

Fencing allows a cluster to enforce performance and availability policies and react to unexpected host failures by automatically rebooting hyperconverged hosts.

Several policies specific to Red Hat Gluster Storage must be enabled to ensure that fencing activities do not disrupt storage services in a Red Hat Hyperconverged (RHHI for Virtualization) Infrastructure deployment.

This requires enabling and configuring fencing at both the cluster level and at the host level. See the following sections for details.

2.1. CONFIGURING FENCING POLICIES IN THE CLUSTER

1. In the Administration Portal, click **Compute** → **Clusters**.
2. Select the cluster and click **Edit**. The *Edit Cluster* window opens.
3. Click the **Fencing policy** tab.
4. Check the **Enable fencing** checkbox.
5. Check the checkboxes for at least the following fencing policies:
 - Skip fencing if gluster bricks are up
 - Skip fencing if gluster quorum not met

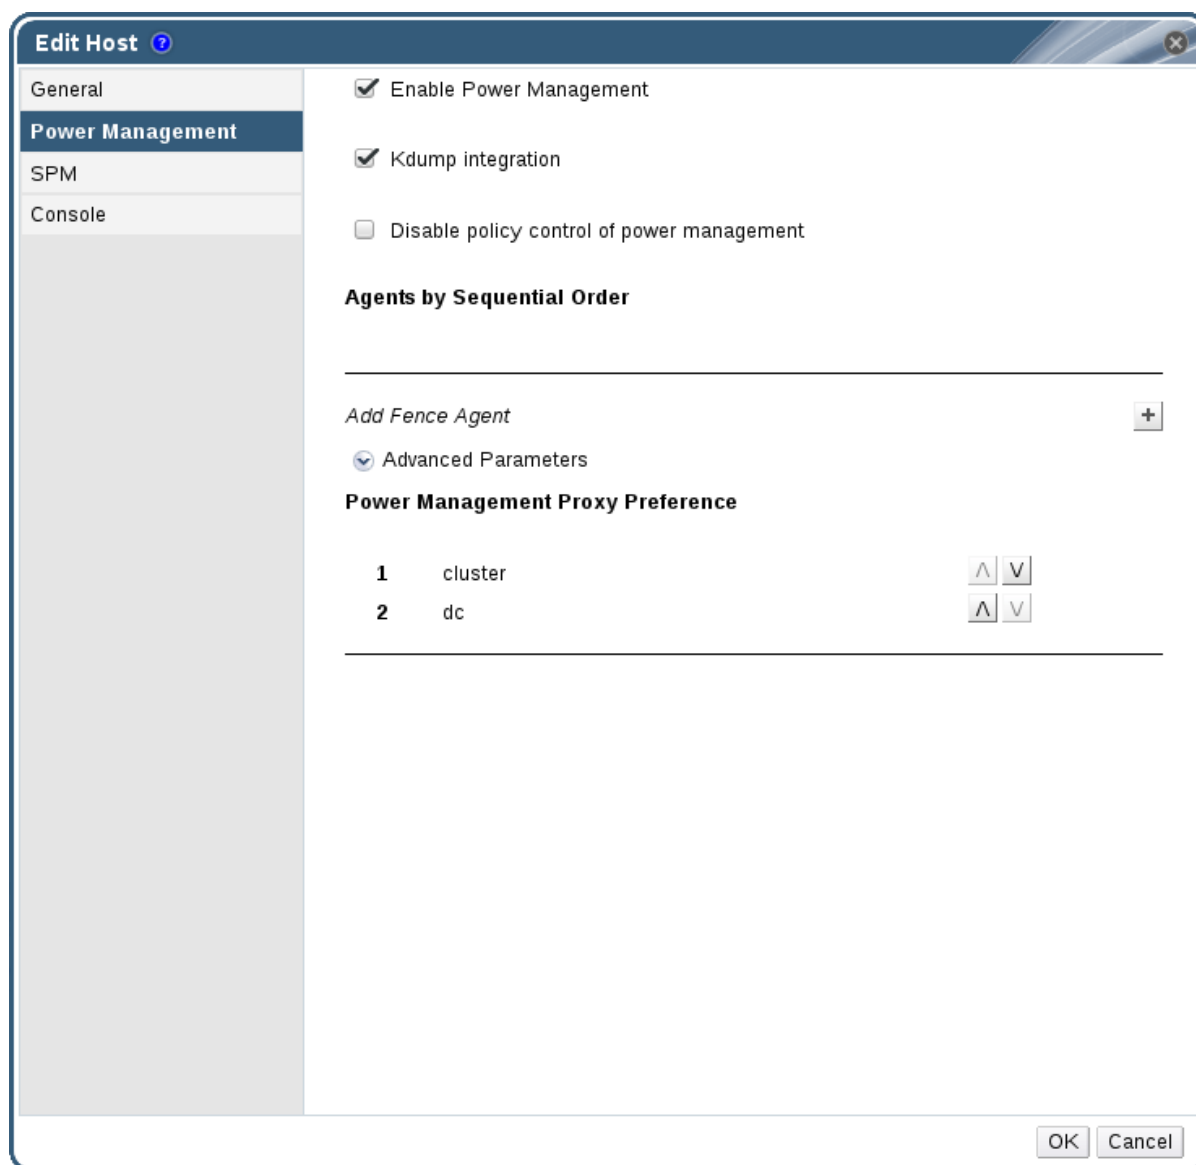
See [Appendix A, *Fencing Policies for Red Hat Gluster Storage*](#) for details on the effects of these policies.

6. Click **OK** to save settings.

2.2. CONFIGURING FENCING PARAMETERS ON THE HOSTS

1. In the Administration Portal, click **Compute** → **Hosts**.
2. Select the host to configure, and click **Edit** to open the **Edit Host** window.
3. Click the **Power Management** tab.

Figure 2.1. Power Management Settings



4. Check the **Enable Power Management** check box. This enables other fields on the tab.
5. Check the **Kdump integration** check box to prevent the host from fencing while performing a kernel crash dump.



IMPORTANT

When you enable Kdump integration on an existing host, the host must be reinstalled for kdump to be configured. See [Chapter 12, Reinstalling a hyperconverged host](#) for instructions on reinstalling a host.

1. Click the plus (+) button to add a new power management device. The **Edit fence agent** window opens.

Figure 2.2. Edit fence agent

Edit fence agent

Address

User Name

Password

Type

SSH Port

Slot

Options

Please use a comma-separated list of 'key=value'

Secure

- a. Enter the **Address**, **User Name**, and **Password** of the power management device.
- b. Select the power management device **Type** from the drop-down list.
- a. Enter the **SSH Port** number used by the power management device to communicate with the host.
- b. Enter the **Slot** number used to identify the blade of the power management device.
- c. Enter the **Options** for the power management device. Use a comma-separated list of *key=value* entries.
- d. Check the **Secure** check box to enable the power management device to connect securely to the host.
- e. Click the **Test** button to ensure the settings are correct. *Test Succeeded, Host Status is: on* displays upon successful verification.

**WARNING**

Power management parameters (userid, password, options, etc.) are tested by Red Hat Virtualization Manager in two situations: during setup, and when parameter values are manually changed in Red Hat Virtualization Manager. If you choose to ignore alerts about incorrect parameters, or if the parameters are changed on the power management hardware without the corresponding change in Red Hat Virtualization Manager, fencing is likely to fail.

- f. Click **OK** to finish adding the fence agent.
1. Click **OK** to save your host configuration.

You are returned to the list of hosts. Note that the exclamation mark (!) next to the host's name has now disappeared, signifying that power management has been successfully configured.

CHAPTER 3. CONFIGURING BACKUP AND RECOVERY OPTIONS

This chapter explains how to add disaster recovery capabilities to your Red Hat Hyperconverged Infrastructure for Virtualization deployment so that you can restore your cluster to a working state after a disk or server failure.

3.1. PREREQUISITES

3.1.1. Prerequisites for geo-replication

Be aware of the following requirements and limitations when configuring geo-replication:

One geo-replicated volume only

Red Hat Hyperconverged Infrastructure for Virtualization (RHHI for Virtualization) supports only one geo-replicated volume. Red Hat recommends backing up the volume that stores the data of your virtual machines, as this is usually contains the most valuable data.

Two different managers required

The source and destination volumes for geo-replication must be managed by different instances of Red Hat Virtualization Manager.

3.1.2. Prerequisites for failover and failback configuration

Versions must match between environments

Ensure that the primary and secondary environments have the same version of Red Hat Virtualization Manager, with identical data center compatibility versions, cluster compatibility versions, and PostgreSQL versions.

No virtual machine disks in the hosted engine storage domain

The storage domain used by the hosted engine virtual machine is not failed over, so any virtual machine disks in this storage domain will be lost.

Execute Ansible playbooks manually from a separate master node

Generate and execute Ansible playbooks manually from a separate machine that acts as an Ansible master node.

3.2. SUPPORTED BACKUP AND RECOVERY CONFIGURATIONS

There are two supported ways to add disaster recovery capabilities to your Red Hat Hyperconverged Infrastructure for Virtualization deployment.

Configure backing up to a secondary volume only

Regularly synchronizing your data to a remote secondary volume helps to ensure that your data is not lost in the event of disk or server failure.

This option is suitable if the following statements are true of your deployment.

- You require only a backup of your data for disaster recovery.
- You do not require highly available storage.
- You do not want to maintain a secondary cluster.

- You are willing to manually restore your data and reconfigure your backup solution after a failure has occurred.

Follow the instructions in [Configuring backup to a secondary volume](#) to configure this option.

Configure failing over to and failing back from a secondary cluster

This option provides failover and failback capabilities in addition to backing up data on a remote volume. Configuring failover of your primary cluster's operations and storage domains to a secondary cluster helps to ensure that your data remains available in event of disk or server failure in the primary cluster.

This option is suitable if the following statements are true of your deployment.

- You require highly available storage.
- You are willing to maintain a secondary cluster.
- You do not want to manually restore your data or reconfigure your backup solution after a failure has occurred.

Follow the instructions in [Configuring failover to and failback from a secondary cluster](#) to configure this option.

Red Hat recommends that you configure at least a backup volume for production deployments.

3.3. CONFIGURING BACKUP TO A SECONDARY VOLUME

This section covers how to back up a gluster volume to a secondary gluster volume using geo-replication.

To do this, you must:

1. Ensure that all [prerequisites](#) are met.
2. [Create a suitable volume to use as a geo-replication target](#) .
3. [Configure a geo-replication session](#) between the source volume and the target volume.
4. [Schedule](#) the geo-replication process.

3.3.1. Prerequisites

3.3.1.1. Enable shared storage on the source volume

Ensure that the volume you want to back up (the source volume) has shared storage enabled. Run the following command on any server that hosts the source volume to enable shared storage.

```
# gluster volume set all cluster.enable-shared-storage enable
```

Ensure that a gluster volume named **gluster_shared_storage** is created in the source cluster, and is mounted at **/var/run/gluster/shared_storage** on all the nodes in the source cluster. See [Setting Up Shared Storage](#) for further information.

3.3.1.2. Match encryption on source and target volumes

If encryption is enabled on the volume that you want to back up, encryption must also be enabled on the volume that will hold your backed up data.

See [Configure Encryption with Transport Layer Security \(TLS/SSL\)](#) for details.

3.3.2. Create a suitable target volume for geo-replication

Prepare a secondary gluster volume to hold the geo-replicated copy of your source volume. This target volume should be in a separate cluster, hosted at a separate site, so that the risk of source and target volumes being affected by the same outages is minimised.


Ensure that the target volume for geo-replication has sharding enabled. Run the following command on any node that hosts the target volume to enable sharding on that volume.

```
# gluster volume set <volname> features.shard enable
```

3.3.3. Configuring geo-replication for backing up volumes

3.3.3.1. Creating a geo-replication session

A geo-replication session is required to replicate data from an active source volume to a passive target volume.



IMPORTANT

Only rsync based geo-replication is supported with Red Hat Hyperconverged Infrastructure for Virtualization.

1. Create a common **pem pub** file.

Run the following command on a source node that has key-based SSH authentication without a password configured to the target nodes.

```
# gluster system:: execute gsec_create
```

2. Create the geo-replication session

Run the following command to create a geo-replication session between the source and target volumes, using the created **pem pub** file for authentication.

```
# gluster volume geo-replication <SOURCE_VOL> <TARGET_NODE>::<TARGET_VOL>
create push-pem
```

For example, the following command creates a geo-replication session from a source volume **prodvol** to a target volume called **backupvol**, which is hosted by **backup.example.com**.

```
# gluster volume geo-replication prodvol backup.example.com::backupvol create push-pem
```

By default this command verifies that the target volume is a valid target with available space. You can append the **force** option to the command to ignore failed verification.

3. Configure a meta-volume

This relies on the source volume having shared storage configured, as described in [Prerequisites](#).

```
# gluster volume geo-replication <SOURCE_VOL> <TARGET_HOST>:::<TARGET_VOL>
config use_meta_volume true
```



IMPORTANT

Do not start the geo-replication session. Starting the geo-replication session begins replication from your source volume to your target volume.

3.3.3.2. Verifying creation of a geo-replication session

1. Log in to the Administration Portal on any source node.
2. Click **Storage** → **Volumes**.
3. Check the **Info** column for the geo-replication icon.
If this icon is present, geo-replication has been configured for that volume.

If this icon is not present, try [synchronizing the volume](#).

3.3.3.3. Synchronizing volume state using the Administration Portal

1. Log in to the Administration Portal.
2. Click **Compute** → **Volumes**.
3. Select the volume that you want to synchronize.
4. Click the **Geo-replication** sub-tab.
5. Click **Sync**.

3.3.4. Scheduling regular backups using geo-replication

1. Log in to the Administration Portal on any source node.
2. Click **Storage** → **Domains**.
3. Click the name of the storage domain that you want to back up.
4. Click the **Remote Data Sync Setup** subtab.
5. Click **Setup**.
The *Setup Remote Data Synchronization* window opens.
 - a. In the **Geo-replicated to** field, select the backup target.
 - b. In the **Recurrence** field, select a recurrence interval type.
Valid values are **WEEKLY** with at least one weekday checkbox selected, or **DAILY**.
 - c. In the **Hours** and **Minutes** field, specify the time to start synchronizing.



NOTE

This time is based on the Hosted Engine's timezone.

- d. Click **OK**.
6. Check the **Events** subtab for the source volume at the time you specified to verify that synchronization works correctly.

3.4. CONFIGURING FAILOVER TO AND FAILBACK FROM A SECONDARY CLUSTER

This section covers how to configure your cluster to fail over to a remote secondary cluster in the event of server failure.

To do this, you must:

1. [Configure backing up to a remote volume](#) .
2. [Create a suitable cluster to use as a failover target](#) .
3. [Prepare a mapping file](#) for the source and target clusters.
4. [Prepare a failover playbook](#) .
5. [Prepare a cleanup playbook](#) for the primary cluster.
6. [Prepare a failback playbook](#) .

3.4.1. Creating a secondary cluster for failover

Install and configure a secondary cluster that can be used in place of the primary cluster in the event of failure.

This secondary cluster can be either of the following configurations:

Red Hat Hyperconverged Infrastructure

See [Deploying Red Hat Hyperconverged Infrastructure](#) for details.

Red Hat Gluster Storage configured for use as a Red Hat Virtualization storage domain

See [Configuring Red Hat Virtualization with Red Hat Gluster Storage](#) for details. Note that creating a storage domain is not necessary for this use case; the storage domain is imported as part of the failover process.

The storage on the secondary cluster must not be attached to a data center, so that it can be added to the secondary site's data center during the failover process.

3.4.2. Creating a mapping file between source and target clusters

Follow this section to create a file that maps the storage in your source cluster to the storage in your target cluster.

Red Hat recommends that you create this file immediately after you first deploy your storage, and keep it up to date as your deployment changes. This helps to ensure that everything in your cluster fails over safely in the event of disaster.

1. Create a playbook to generate the mapping file.
Create a playbook that passes information about your cluster to the **oVirt.disaster-recovery** role, using the **site**, **username**, **password**, and **ca** variables.

Red Hat recommends creating this file in the `/usr/share/ansible/roles/oVirt.disaster-recovery` directory of the server that provides **ansible** and manages failover and failback.

Example playbook file: dr-ovirt-setup.yml

```
---
- name: Collect mapping variables
  hosts: localhost
  connection: local

  vars:
    site: https://example.engine.redhat.com/ovirt-engine/api
    username: admin@internal
    password: my_password
    ca: /etc/pki/ovirt-engine/ca.pem
    var_file: disaster_recovery_vars.yml

  roles:
    - oVirt.disaster-recovery
```

2. Generate the mapping file by running the playbook with the **generate_mapping** tag.

```
# ansible-playbook dr-ovirt-setup.yml --tags "generate_mapping"
```

This creates the mapping file, **disaster_recovery_vars.yml**.

3. Edit **disaster_recovery_vars.yml** and add information about the secondary cluster.
See [Appendix A: Mapping File Attributes](#) in the Red Hat Virtualization *Disaster Recovery Guide* for detailed information about attributes used in the mapping file.

3.4.3. Creating a failover playbook between source and target clusters

Create a playbook file that passes the lists of hyperconverged hosts to use as a failover source and target to the **oVirt.disaster-recovery** role, using the **dr_target_host** and **dr_source_map** variables.

Red Hat recommends creating this file in the `/usr/share/ansible/roles/oVirt.disaster-recovery` directory of the server that provides **ansible** and manages failover and failback.

Example playbook file: dr-rhv-failover.yml

```
---
- name: Failover RHV
  hosts: localhost
  connection: local
  vars:
    dr_target_host: secondary
    dr_source_map: primary
  vars_files:
    - disaster_recovery_vars.yml
    - passwords.yml
  roles:
    - oVirt.disaster-recovery
```

For information about executing failover, see [Failing over to a secondary cluster](#).

3.4.4. Creating a failover cleanup playbook for your primary cluster

Create a playbook file that cleans up your primary cluster so that you can use it as a failback target.

Red Hat recommends creating this file in the `/usr/share/ansible/roles/oVirt.disaster-recovery` directory of the server that provides **ansible** and manages failover and failback.

Example playbook file: `dr-cleanup.yml`

```
---
- name: Clean RHV
  hosts: localhost
  connection: local
  vars:
    dr_source_map: primary
  vars_files:
    - disaster_recovery_vars.yml
  roles:
    - oVirt.disaster-recovery
```

For information about executing failback, see [Failing back to a primary cluster](#).

3.4.5. Create a failback playbook between source and target clusters

Create a playbook file that passes the lists of hyperconverged hosts to use as a failback source and target to the **oVirt.disaster-recovery** role, using the `dr_target_host` and `dr_source_map` variables.

Red Hat recommends creating this file in the `/usr/share/ansible/roles/oVirt.disaster-recovery` directory of the server that provides **ansible** and manages failover and failback.

Example playbook file: `dr-rhv-failback.yml`

```
---
- name: Failback RHV
  hosts: localhost
  connection: local
  vars:
    dr_target_host: primary
    dr_source_map: secondary
  vars_files:
    - disaster_recovery_vars.yml
    - passwords.yml
  roles:
    - oVirt.disaster-recovery
```

For information about executing failback, see [Failing back to a primary cluster](#).

-
- 6. Unmount the hosted engine storage domain from all hyperconverged hosts

```
# hosted-engine --disconnect-storage
```

- 7. Verify that all volumes are unmounted
On each hyperconverged host, verify that all gluster volumes are no longer mounted.

```
# mount
```

- 8. Prepare self-signed certificates
Follow *Procedure 23.1. Preparing a self-signed certificate* in the Red Hat Gluster Storage Administration Guide: [Preparing Certificates](#).

- 9. Stop all volumes

```
# gluster v stop <VolumeName>
```

- 10. Restart glusterd on all nodes

```
# systemctl restart glusterd
```

- 11. Enable TLS/SSL encryption on all volumes

```
# gluster volume set <volname> client.ssl on  
# gluster volume set <volname> server.ssl on
```

- 12. Specify access permissions on all hosts

```
# gluster volume set <volname> auth.ssl-allow "host1,host2,host3"
```

- 13. Start all volumes

```
# gluster v start <VolumeName>
```

- 14. Verify that no TLS/SSL errors occurred
Check the `/var/log/glusterfs/glusterd.log` file on each physical machine to ensure that no TLS/SSL related errors occurred, and setup completed successfully.

- 15. Start all high availability services
Run the following commands on all hyperconverged hosts:

```
# systemctl start ovirt-ha-agent  
# systemctl start ovirt-ha-broker
```

- 16. Move the hosted engine out of Global Maintenance mode

```
# hosted-engine --set-maintenance --mode=none
```

The hosted engine starts automatically after a short wait.

- 17. Wait for nodes to synchronize

Run the following command on the first hyperconverged host to check synchronization status. If engine status is listed as **unknown stale-data**, synchronization requires several more minutes to complete.

The following output indicates completed synchronization.

```
# hosted-engine --vm-status | grep 'Engine status'
Engine status : {"health": "good", "vm": "up", "detail": "up"}
Engine status : {"reason": "vm not running on this host",
  "health": "bad", "vm": "down", "detail": "unknown"}
Engine status : {"reason": "vm not running on this host",
  "health": "bad", "vm": "down", "detail": "unknown"}
```

18. Activate all storage domains

Activate the master storage domain first, followed by all other storage domains.

For details on activating storage domains, see *Activating Storage Domains from Maintenance Mode* in the Red Hat Virtualization documentation:

https://access.redhat.com/documentation/en-us/red_hat_virtualization/4.3/html/administration_guide/sect-storage_tasks.

19. Start all virtual machines

See *Starting a Virtual Machine* in the Red Hat Virtualization documentation for details:

https://access.redhat.com/documentation/en-us/red_hat_virtualization/4.3/html/virtual_machine_management_guide/sect-starting_the_virtual_machine.

4.2. CONFIGURING TLS/SSL USING CERTIFICATE AUTHORITY SIGNED CERTIFICATES



IMPORTANT

Enabling or disabling encryption is a disruptive process that requires virtual machines and the Hosted Engine to be shut down.



IMPORTANT

Ensure that you have appropriate certificates signed by a Certificate Authority before proceeding. Obtaining certificates is outside the scope of this document, but further details are available in the Red Hat Gluster Storage *Administration Guide*:

https://access.redhat.com/documentation/en-us/red_hat_gluster_storage/3.4/html/administration_guide/chap-network_encryption#chap-Network_Encryption-Prereqs.

1. Shut down all virtual machines

See *Shutting Down a Virtual Machine* in the Red Hat Virtualization documentation for details:

https://access.redhat.com/documentation/en-us/red_hat_virtualization/4.3/html/virtual_machine_management_guide/chap-administrative_tasks.

2. Move all storage domains **except the hosted engine storage domain** into Maintenance mode

See *Moving Storage Domains to Maintenance Mode* in the Red Hat Virtualization documentation for details: https://access.redhat.com/documentation/en-us/red_hat_virtualization/4.3/html/administration_guide/sect-storage_tasks.

3. Move the hosted engine into global maintenance mode
Run the following command on the hyperconverged host that hosts the hosted engine:

```
# hosted-engine --set-maintenance --mode=global
```

4. Shut down the hosted engine virtual machine
Run the following command on the hyperconverged host that hosts the hosted engine:

```
# hosted-engine --vm-shutdown
```

Verify that the hosted engine has shut down by running the following command:

```
# hosted-engine --vm-status
```

5. Stop all high availability services
Run the following command on all hyperconverged hosts:

```
# systemctl stop ovirt-ha-agent
# systemctl stop ovirt-ha-broker
```


6. Unmount the hosted engine storage domain from all hyperconverged hosts

```
# hosted-engine --disconnect-storage
```

7. Verify that all volumes are unmounted
On each hyperconverged host, verify that all gluster volumes are no longer mounted.

```
# mount
```

8. Configure Certificate Authority signed encryption



IMPORTANT

Ensure that you have appropriate certificates signed by a Certificate Authority before proceeding. Obtaining certificates is outside the scope of this document.

- a. Place certificates in the following locations on all nodes.

/etc/ssl/glusterfs.key

The node's private key.

/etc/ssl/glusterfs.pem

The certificate signed by the Certificate Authority, which becomes the node's certificate.

/etc/ssl/glusterfs.ca

The Certificate Authority's certificate.

- b. Stop all volumes



```
# gluster v stop <VolumeName>
```

- c. Restart glusterd on all nodes

```
# systemctl restart glusterd
```

- d. Enable TLS/SSL encryption on all volumes

```
# gluster volume set <volname> client.ssl on
# gluster volume set <volname> server.ssl on
```

- e. Specify access permissions on all hosts

```
# gluster volume set <volname> auth.ssl-allow "host1,host2,host3"
```

- f. Start all volumes

```
# gluster v start <VolumeName>
```

9. Verify that no TLS/SSL errors occurred

Check the `/var/log/glusterfs/glusterd.log` file on each physical machine to ensure that no TLS/SSL related errors occurred, and setup completed successfully.

10. Start all high availability services

Run the following commands on all hyperconverged hosts:

```
# systemctl start ovirt-ha-agent
# systemctl start ovirt-ha-broker
```

11. Move the hosted engine out of Global Maintenance mode

```
# hosted-engine --set-maintenance --mode=none
```

The hosted engine starts automatically after a short wait.

12. Wait for nodes to synchronize

Run the following command on the first hyperconverged host to check synchronization status. If engine status is listed as **unknown stale-data**, synchronization requires several more minutes to complete.

The following output indicates completed synchronization.

```
# hosted-engine --vm-status | grep 'Engine status'
Engine status : {"health": "good", "vm": "up", "detail": "up"}
Engine status : {"reason": "vm not running on this host",
  "health": "bad", "vm": "down", "detail": "unknown"}
Engine status : {"reason": "vm not running on this host",
  "health": "bad", "vm": "down", "detail": "unknown"}
```

13. Activate all storage domains

Activate the master storage domain first, followed by all other storage domains.

For details on activating storage domains, see *Activating Storage Domains from Maintenance*

Mode in the Red Hat Virtualization documentation:

https://access.redhat.com/documentation/en-us/red_hat_virtualization/4.3/html/administration_guide/sect-storage_tasks.

14. Start all virtual machines

See *Starting a Virtual Machine* in the Red Hat Virtualization documentation for details:

https://access.redhat.com/documentation/en-us/red_hat_virtualization/4.3/html/virtual_machine_management_guide/sect-starting_the_virtual_machine.

CHAPTER 5. CONFIGURE PERFORMANCE IMPROVEMENTS

Some deployments benefit from additional configuration to achieve optimal performance. This section covers recommended additional configuration for certain deployments.

5.1. IMPROVING VOLUME PERFORMANCE BY CHANGING SHARD SIZE

The default value of the **shard-block-size** parameter changed from **4MB** to **64MB** between Red Hat Hyperconverged Infrastructure for Virtualization version 1.0 and 1.1. This means that all new volumes are created with a **shard-block-size** value of 64MB. However, existing volumes retain the original **shard-block-size** value of 4MB.

There is no safe way to modify the **shard-block-size** value on volumes that contain data. Because shard block size applies only to writes that occur after the value is set, attempting to change the value on a volume that contains data results in a mixed shard block size, which results in poor performance.

This section shows you how to safely modify the shard block size on an existing volume after upgrading to Red Hat Hyperconverged Infrastructure for Virtualization 1.1 or higher, in order to take advantage of the performance benefits of a larger shard size.

5.1.1. Changing shard size on replicated volumes

1. Create an inventory file

Create an inventory file called **normal_replicated_inventory.yml** based on the following example.

Replace **host1**, **host2**, and **host3** with the FQDNs of your hosts, and edit device details to match your environment.

Example **normal_replicated_inventory.yml** inventory file

```
hc_nodes:
  hosts:
    # Host1
    host1:
      # Dedupe & Compression config
      # If logicalsize >= 1000G then slabsize=32G else slabsize=2G
      #gluster_infra_vdo:
      # - { name: 'vdo_sdb', device: '/dev/sdb', logicalsize: '3000G', emulate512: 'on', slabsize:
'32G',
      #   blockmapcachesize: '128M', readcachesize: '20M', readcache: 'enabled',
writepolicy: 'auto' }

      # With Dedupe & Compression
      #gluster_infra_volume_groups:
      # - vgroup: <volgroup_name>
      #   pvname: /dev/mapper/vdo_sdb

      # Without Dedupe & Compression
      gluster_infra_volume_groups:
      - vgroup: <volgroup_name>
        pvname: /dev/sdb

      gluster_infra_mount_devices:
      - path: <brick_mountpoint>
```

```

    lvname: <lv_name>
    vgname: <volgroup_name>

gluster_infra_thinpools:
  - {vgname: '<volgroup_name>', thinpoolname: 'thinpool_<volgroup_name>',
thinpoolsize: '500G', poolmetadatasize: '4G'}

gluster_infra_lv_logicalvols:
  - vgname: <volgroup_name>
    thinpool: thinpool_<volgroup_name>
    lvname: <lv_name>
    lvsize: <size>G

# Mount the devices
gluster_infra_mount_devices:
  - { path: '<brick_mountpoint>', vgname: <volgroup_name>, lvname: <lv_name> }

#Host2
host2:
  # Dedupe & Compression config
  # If logicalsize >= 1000G then slabsize=32G else slabsize=2G
  #gluster_infra_vdo:
  # - { name: 'vdo_sdb', device: '/dev/sdb', logicalsize: '3000G', emulate512: 'on', slabsize:
'32G',
  #   blockmapcachesize: '128M', readcachesize: '20M', readcache: 'enabled',
writepolicy: 'auto' }

  # With Dedupe & Compression
  #gluster_infra_volume_groups:
  # - vgname: <volgroup_name>
  #   pvname: /dev/mapper/vdo_sdb

  # Without Dedupe & Compression
  gluster_infra_volume_groups:
  - vgname: <volgroup_name>
    pvname: /dev/sdb

gluster_infra_mount_devices:
  - path: <brick_mountpoint>
    lvname: <lv_name>
    vgname: <volgroup_name>

gluster_infra_thinpools:
  - {vgname: '<volgroup_name>', thinpoolname: 'thinpool_<volgroup_name>',
thinpoolsize: '500G', poolmetadatasize: '4G'}

gluster_infra_lv_logicalvols:
  - vgname: <volgroup_name>
    thinpool: thinpool_<volgroup_name>
    lvname: <lv_name>
    lvsize: <size>G

# Mount the devices
gluster_infra_mount_devices:
  - { path: '<brick_mountpoint>', vgname: <volgroup_name>, lvname: <lv_name> }

```

```

#Host3
host3:
  # Dedupe & Compression config
  # If logicalsize >= 1000G then slabsize=32G else slabsize=2G
  #gluster_infra_vdo:
  # - { name: 'vdo_sdb', device: '/dev/sdb', logicalsize: '3000G', emulate512: 'on', slabsize:
'32G',
  #   blockmapcachesize: '128M', readcachesize: '20M', readcache: 'enabled',
writepolicy: 'auto' }

  # With Dedupe & Compression
  #gluster_infra_volume_groups:
  # - vgroupname: <volgroup_name>
  #   pvname: /dev/mapper/vdo_sdb

  # Without Dedupe & Compression
  gluster_infra_volume_groups:
  - vgroupname: <volgroup_name>
    pvname: /dev/sdb

  gluster_infra_mount_devices:
  - path: <brick_mountpoint>
    lvname: <lv_name>
    vgroupname: <volgroup_name>

  gluster_infra_thinpools:
  - {vgroupname: '<volgroup_name>', thinpoolname: 'thinpool_<volgroup_name>',
thinpoolsize: '500G', poolmetadatasize: '4G'}

  gluster_infra_lv_logicalvols:
  - vgroupname: <volgroup_name>
    thinpool: thinpool_<volgroup_name>
    lvname: <lv_name>
    lvsize: <size>G

  # Mount the devices
  gluster_infra_mount_devices:
  - { path: '<brick_mountpoint>', vgroupname: <volgroup_name>, lvname: <lv_name> }

# Common configurations
vars:
  cluster_nodes:
  - host1
  - host2
  - host3
  gluster_features_hci_cluster: "{{ cluster_nodes }}"
  gluster_features_hci_volumes:
  - { volname: 'data', brick: '<brick_mountpoint>' }
  gluster_features_hci_volume_options:
  {
    group: 'virt',
    storage.owner-uid: '36',
    storage.owner-gid: '36',
    network.ping-timeout: '30',
    performance.strict-o-direct: 'on',
    network.remote-dio: 'off',
  }

```

```

cluster.granular-entry-heal: 'enable',
features.shard-block-size: '64MB'
}

```

2. Create the `normal_replicated.yml` playbook

Create a `normal_replicated.yml` playbook file using the following example:

Example `normal_replicated.yml` playbook

```

---

# Safely changing the shard block size parameter value for normal replicated volume
- name: Changing the shard block size
  hosts: hc_nodes
  remote_user: root
  gather_facts: no
  any_errors_fatal: true

  roles:
    - gluster.infra
    - gluster.features

```

3. Run the playbook

```
ansible-playbook -i normal_replicated_inventory.yml normal_replicated.yml
```

5.1.2. Changing shard size on arbitrated volumes

1. Create an inventory file

Create an inventory file called `arbitrated_replicated_inventory.yml` based on the following example.

Replace **host1**, **host2**, and **host3** with the FQDNs of your hosts, and edit device details to match your environment.

Example `arbitrated_replicated_inventory.yml` inventory file

```

hc_nodes:
  hosts:
    # Host1
    host1:
      # Dedupe & Compression config
      # If logicalsize >= 1000G then slabsize=32G else slabsize=2G
      #gluster_infra_vdo:
      # - { name: 'vdo_sdb', device: '/dev/sdb', logicalsize: '3000G', emulate512: 'on', slabsize:
'32G',
      #   blockmapcachesize: '128M', readcachesize: '20M', readcache: 'enabled',
writepolicy: 'auto' }

      # With Dedupe & Compression
      #gluster_infra_volume_groups:
      # - vgroup: <volgroup_name>
      #   pvname: /dev/mapper/vdo_sdb

```



```

# Without Dedupe & Compression
gluster_infra_volume_groups:
  - vgroupname: <volgroup_name>
    pvname: /dev/sdb

gluster_infra_mount_devices:
  - path: <brick_mountpoint>
    lvname: <lv_name>
    vgroupname: <volgroup_name>
  - path: <brick_mountpoint>
    lvname: <lv_name>
    vgroupname: <volgroup_name>

gluster_infra_thinpools:
  - {vgroupname: '<volgroup_name>', thinpoolname: 'thinpool_<volgroup_name>',
thinpoolsize: '500G', poolmetadatasize: '4G'}

gluster_infra_lv_logicalvols:
  - vgroupname: <volgroup_name>
    thinpool: thinpool_<volgroup_name>
    lvname: <lv_name>
    lvsize: <size>G
  - vgroupname: <volgroup_name>
    thinpool: thinpool_<volgroup_name>
    lvname: <lv_name>
    lvsize: <size>G

# Mount the devices
gluster_infra_mount_devices:
  - { path: '<brick_mountpoint>', vgroupname: <volgroup_name>, lvname: <lv_name> }
  - { path: '<brick_mountpoint>', vgroupname: <volgroup_name>, lvname: <lv_name> }

#Host2
host2:
  # Dedupe & Compression config
  # If logicalsize >= 1000G then slabsize=32G else slabsize=2G
  #gluster_infra_vdo:
  # - { name: 'vdo_sdb', device: '/dev/sdb', logicalsize: '3000G', emulate512: 'on', slabsize:
'32G',
  #   blockmapcachesize: '128M', readcachesize: '20M', readcache: 'enabled',
writepolicy: 'auto' }

# With Dedupe & Compression
#gluster_infra_volume_groups:
# - vgroupname: <volgroup_name>
#   pvname: /dev/mapper/vdo_sdb

# Without Dedupe & Compression
gluster_infra_volume_groups:
  - vgroupname: <volgroup_name>
    pvname: /dev/sdb

gluster_infra_mount_devices:
  - path: <brick_mountpoint>
    lvname: <lv_name>
    vgroupname: <volgroup_name>

```

```

- path: <brick_mountpoint>
  lvname: <lv_name>
  vgname: <volgroup_name>

gluster_infra_thinpools:
- {vgname: '<volgroup_name>', thinpoolname: 'thinpool_<volgroup_name>',
thinpoolsize: '500G', poolmetadatasize: '4G'}

gluster_infra_lv_logicalvols:
- vgname: <volgroup_name>
  thinpool: thinpool_<volgroup_name>
  lvname: <lv_name>
  lvsize: <size>G
- vgname: <volgroup_name>
  thinpool: thinpool_<volgroup_name>
  lvname: <lv_name>
  lvsize: <size>G

# Mount the devices
gluster_infra_mount_devices:
- { path: '<brick_mountpoint>', vgname: <volgroup_name>, lvname: <lv_name> }
- { path: '<brick_mountpoint>', vgname: <volgroup_name>, lvname: <lv_name> }

#Host3
host3:
# Dedupe & Compression config
# If logicalsize >= 1000G then slabsize=32G else slabsize=2G
#gluster_infra_vdo:
# - { name: 'vdo_sdb', device: '/dev/sdb', logicalsize: '3000G', emulate512: 'on', slabsize:
'32G',
#   blockmapcachesize: '128M', readcachesize: '20M', readcache: 'enabled',
writepolicy: 'auto' }

# With Dedupe & Compression
#gluster_infra_volume_groups:
# - vgname: <volgroup_name>
#   pvname: /dev/mapper/vdo_sdb

# Without Dedupe & Compression
gluster_infra_volume_groups:
- vgname: <volgroup_name>
  pvname: /dev/sdb

gluster_infra_mount_devices:
- path: <brick_mountpoint>
  lvname: <lv_name>
  vgname: <volgroup_name>

gluster_infra_thinpools:
- {vgname: '<volgroup_name>', thinpoolname: 'thinpool_<volgroup_name>',
thinpoolsize: '500G', poolmetadatasize: '4G'}

gluster_infra_lv_logicalvols:
- vgname: <volgroup_name>
  thinpool: thinpool_<volgroup_name>
  lvname: <lv_name>

```

```

    lvsizes: <size>G

    # Mount the devices
    gluster_infra_mount_devices:
      - { path: '<brick_mountpoint>', vgroupname: <volgroup_name>, lvname: <lv_name> }

    # Common configurations
    vars:
      cluster_nodes:
        - host1
        - host2
        - host3
      gluster_features_hci_cluster: "{{ cluster_nodes }}"
      gluster_features_hci_volumes:
        - { volname: 'data_one', brick: '<brick_mountpoint>', arbiter: 1 }
      gluster_features_hci_volume_options:
        {
          group: 'virt',
          storage.owner-uid: '36',
          storage.owner-gid: '36',
          network.ping-timeout: '30',
          performance.strict-o-direct: 'on',
          network.remote-dio: 'off',
          cluster.granular-entry-heal: 'enable',
          features.shard-block-size: '64MB',
          server.ssl: 'on',
          client.ssl: 'on',
          auth.ssl-allow: '<host1>;<host2>;<host3>'
        }

```

2. Create the `arbitrated_replicated.yml` playbook

Create a **arbitrated_replicated.yml** playbook file using the following example:

Example `arbitrated_replicated.yml` playbook

```

---

# Safely changing the shard block size parameter value for arbitrated replicated volume
- name: Changing the shard block size
  hosts: hc_nodes
  remote_user: root
  gather_facts: no
  any_errors_fatal: true

  roles:
    - gluster.infra
    - gluster.features

```

3. Run the playbook

```
ansible-playbook -i arbitrated_replicated_inventory.yml arbitrated_replicated.yml
```

5.2. CONFIGURING A LOGICAL VOLUME CACHE (LVMCACHE) FOR AN EXISTING VOLUME

If your main storage devices are not Solid State Disks (SSDs), Red Hat recommends configuring a logical volume cache (lvmcache) to achieve the required performance for Red Hat Hyperconverged Infrastructure for Virtualization deployments.

1. Create inventory file

Create an inventory file called **cache_inventory.yml** based on the example below.

Replace **<host1>**, **<host2>**, and **<host3>** with the FQDNs of the hosts on which to configure the cache.

Replace the following values throughout the file.

<slow_device>,**<fast_device>**

Specify the device to which the cache should attach, followed by the cache device, as a comma-delimited list, for example, **cachedisk: '/dev/sdb,/dev/sde'**.

<fast_device_name>

Specify the name of the cache logical volume to create, for example, **cachelv_thinpool_gluster_vg_sde**

<fast_device_thinpool>

Specify the name of the cache thin pool to create, for example, **gluster_thinpool_gluster_vg_sde**.

Example cache_inventory.yml file

```
hc_nodes:
  hosts:
    # Host1
    <host1>:
      gluster_infra_cache_vars:
        - vgname: gluster_vg_sdb
          cachedisk: '<slow_device>,<fast_device>'
          cachelvname: <fast_device_name>
          cachethinpoolname: <fast_device_thinpool>
          cachelvsize: '10G'
          cachemode: writethrough

    #Host2
    <host2>:
      gluster_infra_cache_vars:
        - vgname: gluster_vg_sdb
          cachedisk: '<slow_device>,<fast_device>'
          cachelvname: <fast_device_name>
          cachethinpoolname: <fast_device_thinpool>
          cachelvsize: '10G'
          cachemode: writethrough

    #Host3
    <host3>:
      gluster_infra_cache_vars:
        - vgname: gluster_vg_sdb
          cachedisk: '<slow_device>,<fast_device>'
          cachelvname: <fast_device_name>
```

```
cachethinpoolname: <fast_device_thinpool>
cachelsize: '10G'
cachemode: writethrough
```

2. Create a playbook file

Create an ansible playbook file named **lvm_cache.yml**.

Example lvm_cache.yml file

```
---
# Create LVM Cache
- name: Setup LVM Cache
  hosts: hc_nodes
  remote_user: root
  gather_facts: no
  any_errors_fatal: true

  roles:
    - gluster.infra
```

3. Run the playbook with thecachesetup tag

Run the following command to apply the configuration specified in **lvm_cache.yml** to the hosts and devices specified in **cache_inventory.yml**.

```
ansible-playbook -i cache_inventory.yml lvm_cache.yml --tags cachesetup
```

CHAPTER 6. CONFIGURE MONITORING

6.1. CONFIGURING EVENT NOTIFICATIONS

To configure which notifications you want to be displayed in the Administration Portal, see [Configuring Event Notifications in the Administration Portal](#) in the Red Hat Virtualization 4.3 *Administration Guide*.

PART II. MAINTENANCE TASKS

CHAPTER 7. BASIC OPERATIONS

Some basic operations are required for many administrative and troubleshooting tasks. This section covers how to safely perform basic tasks like shutting down and starting up the hyperconverged cluster.

7.1. CREATING A SHUTDOWN PLAYBOOK

A hyperconverged environment must be shut down in a particular order. The simplest way to do this is to create a shutdown playbook that can be run from the Hosted Engine virtual machine.

The `ovirt.shutdown_env` role enables Global Maintenance Mode, and initiates shutdown for all virtual machines and hosts in the cluster. Host shutdown is asynchronous. The playbook terminates before hyperconverged hosts are actually shut down.

Prerequisites

- Ensure that the **`ovirt.shutdown_env`** ansible role is available on the Hosted Engine virtual machine.

```
# yum install ovirt-ansible-shutdown-env -y
```

Procedure

1. Log in to the Hosted Engine virtual machine.
2. Create a shutdown playbook for your environment. Use the following template to create the playbook file.
 - Replace **`ovirt-engine.example.com`** with the FQDN of your Hosted Engine virtual machine.
 - Replace **`123456`** with the password for the **`admin@internal`** account.

Example playbook file: `shutdown_rhhi-v.yml`

```
---
- name: oVirt shutdown environment
  hosts: localhost
  connection: local
  gather_facts: false

  vars:
    engine_url: https://ovirt-engine.example.com/ovirt-engine/api
    engine_user: admin@internal
    engine_password: 123456
    engine_cafile: /etc/pki/ovirt-engine/ca.pem

  roles:
    - ovirt.shutdown_env
```

7.2. SHUTTING DOWN RHHI FOR VIRTUALIZATION

A hyperconverged environment must be shut down in a particular order. Use an Ansible playbook to automate this process and ensure that your environment is shut down safely.

Prerequisites

- Create a shutdown playbook as described in [Creating a shutdown playbook](#)
- Ensure that the **ovirt.shutdown_env** ansible role is available on the Hosted Engine virtual machine.

```
# yum install ovirt-ansible-shutdown-env -y
```

Procedure

1. Run the shutdown playbook against the Hosted Engine virtual machine.

```
# ansible-playbook -i localhost <shutdown_rhhi-v.yml>
```

7.3. STARTING UP A HYPERCONVERGED CLUSTER

Starting up a hyperconverged cluster is more complex than starting up a traditional compute or storage cluster. Follow these instructions to start up your hyperconverged cluster safely.

1. Power on all hosts in the cluster.
2. Ensure that the required services are available.
 - a. Verify that the **glusterd** service started correctly on all hosts.

```
# systemctl status glusterd
● glusterd.service - GlusterFS, a clustered file-system server
   Loaded: loaded (/usr/lib/systemd/system/glusterd.service; enabled; vendor preset: disabled)
   Drop-In: /etc/systemd/system/glusterd.service.d
            └─99-cpu.conf
   Active: active (running) since Wed 2018-07-18 11:15:03 IST; 3min 48s ago
   [...]

```

If glusterd is not started, start it.

```
# systemctl start glusterd
```

- b. Verify that host networks are available and hosts have IP addresses assigned to the required interfaces.

```
# ip addr show
```

- c. Verify that all hosts are part of the storage cluster (listed as *Peer in Cluster (Connected)*).

```
# gluster peer status

Number of Peers: 2

Hostname: 10.70.37.101
Uuid: 773f1140-68f7-4861-a996-b1ba97586257
State: Peer in Cluster (Connected)

```

```

Hostname: 10.70.37.102
Uuid: fc4e7339-9a09-4a44-aa91-64dde2fe8d15
State: Peer in Cluster (Connected)

```

- d. Verify that all bricks are shown as online.

```

# gluster volume status engine
Status of volume: engine
Gluster process                TCP Port  RDMA Port  Online  Pid
-----
Brick 10.70.37.28:/gluster_bricks/engine/engine 49153    0          Y      23160
Brick 10.70.37.29:/gluster_bricks/engine/engine 49160    0          Y      12392
Brick 10.70.37.30:/gluster_bricks/engine/engine 49157    0          Y      15200
Self-heal Daemon on localhost                N/A      N/A        Y      23008
Self-heal Daemon on 10.70.37.30              N/A      N/A        Y      10905
Self-heal Daemon on 10.70.37.29              N/A      N/A        Y      13568

Task Status of Volume engine
-----
There are no active volume tasks

```

3. Start the hosted engine virtual machine.
 - a. Run the following command on the host that you want to be the hosted engine node.

```
# hosted-engine --vm-start
```

- b. Verify that the hosted engine virtual machine has started correctly.

```
# hosted-engine --vm-status
```

4. Take the hosted engine virtual machine out of Global Maintenance mode.
 - a. Log in to the Administration Portal.
 - b. Click **Compute** → **Hosts** and select the hosted engine node.
 - c. Click **⋮** → **Disable Global HA Maintenance**.
5. Start any other virtual machines using the Web Console.
 - a. Click **Compute** → **Virtualization**.
 - b. Select any virtual machines you want to start and click **Run**.

CHAPTER 8. UPGRADING TO RED HAT HYPERCONVERGED INFRASTRUCTURE FOR VIRTUALIZATION 1.6

Upgrading involves moving from one version of a product to a newer major release of the same product. This section shows you how to upgrade to Red Hat Hyperconverged Infrastructure for Virtualization 1.6 from version 1.5.

From a component standpoint, this involves the following:

- Upgrading the Hosted Engine virtual machine to Red Hat Virtualization Manager version 4.3.
- Upgrading the physical hosts to Red Hat Virtualization 4.3.

8.1. 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.

8.2. UPGRADE WORKFLOW

Red Hat Hyperconverged Infrastructure for Virtualization is a software solution comprised of several different components. Upgrade the components in the following order to minimize disruption to your deployment:

1. [Prepare the systems to be upgraded](#).
2. [Upgrade the Hosted Engine virtual machine](#).
3. [Upgrade the hyperconverged hosts](#).

8.3. PREPARING TO UPGRADE

8.3.1. Update to the latest version of the previous release

Ensure that you are using the latest version (4.2.8) of Red Hat Virtualization Manager 4.2 on the hosted engine virtual machine, and the latest version of Red Hat Virtualization 4.2 on the hosted engine node.

See the Red Hat Virtualization [Self-Hosted Engine Guide](#) for the Red Hat Virtualization 4.3 update process.



IMPORTANT

Do not proceed with the following prerequisites until you have updated to the latest version of Red Hat Virtualization 4.2.

8.3.2. Update subscriptions

You can check which repositories a machine has access to by running the following command as the root user:

```
# subscription-manager repos --list-enabled
```

- Verify that the Hosted Engine virtual machine is subscribed to the following repositories:
 - **rhel-7-server-rhv-4.3-manager-rpms**
 - **rhel-7-server-rhv-4-manager-tools-rpms**
 - **rhel-7-server-rpms**
 - **rhel-7-server-supplementary-rpms**
 - **jb-eap-7-for-rhel-7-server-rpms**
 - **rhel-7-server-ansible-2-rpms**
- Verify that the Hosted Engine virtual machine is not subscribed to previous versions of the required repositories.
 - **rhel-7-server-rhv-4.3-manager-rpms** replaces the **rhel-7-server-rhv-4.2-manager-rpms** repository

Subscribe a machine to a repository by running the following command on that machine:

```
# subscription-manager repos --enable=<repository>
```

8.3.3. Verify that data is not currently being synchronized using geo-replication

- Click the **Tasks** tab at the bottom right of the Manager. Ensure that there are no ongoing tasks related to Data Synchronization. If data synchronization tasks are present, wait until they are complete before beginning the update.
- Stop all geo-replication sessions so that synchronization will not occur during the update. Click the **Geo-replication** subtab and select the session that you want to stop, then click **Stop**. Alternatively, run the following command to stop a geo-replication session:

```
# gluster volume geo-replication <MASTER_VOL> <SLAVE_HOST>::<<SLAVE_VOL> stop
```

8.4. UPGRADING RED HAT HYPERCONVERGED INFRASTRUCTURE FOR VIRTUALIZATION

8.4.1. Upgrading the Hosted Engine virtual machine

1. Place the cluster into Global Maintenance mode

- a. Log in to the Web Console.
- b. Click **Virtualization** → **Hosted Engine**.
- c. Click **Put this cluster into global maintenance**

2. Upgrade Red Hat Virtualization Manager.

- a. Log in to the Hosted Engine virtual machine.
- b. Upgrade the setup packages:

```
# yum update ovirt*setup*
```

- c. Run **engine-setup** and follow the prompts to upgrade the Manager.
This process can take a while and cannot be aborted, so Red Hat recommends running it inside a **screen** session. See [How to use the screen command](#) for more information about this function.
- d. Upgrade all other packages.

```
# yum update
```

- e. Reboot the Hosted Engine virtual machine to ensure all updates are applied.

```
# reboot
```

3. Restart the Hosted Engine virtual machine.

- a. Log in to any hyperconverged host.
- b. Start the Hosted Engine virtual machine.

```
# hosted-engine --vm-start
```

- c. Verify the status of the Hosted Engine virtual machine.

```
# hosted-engine --vm-status
```

4. Remove the cluster from Global Maintenance mode.

- a. Log in to the Web Console.
- b. Click **Virtualization** → **Hosted Engine**.
- c. Click **Remove this cluster from global maintenance**

8.4.2. Upgrading the hyperconverged hosts



IMPORTANT

If you are upgrading a host from Red Hat Virtualization 4.2.7 or 4.2.7-1, ensure that the hosted engine virtual machine is not running on that host during the upgrade process. This is related to a bug introduced in Red Hat Enterprise Linux 7.6, [BZ#1641798](#), which affects these versions of Red Hat Virtualization.

To work around this issue, stop the hosted engine virtual machine before upgrading a host, and start it on another host.

```
[root@host1] # hosted-engine --vm-shutdown
```

```
[root@host2] # hosted-engine --vm-start
```

1. Upgrade one host at a time

Perform the following steps on one hyperconverged host at a time.

a. Upgrade the hyperconverged host.

- i. In the Manager, click **Compute** → **Hosts** and select a node.
- ii. Click **Installation** → **Upgrade**.
- iii. Click **OK** to confirm the upgrade.
Wait for the upgrade to complete, and for the host to become available again.

b. Verify self-healing is complete before upgrading the next host.

- i. Click the name of the host.
- ii. Click the **Bricks** tab.
- iii. Verify that the **Self-Heal Info** column of all bricks is listed as **OK** before upgrading the next host.

2. Update cluster compatibility

When all hosts are upgraded, update cluster compatibility setting.

- a. In the Manager, click **Cluster** and select the cluster (**Default**).
- b. Click **Edit**.
- c. Update the value of **Cluster compatibility** to **4.3** and save.

d. Restart all virtual machines

This ensures that the new cluster compatibility setting takes effect.

- i. Click **Compute** → **Virtual Machines** and select a running virtual machine.
- ii. Click **Reboot**.
- iii. Click **OK** in the *Reboot Virtual Machine(s)* confirmation window.
The **Status** of the virtual machine changes to **Reboot In Progress** before returning to **Up**.

- If upgrading a hyperconverged host fails because of a conflict with the **rhvm-appliance** package, log in to the hyperconverged host and follow the steps in [RHV: RHV-H Upgrade failed](#) before continuing.

CHAPTER 9. MONITORING RED HAT HYPERCONVERGED INFRASTRUCTURE FOR VIRTUALIZATION

9.1. MONITORING VIRTUAL DATA OPTIMIZER (VDO)

Monitoring VDO helps in understanding when the physical storage is running out of space. Physical space in VDO needs to be monitored like thin provisioned storage. VDO devices should use thin provisioning because more logical space will be available and VDO space will be used in a more effective way. By default thin provisioning is enabled and it can be unchecked as required.

You can check available blocks, used space, and device information by clicking on **View Details**.

9.1.1. Monitoring VDO using the command line interface

There are several options for monitoring VDO using the command line interface.

The `vdostats` command

This command displays volume statistics including available blocks, number of blocks used, device name, percentage of physical blocks saved, and percentage of physical blocks on a VDO volume. For more information on `vdostats`, see the manual page: **man `vdostats`**.

The `vdo status` command

This command reports VDO system and volume status in YAML format.

The `/sys/kvdo/<vdo_volume>/statistics` directory

Files in this directory include volume statistics for VDO. You can read these files instead of using the **`vdostats`** command.

9.1.2. Monitoring VDO using the Web Console

Events related to VDO usage are displayed under the **Notifications** tab. Events provide information about the physical space remaining on the VDO volume, and keep you up to date about whether more physical space is needed.

Table 9.1. Types of Event Notification

Type	Text	Actions
Warn	Warning, low confirmed disk space. StorageDomainName domain has DiskSpace GB of confirmed free space.	<ul style="list-style-type: none"> • Delete data that is not required. • Replace existing disks with larger disks. • Add more Red Hat Gluster Storage servers and expand the volume across the new servers. • Add more disks and expand the logical volumes underlying the Gluster volume.

CHAPTER 10. FREEING SPACE ON THINLY-PROVISIONED LOGICAL VOLUMES USING FSTRIM

You can manually run **fstrim** to return unused logical volume space to the thin pool so that it is available for other logical volumes.

Red Hat recommends running **fstrim** daily.

Prerequisites

- Verify that the thin pool logical volume supports discard behavior. Discard is supported if the output of the following command for the underlying device is not zero.

```
# cat /sys/block/<device>/queue/discard_max_bytes
```

Procedure

1. Run **fstrim** to restore physical space to the thin pool.

```
# fstrim -v <mountpoint>
```

For example, the following command discards any unused space it finds on the logical volume mounted at **/gluster_bricks/data/data**, and provides verbose output (**-v**).

```
# fstrim -v /gluster_bricks/data/data
```

Additional resources

- See [Scheduling a recurring job using cron](#) for information on configuring an automatically recurring task.

CHAPTER 11. ADD HYPERCONVERGED HOSTS TO RED HAT VIRTUALIZATION MANAGER

Follow this process to allow Red Hat Virtualization Manager to manage an existing hyperconverged host.

1. Log in to the Administration Portal.
2. Click **Compute** → **Hosts**.
3. Click **New**. The **New Host** window opens.
4. On the **General** tab, specify the following details about your hyperconverged host.
 - **Host Cluster**
 - **Name**
 - **Hostname**
 - **Password**
5. On the **General** tab, click the **Advanced Parameters** dropdown, and uncheck the **Automatically configure host firewall** checkbox.
6. Click **OK**.

CHAPTER 12. REINSTALLING A HYPERCONVERGED HOST


Some configuration changes require a hyperconverged host to be reinstalled before the configuration change can take effect. Follow these steps to reinstall a hyperconverged host.

1. Log in to the Administration Portal.
2. Click **Compute** → **Hosts**.
3. Select the host and click **Management** > **Maintenance** > **OK** to place this host in Maintenance mode.
4. Click **Installation** > **Reinstall** to open the Reinstall window.
5. On the General tab, uncheck the **Automatically Configure Host firewall** checkbox.
6. On the **Hosted Engine** tab, set the value of **Choose hosted engine deployment action** to **Deploy**.
7. Click **OK** to reinstall the host.

CHAPTER 13. REPLACING HOSTS

13.1. REPLACING THE PRIMARY HYPERCONVERGED HOST USING ANSIBLE

Follow this section to replace the hyperconverged host that you used to perform all deployment operations.



IMPORTANT

When self-signed encryption is enabled, replacing a node is a disruptive process that requires virtual machines and the Hosted Engine to be shut down.

1. (Optional) If encryption using a Certificate Authority is enabled, follow the steps under *Expanding Volumes* in the [Network Encryption](#) chapter of the Red Hat Gluster Storage 3.4 *Administration Guide*.
2. Move the server to be replaced into Maintenance mode.
 - a. In the Administration Portal, click **Compute** → **Hosts** and select the host to replace.
 - b. Click **Management** → **Maintenance** and click **OK** to move the host to Maintenance mode.
3. Install the replacement host
Follow the instructions in *Deploying Red Hat Hyperconverged Infrastructure for Virtualization for Virtualization* to install the physical machine and configure storage on the new host.
 - a. [Installing hyperconverged hosts](#)
 - b. [Configuring Public Key based SSH Authentication](#)
 - c. [Configuring RHGS for Hosted Engine using the Web Console](#)
4. Configure the replacement host
Follow the instructions in [Section 13.3, "Preparing a replacement hyperconverged host using ansible"](#).
5. (Optional) If encryption with self-signed certificates is enabled:
 - a. Generate the private key and self-signed certificate on the replacement host. See the Red Hat Gluster Storage *Administration Guide* for details:
https://access.redhat.com/documentation/en-us/red_hat_gluster_storage/3.4/html/administration_guide/chap-network_encryption#chap-Network_Encryption-Prereqs.
 - b. On a healthy host, create a copy of the `/etc/ssl/glusterfs.ca` file.


```
# cp /etc/ssl/glusterfs.ca /etc/ssl/glusterfs.ca.bk
```
 - c. Append the new host's certificate to the content of the original `/etc/ssl/glusterfs.ca` file.
 - d. Distribute the `/etc/ssl/glusterfs.ca` file to all hosts in the cluster, including the new host.
 - e. Run the following command on the replacement host to enable management encryption:


```
■
```

```
# touch /var/lib/glusterd/secure-access
```

- f. Include the new host in the value of the **auth.ssl-allow** volume option by running the following command for each volume.

```
# gluster volume set <volname> auth.ssl-allow "<old_host1>,<old_host2>,<new_host>"
```

- g. Restart the glusterd service on all hosts.

```
# systemctl restart glusterd
```

- h. Follow the steps in [Section 4.1, "Configuring TLS/SSL using self-signed certificates"](#) to remount all gluster processes.

6. Add the replacement host to the cluster.

Run the following command from any host already in the cluster.

```
# gluster peer probe <new_host>
```

7. Move the Hosted Engine into Maintenance mode:

```
# hosted-engine --set-maintenance --mode=global
```

8. Stop the ovirt-engine service.

```
# systemctl stop ovirt-engine
```

9. Update the database.

```
# hosted-engine --set-shared-config storage <new_host_IP>:/engine
```

10. Start the ovirt-engine service.

```
# systemctl start ovirt-engine
```

11. Stop all virtual machines except the Hosted Engine.

12. Move all storage domains **except** the Hosted Engine domain into Maintenance mode.

13. Stop the Hosted Engine virtual machine.

Run the following command on the existing server that hosts the Hosted Engine.

```
# hosted-engine --vm-shutdown
```

14. Stop high availability services on all hosts.

```
# systemctl stop ovirt-ha-agent
# systemctl stop ovirt-ha-broker
```

15. Disconnect Hosted Engine storage from the hyperconverged host.

Run the following command on the existing server that hosts the Hosted Engine.

■

```
# hosted-engine --disconnect-storage
```

16. Update the Hosted Engine configuration file.

Edit the storage parameter in the `/etc/ovirt-hosted-engine/hosted-engine.conf` file to use the replacement host.

```
storage=<new_server_IP>:/engine
```



NOTE

To configure the Hosted Engine for new hosts, use the command:

```
# hosted-engine --set-shared-config storage <new_server_IP>:/engine
```

17. Restart high availability services on all hosts.

```
# systemctl restart ovirt-ha-agent
# systemctl restart ovirt-ha-broker
```

18. Reboot the existing and replacement hosts.
Wait until all hosts are available before continuing.

19. Take the Hosted Engine out of Maintenance mode.

```
# hosted-engine --set-maintenance --mode=none
```

20. Verify that the replacement host is used.
On all hyperconverged hosts, verify that the **engine** volume is mounted from the replacement host by checking the IP address in the output of the **mount** command.

21. Activate storage domains.
Verify that storage domains mount using the IP address of the replacement host.

22. Using the RHV Management UI, add the replacement host.
Specify that the replacement host be used to host the Hosted Engine.

23. Move the replacement host into Maintenance mode.

```
# hosted-engine --set-maintenance --mode=global
```

24. Reboot the replacement host.
Wait until the host is back online before continuing.

25. Activate the replacement host from the RHV Management UI.
Ensure that all volumes are mounted using the IP address of the replacement host.

26. Replace engine volume brick.
Replace the brick on the old host that belongs to the **engine** volume with a new brick on the replacement host.

- a. Click **Storage** → **Volumes** and select the volume.

- b. Click the **Bricks** subtab.
 - c. Select the brick to replace, and then click **Replace brick**
 - d. Select the host that hosts the brick being replaced.
 - e. In the *Replace brick* window, provide the path to the new brick.
27. Remove the old host.
- a. Click **Compute** → **Hosts** and select the old host.
 - b. Click **Management** → **Maintenance** to move the host to maintenance mode.
 - c. Click **Remove**. The *Remove Host(s)* confirmation dialog appears.
 - d. If there are still volume bricks on this host, or the host is non-responsive, check the **Force Remove** checkbox.
 - e. Click **OK**.
 - f. Detach the old host from the cluster.

```
# gluster peer detach <old_host_IP> force
```

28. On the replacement host, run the following command to remove metadata from the previous host.

```
# hosted-engine --clean-metadata --host-id=<old_host_id> --force-clean
```

13.2. REPLACING OTHER HYPERCONVERGED HOSTS USING ANSIBLE

There are two options for replacing a hyperconverged host that is not the first host:

1. Replace the host with a new host that has a different fully-qualified domain name by following the instructions in [Section 13.2.1, “Replacing a hyperconverged host to use a different FQDN”](#) .
2. Replace the host with a new host that has the same fully-qualified domain name by following the instructions in [Section 13.2.2, “Replacing a hyperconverged host to use the same FQDN”](#) .

Follow the instructions in whichever section is appropriate for your deployment.

13.2.1. Replacing a hyperconverged host to use a different FQDN



IMPORTANT

When self-signed encryption is enabled, replacing a node is a disruptive process that requires virtual machines and the Hosted Engine to be shut down.

1. Install the replacement host
Follow the instructions in *Deploying Red Hat Hyperconverged Infrastructure for Virtualization for Virtualization* to install the physical machine.
 - a. [Installing host physical machines](#)

b. [Configuring Public Key based SSH Authentication](#)

2. Stop any existing geo-replication sessions

```
# gluster volume geo-replication MASTER_VOL SLAVE_HOST::SLAVE_VOL stop
```

For further information, see the Red Hat Gluster Storage *Administration Guide*: https://access.redhat.com/documentation/en-us/red_hat_gluster_storage/3.4/html/administration_guide/sect-starting_geo-replication#Stopping_a_Geo-replication_Session.

3. Move the host to be replaced into Maintenance mode

Perform the following steps in the Administration Portal:

- a. Click **Compute** → **Hosts** and select the hyperconverged host in the results list.
- b. Click **Management** → **Maintenance** and click **OK** to move the host to Maintenance mode.

4. Prepare the replacement host

- a. Configure key-based SSH authentication without a password

Configure key-based SSH authentication without a password from a physical machine still in the cluster to the replacement host. For details, see

https://access.redhat.com/documentation/en-us/red_hat_hyperconverged_infrastructure_for_virtualization/1.6/html/deploying_red_hat_hyp_configure-key-based-ssh-auth.

- b. Prepare the replacement host

Follow the instructions in [Section 13.3, "Preparing a replacement hyperconverged host using ansible"](#).

5. Create replacement brick directories

Ensure the new directories are owned by the **vds** user and the **kvm** group.

```
# mkdir /gluster_bricks/engine/engine
# chmod vds:kvm /gluster_bricks/engine/engine
# mkdir /gluster_bricks/data/data
# chmod vds:kvm /gluster_bricks/data/data
# mkdir /gluster_bricks/vmstore/vmstore
# chmod vds:kvm /gluster_bricks/vmstore/vmstore
```

6. (Optional) If encryption is enabled

- a. Generate the private key and self-signed certificate on the new server using the steps in the Red Hat Gluster Storage *Administration Guide*:

https://access.redhat.com/documentation/en-us/red_hat_gluster_storage/3.4/html/administration_guide/chap-network_encryption#chap-Network_Encryption-Prereqs.

If encryption using a Certificate Authority is enabled, follow the steps under *Expanding Volumes* in the [Network Encryption](#) chapter of the Red Hat Gluster Storage 3.4 *Administration Guide*.

- b. Add the new host's certificate to existing certificates.
 - i. On a healthy host, make a backup copy of the **/etc/ssl/glusterfs.ca** file.

- ii. Add the new host's certificate to the `/etc/ssl/glusterfs.ca` file on the healthy host.
- iii. Distribute the updated `/etc/ssl/glusterfs.ca` file to all other hosts, including the new host.

c. Enable management encryption

Run the following command on the new host to enable management encryption:

```
# touch /var/lib/glusterd/secure-access
```

d. Include the new host in the value of the `auth.ssl-allow` volume option by running the following command for each volume.

```
# gluster volume set <volname> auth.ssl-allow "<old_host1>,<old_host2>,<new_host>"
```

e. Restart the glusterd service on all hosts

```
# systemctl restart glusterd
```

f. If encryption uses self-signed certificates, follow the steps in [Section 4.1, "Configuring TLS/SSL using self-signed certificates"](#) to remount all gluster processes.

7. Add the new host to the existing cluster

a. Run the following command from one of the healthy hosts:

```
# gluster peer probe <new_host>
```

b. Add the new host to the existing cluster

- i. Click **Compute** → **Hosts** and then click **New** to open the *New Host* dialog.
- ii. Provide a **Name**, **Address**, and **Password** for the new host.
- iii. Uncheck the **Automatically configure host firewall** checkbox, as firewall rules are already configured by gdeploy.
- iv. In the **Hosted Engine** tab of the *New Host* dialog, set the value of **Choose hosted engine deployment action** to **Deploy**.
- v. Click **OK**.
- vi. When the host is available, click the name of the new host.
- vii. Click the **Network Interfaces** subtab and then click **Setup Host Networks**. The *Setup Host Networks* dialog appears.
- viii. Drag and drop the network you created for gluster to the IP associated with this host, and click **OK**.
See the Red Hat Virtualization 4.3 Self-Hosted Engine Guide for further details:
https://access.redhat.com/documentation/en-us/red_hat_virtualization/4.3/html/self-hosted_engine_guide/chap-installing_additional_hosts_to_a_self-hosted_environment.

8. Configure and mount shared storage on the new host

```
# cp /etc/fstab /etc/fstab.bk
# echo "<new_host>:/gluster_shared_storage /var/run/gluster/shared_storage/ glusterfs
defaults 0 0" >> /etc/fstab
# mount /gluster_shared_storage
```

9. Replace the old brick with the brick on the new host
 - a. In the Administration Portal, click **Storage** → **Volumes** and select the volume.
 - b. Click the **Bricks** subtab.
 - c. Select the brick that you want to replace and click **Replace Brick**. The *Replace Brick* dialog appears.
 - d. Specify the **Host** and the **Brick Directory** of the new brick.
 - e. Verify that brick heal completes successfully.
10. Click **Compute** → **Hosts**.
 11. Select the old host and click **Remove**.
Use **gluster peer status** to verify that the old host is no longer part of the cluster. If the old host is still present in the status output, run the following command to forcibly remove it:

```
# gluster peer detach <old_host> force
```

12. Clean old host metadata.

```
# hosted-engine --clean-metadata --host-id=<old_host_id> --force-clean
```

13. Set up new SSH keys for geo-replication of new brick.

```
# gluster system:: execute gsec_create
```

14. Recreate geo-replication session and distribute new SSH keys.

```
# gluster volume geo-replication <MASTER_VOL> <SLAVE_HOST>::<SLAVE_VOL> create
push-pem force
```

15. Start the geo-replication session.

```
# gluster volume geo-replication <MASTER_VOL> <SLAVE_HOST>::<SLAVE_VOL> start
```

13.2.2. Replacing a hyperconverged host to use the same FQDN



IMPORTANT

When self-signed encryption is enabled, replacing a node is a disruptive process that requires virtual machines and the Hosted Engine to be shut down.

1. (Optional) If encryption using a Certificate Authority is enabled, follow the steps under *Expanding Volumes* in the [Network Encryption](#) chapter of the Red Hat Gluster Storage 3.4 *Administration Guide*.
2. Move the host to be replaced into Maintenance mode
 - a. In the Administration Portal, click **Compute** → **Hosts** and select the hyperconverged host.
 - b. Click **Management** → **Maintenance**.
 - c. Click **OK** to move the host to Maintenance mode.
3. Install the replacement host

Follow the instructions in *Deploying Red Hat Hyperconverged Infrastructure for Virtualization for Virtualization* to install the physical machine and configure storage on the new host.

 - a. [Installing host physical machines](#)
 - b. [Configuring Public Key based SSH Authentication](#)
 - c. [Configuring RHGS for Hosted Engine using the Web Console](#)
4. Configure the replacement host

Follow the instructions in [Section 13.3, "Preparing a replacement hyperconverged host using ansible"](#).
5. (Optional) If encryption with self-signed certificates is enabled
 - a. Generate the private key and self-signed certificate on the replacement host. See the Red Hat Gluster Storage *Administration Guide* for details: https://access.redhat.com/documentation/en-us/red_hat_gluster_storage/3.4/html/administration_guide/chap-network_encryption#chap-Network_Encryption-Prereqs.
 - b. On a healthy host, make a backup copy of the **/etc/ssl/glusterfs.ca** file:


```
# cp /etc/ssl/glusterfs.ca /etc/ssl/glusterfs.ca.bk
```
 - c. Append the new host's certificate to the content of the **/etc/ssl/glusterfs.ca** file.
 - d. Distribute the **/etc/ssl/glusterfs.ca** file to all hosts in the cluster, including the new host.
 - e. Run the following command on the replacement host to enable management encryption:


```
# touch /var/lib/glusterd/secure-access
```
6. Replace the host machine

Follow the instructions in the Red Hat Gluster Storage *Administration Guide* to replace the host: https://access.redhat.com/documentation/en-us/red_hat_gluster_storage/3.4/html/administration_guide/sect-replacing_hosts#Replacing_a_Host_Machine_with_the_Same_Hostname.
7. Restart the glusterd service on all hosts

```
# systemctl restart glusterd
```

8. Verify that all hosts reconnect

```
# gluster peer status
```

9. **(Optional)** If encryption uses self-signed certificates, follow the steps in [Section 4.1, "Configuring TLS/SSL using self-signed certificates"](#) to remount all gluster processes.

10. Verify that all hosts reconnect and that brick heal completes successfully

```
# gluster peer status
```

11. Refresh fingerprint

- a. In the Administration Portal, click **Compute → Hosts** and select the new host.
- b. Click **Edit**.
- c. Click **Advanced Parameters** on the **General** tab.
- d. Click **fetch** to fetch the fingerprint from the host.
- e. Click **OK**.

12. Click **Installation → Reinstall** and provide the root password when prompted.

13. On the **Hosted Engine** tab set the value of **Choose hosted engine deployment action** to **Deploy**.

14. Attach the gluster network to the host

- a. Click **Compute → Hosts** and click the name of the host.
- b. Click the **Network Interfaces** subtab and then click **Setup Host Networks**.
- c. Drag and drop the newly created network to the correct interface.
- d. Ensure that the **Verify connectivity between Host and Engine** checkbox is checked.
- e. Ensure that the **Save network configuration** checkbox is checked.
- f. Click **OK** to save.

15. **Verify the health of the network**

Click the **Network Interfaces** tab and check the state of the host's network. If the network interface enters an "Out of sync" state or does not have an IP Address, click **Management → Refresh Capabilities**.

13.3. PREPARING A REPLACEMENT HYPERCONVERGED HOST USING ANSIBLE

Follow this process to replace a hyperconverged host in the cluster.

Prerequisites

- Ensure that the host you intend to replace is not associated with the FQDN that you want to use for the new host.
- Ensure that the new host is associated with the FQDN you want it to use.

Procedure

1. Create `node_prep_inventory.yml` inventory file

Create an inventory file called `node_prep_inventory.yml`, based on the following example.

Replace `host1` with the FQDN that you want to use for the new host, and device details with details appropriate for your host.

Example `node_prep_inventory.yml` file

```

hc_nodes:
  hosts:
    # New host
    newhost.example.com:

    # Dedupe & Compression config
    # If logicalsize >= 1000G then slabsize=32G else slabsize=2G
    #gluster_infra_vdo:
    # - { name: 'vdo_sdc', device: '/dev/sdc', logicalsize: '3000G', emulate512: 'on', slabsize:
'32G',
    #   blockmapcachesize: '128M', readcachesize: '20M', readcache: 'enabled',
writepolicy: 'auto' }

    # With Dedupe & Compression
    #gluster_infra_volume_groups:
    # - vgroup: gluster_vg_sdc
    #   pvname: /dev/mapper/vdo_sdc

    # Without Dedupe & Compression
    gluster_infra_volume_groups:
    - vgroup: gluster_vg_sdc
      pvname: /dev/sdc

    gluster_infra_mount_devices:
    - path: /gluster_bricks/engine
      lvname: gluster_lv_engine
      vgroup: gluster_vg_sdc
    - path: /gluster_bricks/data
      lvname: gluster_lv_data
      vgroup: gluster_vg_sdc
    - path: /gluster_bricks/vmstore
      lvname: gluster_lv_vmstore
      vgroup: gluster_vg_sdc

    gluster_infra_thinpools:
    - {vgroup: 'gluster_vg_sdc', thinpoolname: 'thinpool_gluster_vg_sdc', thinpoolsizesize:
'500G', poolmetadatasize: '4G'}

    # This is optional
    gluster_infra_cache_vars:
    - vgroup: gluster_vg_sdc

```

```

    cachedisk: /dev/sde
    cachelvname: cachelv_thinpool_vg_sdc
    cachethinpoolname: thinpool_gluster_vg_sdc # cachethinpoolname is equal to the
already created thinpool which you want to attach
    cachelvsize: '10G'
    cachemetalvsize: '2G'
    cachemetalvname: cache_thinpool_vg_sdc
    cachemode: writethrough

gluster_infra_thick_lvs:
  - vgname: gluster_vg_sdc
    lvname: gluster_lv_engine
    size: 100G

gluster_infra_lv_logicalvols:
  - vgname: gluster_vg_sdc
    thinpool: thinpool_gluster_vg_sdc
    lvname: gluster_lv_data
    lvsize: 500G
  - vgname: gluster_vg_sdc
    thinpool: thinpool_gluster_vg_sdc
    lvname: gluster_lv_vmstore
    lvsize: 500G

# Mount the devices
gluster_infra_mount_devices:
  - { path: '/gluster_bricks/data', vgname: gluster_vg_sdc, lvname: gluster_lv_data }
  - { path: '/gluster_bricks/vmstore', vgname: gluster_vg_sdc, lvname: gluster_lv_vmstore
}
  - { path: '/gluster_bricks/engine', vgname: gluster_vg_sdc, lvname: gluster_lv_engine }

# Common configurations
vars:
  # Firewall setup
  gluster_infra_fw_ports:
    - 2049/tcp
    - 54321/tcp
    - 5900/tcp
    - 5900-6923/tcp
    - 5666/tcp
    - 16514/tcp
  gluster_infra_fw_permanent: true
  gluster_infra_fw_state: enabled
  gluster_infra_fw_zone: public
  gluster_infra_fw_services:
    - glusterfs
  gluster_infra_disktype: RAID6
  gluster_infra_diskcount: 12
  gluster_infra_stripe_unit_size: 128

```

2. Create `node_prep.yml` playbook

Create a `node_prep.yml` playbook file based on the following example.

Example `node_prep.yml` playbook

```
---  
  
# Prepare Node for replace  
- name: Setup backend  
  hosts: hc_nodes  
  remote_user: root  
  gather_facts: no  
  any_errors_fatal: true  
  
  roles:  
    - gluster.infra  
    - gluster.features
```

3. Run `node_prep.yml` playbook

```
# ansible-playbook -i node_prep_inventory.yml node_prep.yml
```

CHAPTER 14. RECOVERING FROM DISASTER

This chapter explains how to restore your cluster to a working state after a disk or server failure.

You must have configured disaster recovery options previously in order to use this chapter. See [Configuring backup and recovery options](#) for details.

14.1. MANUALLY RESTORING DATA FROM A BACKUP VOLUME

This section covers how to restore data from a remote backup volume to a freshly installed replacement deployment of Red Hat Hyperconverged Infrastructure for Virtualization.

To do this, you must:

1. Install and configure a replacement deployment according to the instructions in [Deploying Red Hat Hyperconverged Infrastructure for Virtualization](#).

14.1.1. Restoring a volume from a geo-replicated backup

1. Install and configure a replacement Hyperconverged Infrastructure deployment
For instructions, refer to *Deploying Red Hat Hyperconverged Infrastructure for Virtualization* : https://access.redhat.com/documentation/en-us/red_hat_hyperconverged_infrastructure_for_virtualization/1.6/html/deploying_red_hat_hypercc

2. Import the backup of the storage domain

From the new Hyperconverged Infrastructure deployment, in the Administration Portal:

- a. Click **Storage** → **Domains**.
- b. Click **Import Domain**. The *Import Pre-Configured Domain* window opens.
- c. In the **Storage Type** field, specify **GlusterFS**.
- d. In the **Name** field, specify a name for the new volume that will be created from the backup volume.
- e. In the **Path** field, specify the path to the backup volume.
- f. Click **OK**. The following warning appears, with any active data centers listed below:

This operation might be unrecoverable and destructive!

Storage Domain(s) are already attached to a Data Center.
Approving this operation might cause data corruption if both Data Centers are active.

- g. Check the **Approve operation** checkbox and click **OK**.
3. Determine a list of virtual machines to import
 - a. Determine the imported domain's identifier by running the following command:

```
# curl -v -k -X GET -u "admin@internal:password" -H "Accept: application/xml"
https://$ENGINE_FQDN/ovirt-engine/api/storagedomains/
```


For example:

```
# curl -v -k -X GET -u "admin@example.com:mybadpassword" -H "Accept: application/xml" https://10.0.2.1/ovirt-engine/api/storagedomains/
```

- b. Determine the list of unregistered disks by running the following command:

```
# curl -v -k -X GET -u "admin@internal:password" -H "Accept: application/xml" "https://$ENGINE_FQDN/ovirt-engine/api/storagedomains/DOMAIN_ID/vms;unregistered"
```

For example:

```
# curl -v -k -X GET -u "admin@example.com:mybadpassword" -H "Accept: application/xml" "https://10.0.2.1/ovirt-engine/api/storagedomains/5e1a37cf-933d-424c-8e3d-eb9e40b690a7/vms;unregistered"
```

4. Perform a partial import of each virtual machine to the storage domain

- a. Determine cluster identifier

The following command returns the cluster identifier.

```
# curl -v -k -X GET -u "admin@internal:password" -H "Accept: application/xml" https://$ENGINE_FQDN/ovirt-engine/api/clusters/
```

For example:

```
# curl -v -k -X GET -u "admin@example.com:mybadpassword" -H "Accept: application/xml" https://10.0.2.1/ovirt-engine/api/clusters/
```

- b. Import the virtual machines

The following command imports a virtual machine without requiring all disks to be available in the storage domain.

```
# curl -v -k -u 'admin@internal:password' -H "Content-type: application/xml" -d '<action> <cluster id="CLUSTER_ID"></cluster> <allow_partial_import>true</allow_partial_import> </action>' "https://ENGINE_FQDN/ovirt-engine/api/storagedomains/DOMAIN_ID/vms/VM_ID/register"
```

For example:

```
# curl -v -k -u 'admin@example.com:mybadpassword' -H "Content-type: application/xml" -d '<action> <cluster id="bf5a9e9e-5b52-4b0d-aeba-4ee4493f1072"></cluster> <allow_partial_import>true</allow_partial_import> </action>' "https://10.0.2.1/ovirt-engine/api/storagedomains/8d21980a-a50b-45e9-9f32-cd8d2424882e/e164f8c6-769a-4cbd-ac2a-ef322c2c5f30/register"
```

For further information, see the Red Hat Virtualization *REST API Guide*:

https://access.redhat.com/documentation/en-us/red_hat_virtualization/4.3/html/rest_api_guide/.

5. Migrate the partially imported disks to the new storage domain

In the Administration Portal, click **Storage** → **Disks**, and Click the **Move Disk** option. Move the imported disks from the synced volume to the replacement cluster's storage domain. For further information, see the Red Hat Virtualization [Administration Guide](#).

6. Attach the restored disks to the new virtual machines
Follow the instructions in the Red Hat Virtualization [Virtual Machine Management Guide](#) to attach the replacement disks to each virtual machine.

14.2. FAILING OVER TO A SECONDARY CLUSTER

This section covers how to fail over from your primary cluster to a remote secondary cluster in the event of server failure.

1. [Configure failover to a remote cluster](#) .
2. Verify that the mapping file for the source and target clusters remains accurate.
3. Run the failover playbook with the **fail_over** tag.

```
# ansible-playbook dr-rhv-failover.yml --tags "fail_over"
```

14.3. FAILING BACK TO A PRIMARY CLUSTER

This section covers how to fail back from your secondary cluster to the primary cluster after you have corrected the cause of a server failure.

1. Prepare the primary cluster for failback by running the cleanup playbook with the **clean_engine** tag.

```
# ansible-playbook dr-cleanup.yml --tags "clean_engine"
```

2. Verify that the mapping file for the source and target clusters remains accurate.
3. Execute failback by running the failback playbook with the **fail_back** tag.

```
# ansible-playbook dr-cleanup.yml --tags "fail_back"
```

14.4. STOPPING A GEO-REPLICATION SESSION USING RHV MANAGER

Stop a geo-replication session when you want to prevent data being replicated from an active source volume to a passive target volume via geo-replication.

1. **Verify that data is not currently being synchronized**
Click the **Tasks** icon at the top right of the Manager, and review the Tasks page.

Ensure that there are no ongoing tasks related to Data Synchronization.

If data synchronization tasks are present, wait until they are complete.

2. **Stop the geo-replication session**
 - a. Click **Storage** → **Volumes**.

- b. Click the name of the volume that you want to prevent geo-replicating.
- c. Click the **Geo-replication** subtab.
- d. Select the session that you want to stop, then click **Stop**.

14.5. TURNING OFF SCHEDULED BACKUPS BY DELETING THE GEO-REPLICATION SCHEDULE

You can stop scheduled backups via geo-replication by deleting the geo-replication schedule.

1. Log in to the Administration Portal on any source node.
2. Click **Storage → Domains**.
3. Click the name of the storage domain that you want to back up.
4. Click the **Remote Data Sync Setup** subtab.
5. Click **Setup**.
The *Setup Remote Data Synchronization* window opens.
6. In the **Recurrence** field, select a recurrence interval type of **NONE** and click **OK**.
7. (Optional) Remove the geo-replication session
Run the following command from the geo-replication master node:

```
# gluster volume geo-replication MASTER_VOL SLAVE_HOST::SLAVE_VOL delete
```

You can also run this command with the **reset-sync-time** parameter. For further information about this parameter and deleting a geo-replication session, see [Deleting a Geo-replication Session](#) in the Red Hat Gluster Storage 3.4 *Administration Guide*.

PART III. REFERENCE MATERIAL

APPENDIX A. FENCING POLICIES FOR RED HAT GLUSTER STORAGE

The following fencing policies are required for Red Hat Hyperconverged Infrastructure for Virtualization (RHHI for Virtualization) deployments. They ensure that hosts are not shut down in situations where brick processes are still running, or when shutting down the host would remove the volume's ability to reach a quorum.

These policies can be set in the **New Cluster** or **Edit Cluster** window in the Administration Portal when Red Hat Gluster Storage functionality is enabled.

Skip fencing if gluster bricks are up

Fencing is skipped if bricks are running and can be reached from other peers.

Skip fencing if gluster quorum not met

Fencing is skipped if bricks are running and shutting down the host will cause loss of quorum

These policies are checked after all other fencing policies when determining whether a node is fenced.

Additional fencing policies may be useful for your deployment. For further details about fencing, see the Red Hat Virtualization *Technical Reference*: https://access.redhat.com/documentation/en-us/red_hat_virtualization/4.3/html/technical_reference/fencing.

APPENDIX B. GLOSSARY OF TERMS

B.1. VIRTUALIZATION TERMS

Administration Portal

A web user interface provided by Red Hat Virtualization Manager, based on the oVirt engine web user interface. It allows administrators to manage and monitor cluster resources like networks, storage domains, and virtual machine templates.

Hosted Engine

The instance of Red Hat Virtualization Manager that manages RHHI for Virtualization.

Hosted Engine virtual machine

The virtual machine that acts as Red Hat Virtualization Manager. The Hosted Engine virtual machine runs on a virtualization host that is managed by the instance of Red Hat Virtualization Manager that is running on the Hosted Engine virtual machine.

Manager node

A virtualization host that runs Red Hat Virtualization Manager directly, rather than running it in a Hosted Engine virtual machine.

Red Hat Enterprise Linux host

A physical machine installed with Red Hat Enterprise Linux plus additional packages to provide the same capabilities as a Red Hat Virtualization host. This type of host is not supported for use with RHHI for Virtualization.

Red Hat Virtualization

An operating system and management interface for virtualizing resources, processes, and applications for Linux and Microsoft Windows workloads.

Red Hat Virtualization host

A physical machine installed with Red Hat Virtualization that provides the physical resources to support the virtualization of resources, processes, and applications for Linux and Microsoft Windows workloads. This is the only type of host supported with RHHI for Virtualization.

Red Hat Virtualization Manager

A server that runs the management and monitoring capabilities of Red Hat Virtualization.

Self-Hosted Engine node

A virtualization host that contains the Hosted Engine virtual machine. All hosts in a RHHI for Virtualization deployment are capable of becoming Self-Hosted Engine nodes, but there is only one Self-Hosted Engine node at a time.

storage domain

A named collection of images, templates, snapshots, and metadata. A storage domain can be comprised of block devices or file systems. Storage domains are attached to data centers in order to provide access to the collection of images, templates, and so on to hosts in the data center.

virtualization host

A physical machine with the ability to virtualize physical resources, processes, and applications for client access.

VM Portal

A web user interface provided by Red Hat Virtualization Manager. It allows users to manage and monitor virtual machines.

B.2. STORAGE TERMS

brick

An exported directory on a server in a trusted storage pool.

cache logical volume

A small, fast logical volume used to improve the performance of a large, slow logical volume.

geo-replication

One way asynchronous replication of data from a source Gluster volume to a target volume. Geo-replication works across local and wide area networks as well as the Internet. The target volume can be a Gluster volume in a different trusted storage pool, or another type of storage.

gluster volume

A logical group of bricks that can be configured to distribute, replicate, or disperse data according to workload requirements.

logical volume management (LVM)

A method of combining physical disks into larger virtual partitions. Physical volumes are placed in volume groups to form a pool of storage that can be divided into logical volumes as needed.

Red Hat Gluster Storage

An operating system based on Red Hat Enterprise Linux with additional packages that provide support for distributed, software-defined storage.

source volume

The Gluster volume that data is being copied from during geo-replication.

storage host

A physical machine that provides storage for client access.

target volume

The Gluster volume or other storage volume that data is being copied to during geo-replication.

thin provisioning

Provisioning storage such that only the space that is required is allocated at creation time, with further space being allocated dynamically according to need over time.

thick provisioning

Provisioning storage such that all space is allocated at creation time, regardless of whether that space is required immediately.

trusted storage pool

A group of Red Hat Gluster Storage servers that recognise each other as trusted peers.

B.3. HYPERCONVERGED INFRASTRUCTURE TERMS

Red Hat Hyperconverged Infrastructure (RHHI) for Virtualization

RHHI for Virtualization is a single product that provides both virtual compute and virtual storage resources. Red Hat Virtualization and Red Hat Gluster Storage are installed in a converged configuration, where the services of both products are available on each physical machine in a cluster.

hyperconverged host

A physical machine that provides physical storage, which is virtualized and consumed by virtualized processes and applications run on the same host. All hosts installed with RHHI for Virtualization are hyperconverged hosts.

Web Console

The web user interface for deploying, managing, and monitoring RHHI for Virtualization. The Web Console is provided by the the Web Console service and plugins for Red Hat Virtualization Manager.

