



Red Hat Hyperconverged Infrastructure for Virtualization 1.6

Deploying Red Hat Hyperconverged Infrastructure for Virtualization

Instructions for deploying Red Hat Hyperconverged Infrastructure for Virtualization

Red Hat Hyperconverged Infrastructure for Virtualization 1.6 Deploying Red Hat Hyperconverged Infrastructure for Virtualization

Instructions for deploying Red Hat Hyperconverged Infrastructure for Virtualization

Laura Bailey

lbailey@redhat.com

Legal Notice

Copyright © 2021 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, the Red Hat logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux[®] is the registered trademark of Linus Torvalds in the United States and other countries.

Java[®] is a registered trademark of Oracle and/or its affiliates.

XFS[®] is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL[®] is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js[®] is an official trademark of Joyent. Red Hat is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack[®] Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

This document outlines how to deploy Red Hat Hyperconverged Infrastructure for Virtualization (RHVI for Virtualization) across three physical machines, using Red Hat Gluster Storage 3.4 and Red Hat Virtualization 4.3. This creates a discrete cluster for use in remote office branch office (ROBO) environments, where a remote office synchronizes data to a central data center on a regular basis, but can remain fully functional if connectivity to the central data center is lost.

Table of Contents

PART I. PLAN	4
CHAPTER 1. ARCHITECTURE	5
1.1. UNDERSTANDING VDO	5
CHAPTER 2. SUPPORT REQUIREMENTS	7
2.1. OPERATING SYSTEM	7
2.2. PHYSICAL MACHINES	7
2.3. VIRTUAL MACHINES	8
2.4. HOSTED ENGINE VIRTUAL MACHINE	8
2.5. NETWORKING	8
2.6. STORAGE	9
2.6.1. Disks	10
2.6.2. RAID	10
2.6.3. JBOD	11
2.6.4. Logical volumes	11
2.6.5. Red Hat Gluster Storage volumes	11
2.6.6. Volume types	11
2.7. VIRTUAL DATA OPTIMIZER (VDO)	12
2.8. SCALING	12
2.9. EXISTING RED HAT GLUSTER STORAGE CONFIGURATIONS	12
2.10. DISASTER RECOVERY	12
2.10.1. Prerequisites for geo-replication	13
2.10.2. Prerequisites for failover and failback configuration	13
2.11. ADDITIONAL REQUIREMENTS FOR SINGLE NODE DEPLOYMENTS	13
CHAPTER 3. RECOMMENDATIONS	14
3.1. GENERAL RECOMMENDATIONS	14
3.2. SECURITY RECOMMENDATIONS	14
3.3. HOST RECOMMENDATIONS	14
3.4. NETWORKING RECOMMENDATIONS	15
3.4.1. Recommended practices for configuring host networks	15
3.5. SELF-HOSTED ENGINE RECOMMENDATIONS	16
PART II. DEPLOY	17
CHAPTER 4. DEPLOYMENT WORKFLOW	18
4.1. INSTALL HOST PHYSICAL MACHINES	18
4.1.1. Installing Red Hat Virtualization Host	18
4.1.2. Enabling software repositories	20
CHAPTER 5. CONFIGURE PUBLIC KEY BASED SSH AUTHENTICATION WITHOUT A PASSWORD	21
5.1. ADDING KNOWN HOSTS TO THE FIRST HOST	21
5.2. GENERATING SSH KEY PAIRS WITHOUT A PASSWORD	22
5.3. COPYING SSH KEYS	23
CHAPTER 6. CONFIGURE RED HAT GLUSTER STORAGE FOR HOSTED ENGINE USING THE WEB CONSOLE	24
CHAPTER 7. DEPLOY THE HOSTED ENGINE USING THE WEB CONSOLE	31
CHAPTER 8. CONFIGURE RED HAT GLUSTER STORAGE AS A RED HAT VIRTUALIZATION STORAGE DOMAIN	40
8.1. CREATE THE LOGICAL NETWORK FOR GLUSTER TRAFFIC	40

8.2. CONFIGURE ADDITIONAL HYPERCONVERGED HOSTS	40
PART III. VERIFY	42
CHAPTER 9. VERIFY YOUR DEPLOYMENT	43
PART IV. NEXT STEPS	46
CHAPTER 10. POST-DEPLOYMENT CONFIGURATION SUGGESTIONS	47
10.1. CONFIGURE FENCING FOR HIGH AVAILABILITY	47
10.2. CONFIGURE BACKUP AND RECOVERY OPTIONS	47
PART V. TROUBLESHOOT	48
CHAPTER 11. LOG FILE LOCATIONS	49
CHAPTER 12. VIEWING DEPLOYMENT ERRORS	50
12.1. FAILED TO DEPLOY STORAGE	50
12.2. FAILED TO PREPARE VIRTUAL MACHINE	50
12.3. FAILED TO DEPLOY HOSTED ENGINE	51
PART VI. REFERENCE MATERIAL	53
APPENDIX A. GLOSSARY OF TERMS	54
A.1. VIRTUALIZATION TERMS	54
A.2. STORAGE TERMS	54
A.3. HYPERCONVERGED INFRASTRUCTURE TERMS	55

PART I. PLAN

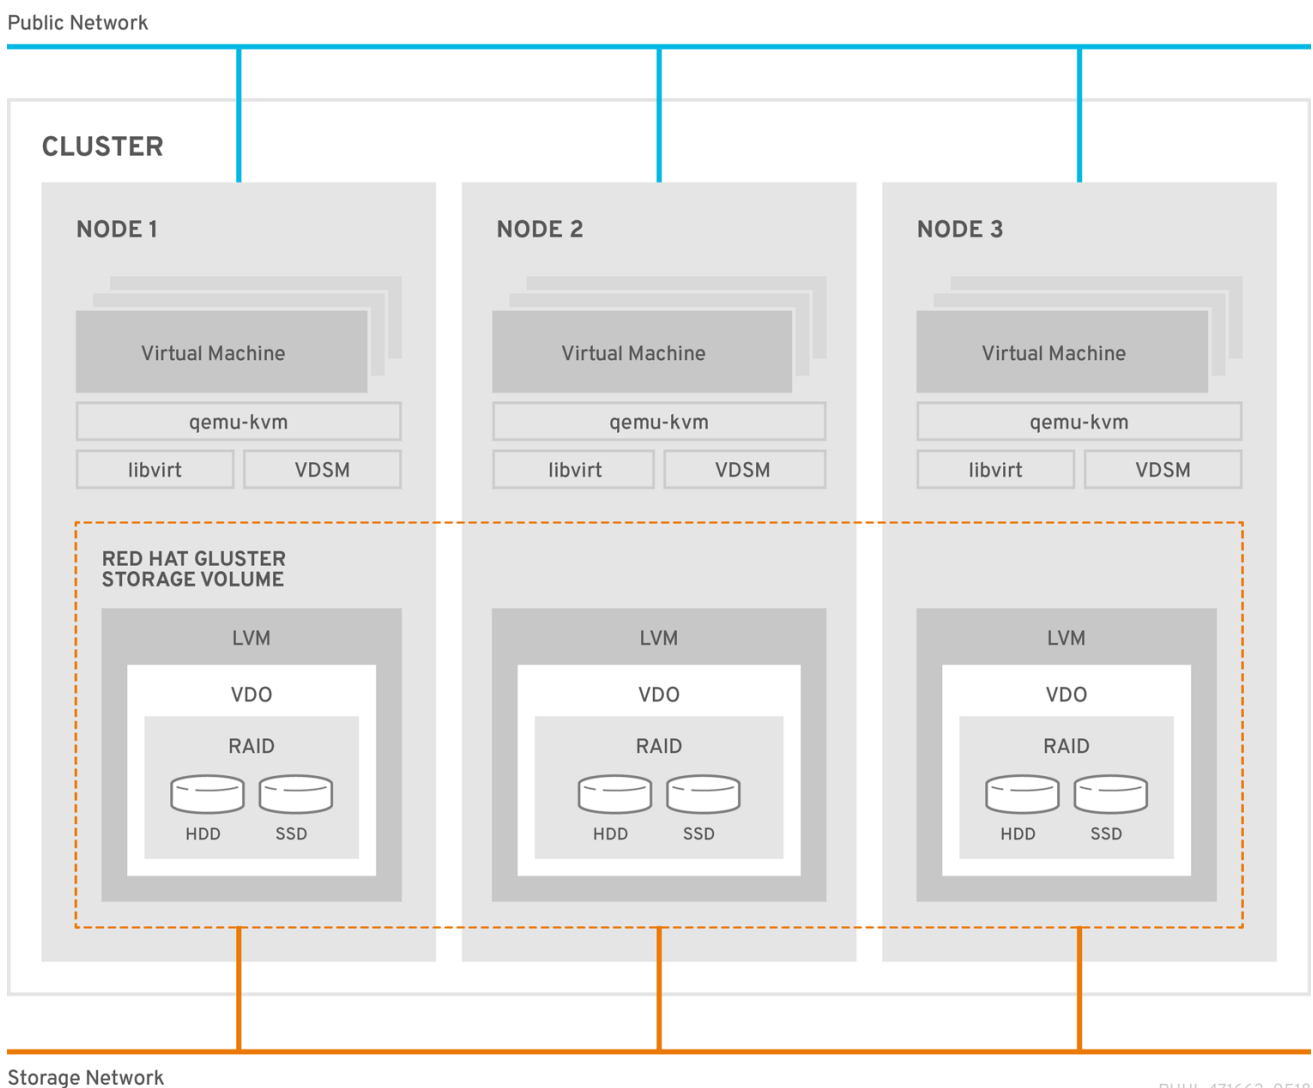
CHAPTER 1. ARCHITECTURE

Red Hat Hyperconverged Infrastructure for Virtualization (RHHI for Virtualization) combines compute, storage, networking, and management capabilities in one deployment.

RHHI for Virtualization is deployed across three physical machines to create a discrete cluster or *pod* using Red Hat Gluster Storage 3.4 and Red Hat Virtualization 4.3.

The dominant use case for this deployment is in remote office branch office (ROBO) environments, where a remote office synchronizes data to a central data center on a regular basis, but does not require connectivity to the central data center to function.

The following diagram shows the basic architecture of a single cluster.



RHHI_471663_0518

1.1. UNDERSTANDING VDO

As of Red Hat Hyperconverged Infrastructure for Virtualization 1.6, you can configure a Virtual Data Optimizer (VDO) layer to provide data reduction and deduplication for your storage.

VDO is supported only when enabled on new installations at deployment time, and cannot be enabled on deployments upgraded from earlier versions of RHHI for Virtualization.

VDO performs following types of data reduction to reduce the space required by data.

Deduplication

Eliminates zero and duplicate data blocks. VDO finds duplicated data using the UDS (Universal Deduplication Service) Kernel Module. Instead of writing the duplicated data, VDO records it as a reference to the original block. The logical block address is mapped to the physical block address by VDO.

Compression

Reduces the size of the data by packing non-duplicate blocks together into fixed length (4 KB) blocks before writing to disk. This helps to speed up the performance for reading data from storage.

At best, data can be reduced to 15% of its original size.

Because reducing data has additional processing costs, enabling compression and deduplication reduces write performance. As a result, VDO is not recommended for performance sensitive workloads. Red Hat strongly recommends that you test and verify that your workload achieves the required level of performance with VDO enabled before deploying VDO in production.

CHAPTER 2. SUPPORT REQUIREMENTS

Review this section to ensure that your planned deployment meets the requirements for support by Red Hat.

2.1. OPERATING SYSTEM

Red Hat Hyperconverged Infrastructure for Virtualization (RHHI for Virtualization) uses Red Hat Virtualization Host 4.3 as a base for all other configuration. Red Hat Enterprise Linux hosts are not supported.

The following table shows the the supported versions of each product to use for a supported RHHI for Virtualization deployment.

Table 2.1. Version compatibility

RHHI version	RHGS version	RHV version
1.0	3.2	4.1.0 to 4.1.7
1.1	3.3.1	4.1.8 to 4.2.0
1.5	3.4 Batch 1 Update	4.2.7
1.5.1	3.4 Batch 2 Update	4.2.8
1.6	3.4 Batch 4 Update	4.3 to current

See [Requirements](#) in the Red Hat Virtualization *Planning and Prerequisites Guide* for details on requirements of Red Hat Virtualization.

2.2. PHYSICAL MACHINES

Red Hat Hyperconverged Infrastructure for Virtualization (RHHI for Virtualization) requires **at least 3 physical machines**. Scaling to 6, 9, or 12 physical machines is also supported; see [Scaling](#) for more detailed requirements.

Each physical machine must have the following capabilities:

- at least 2 NICs (Network Interface Controllers) per physical machine, for separation of data and management traffic (see [Section 2.5, “Networking”](#) for details)
- for small deployments:
 - at least 12 cores
 - at least 64GB RAM
 - at most 48TB storage
- for medium deployments:
 - at least 12 cores

- at least 128GB RAM
- at most 64TB storage
- for large deployments:
 - at least 16 cores
 - at least 256GB RAM
 - at most 80TB storage

2.3. VIRTUAL MACHINES

The number of virtual machines that you are able to run on your hyperconverged deployment depends greatly on what those virtual machines do, and what load they are under. Test your workload's CPU, memory, and throughput requirements and provision your hyperconverged environment accordingly.

2.4. HOSTED ENGINE VIRTUAL MACHINE

The Hosted Engine virtual machine requires at least the following:


- 1 dual core CPU (1 quad core or multiple dual core CPUs recommended)
- 4GB RAM that is not shared with other processes (16GB recommended)
- 25GB of local, writable disk space (50GB recommended)
- 1 NIC with at least 1Gbps bandwidth

For more information, see [Requirements](#) in the Red Hat Virtualization 4.3 *Planning and Prerequisites Guide*.

2.5. NETWORKING

Fully-qualified domain names that are forward and reverse resolvable by DNS are required for all hyperconverged hosts and for the Hosted Engine virtual machine that provides Red Hat Virtualization Manager.

IPv6 is supported as a Technology Preview in IPv6-only environments (including DNS and gateway addresses). Environments with both IPv4 and IPv6 addresses are not supported.



IMPORTANT

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

Client storage traffic and management traffic in the cluster must use separate networks: a **front-end management network** and a **back-end storage network**.

Each node requires two Ethernet ports, one for each network This ensures optimal performance. For high availability, place each network on a separate network switch. For improved fault tolerance, provide a separate power supply for each switch.

Front-end management network

- Used by Red Hat Virtualization and virtual machines.
- Requires at least one 1Gbps Ethernet connection.
- IP addresses assigned to this network must be on the same subnet as each other, and on a different subnet to the back-end storage network.
- IP addresses on this network can be selected by the administrator.

Back-end storage network

- Used by storage and migration traffic between hyperconverged nodes.
- Requires at least one 10Gbps Ethernet connection.
- Requires maximum latency of 5 milliseconds between peers.

Network fencing devices that use Intelligent Platform Management Interfaces (IPMI) require a separate network.

If you want to use DHCP network configuration for the Hosted Engine virtual machine, then you must have a DHCP server configured prior to configuring Red Hat Hyperconverged Infrastructure for Virtualization.

If you want to configure disaster recovery by using geo-replication to store copies of data:

- Configure a reliable time source.
- Do not use IPv6 addresses.



WARNING

[Bug 1688239](#) currently prevents IPv6 based geo-replication from working correctly. Do not use IPv6 addresses if you require disaster recovery functionality using geo-replication.

Before you begin the deployment process, determine the following details:

- IP address for a gateway to the hyperconverged host. This address must respond to ping requests.
- IP address of the front-end management network.
- Fully-qualified domain name (FQDN) for the Hosted Engine virtual machine.
- MAC address that resolves to the static FQDN and IP address of the Hosted Engine.

2.6. STORAGE

A hyperconverged host stores configuration, logs and kernel dumps, and uses its storage as swap space. This section lists the minimum directory sizes for hyperconverged hosts. Red Hat recommends using the default allocations, which use more storage space than these minimums.

- `/` (root) - 6GB
- `/home` - 1GB
- `/tmp` - 1GB
- `/boot` - 1GB
- `/var` - 15GB
- `/var/crash` - 10GB
- `/var/log` - 8GB



IMPORTANT

Red Hat recommends increasing the size of `/var/log` to at least 15GB to provide sufficient space for the additional logging requirements of Red Hat Gluster Storage.

Follow the instructions in [Growing a logical volume using the Web Console](#) to increase the size of this partition after installing the operating system.

- `/var/log/audit` - 2GB
- `swap` - 1GB (see [Recommended swap size](#) for details)
- Anaconda reserves 20% of the thin pool size within the volume group for future metadata expansion. This is to prevent an out-of-the-box configuration from running out of space under normal usage conditions. Overprovisioning of thin pools during installation is also not supported.
- **Minimum Total - 55GB**

2.6.1. Disks

Red Hat recommends Solid State Disks (SSDs) for best performance. If you use Hard Drive Disks (HDDs), you should also configure a smaller, faster SSD as an LVM cache volume.

4K native devices are not supported with Red Hat Hyperconverged Infrastructure for Virtualization, as Red Hat Virtualization requires 512 byte emulation (512e) support.

2.6.2. RAID

RAID5 and RAID6 configurations are supported. However, RAID configuration limits depend on the technology in use.

- SAS/SATA 7k disks are supported with RAID6 (at most 10+2)
- SAS 10k and 15k disks are supported with the following:
 - RAID5 (at most 7+1)

- RAID6 (at most 10+2)

RAID cards must use flash backed write cache.

Red Hat further recommends providing at least one hot spare drive local to each server.

2.6.3. JBOD

As of Red Hat Hyperconverged Infrastructure for Virtualization 1.6, JBOD configurations are fully supported and no longer require architecture review.

2.6.4. Logical volumes

The logical volumes that comprise the **engine** gluster volume must be thick provisioned. This protects the Hosted Engine from out of space conditions, disruptive volume configuration changes, I/O overhead, and migration activity.

The logical volumes that comprise the **vmstore** and optional **data** gluster volumes must be thin provisioned. This allows greater flexibility in underlying volume configuration. If your thin provisioned volumes are on Hard Drive Disks (HDDs), configure a smaller, faster Solid State Disk (SSD) as an lvmcache for improved performance.

2.6.5. Red Hat Gluster Storage volumes

Red Hat Hyperconverged Infrastructure for Virtualization is expected to have 3–4 Red Hat Gluster Storage volumes.

- 1 **engine** volume for the Hosted Engine
- 1 **vmstore** volume for virtual machine operating system disk images
- 1 optional **data** volume for other virtual machine disk images
- 1 **shared_storage** volume for geo-replication metadata

Separate **vmstore** and **data** volumes are recommended to minimize backup storage requirements. Storing virtual machine data separate from operating system images means that only the **data** volume needs to be backed up when storage space is at a premium, since operating system images on the **vmstore** volume can be more easily rebuilt.

A Red Hat Hyperconverged Infrastructure for Virtualization deployment can contain at most 1 geo-replicated volume.

2.6.6. Volume types

Red Hat Hyperconverged Infrastructure for Virtualization (RHVI for Virtualization) supports only the following volume types at deployment time:

- **Replicated volumes** (3 copies of the same data on 3 bricks, across 3 nodes).
These volumes can be expanded into distributed-replicated volumes after deployment.
- **Arbitrated replicated volumes** (2 full copies of the same data on 2 bricks and 1 arbiter brick that contains metadata, across three nodes).
These volumes can be expanded into arbitrated distributed-replicated volumes after deployment.

- [Distributed volumes](#) (1 copy of the data, no replication to other bricks).

Note that arbiter bricks store only file names, structure, and metadata. This means that a three-way arbitrated replicated volume requires about 75% of the storage space that a three-way replicated volume would require to achieve the same level of consistency. However, because the arbiter brick stores only metadata, a three-way arbitrated replicated volume only provides the availability of a two-way replicated volume.

For more information on laying out arbitrated replicated volumes, see [Creating multiple arbitrated replicated volumes across fewer total nodes](#) in the Red Hat Gluster Storage *Administration Guide*.

2.7. VIRTUAL DATA OPTIMIZER (VDO)

A Virtual Data Optimizer (VDO) layer is supported as of Red Hat Hyperconverged Infrastructure for Virtualization 1.6.

VDO support is limited to new deployments only; do not attempt to add a VDO layer to an existing deployment.

2.8. SCALING

Initial deployments of Red Hat Hyperconverged Infrastructure for Virtualization are either 1 node or 3 nodes.

1 node deployments cannot be scaled.

3 node deployments can be scaled to 6, 9, or 12 nodes using one of the following methods:

1. Add new hyperconverged nodes to the cluster, in sets of three, up to the maximum of 12 hyperconverged nodes.
2. Create new Gluster volumes using new disks on new or existing nodes.
3. Expand existing Gluster volumes to span 6, 9, or 12 nodes using new disks on new or existing nodes.

You cannot create a volume that spans more than 3 nodes at creation time; you must create a 3-node volume first and then expand it across more nodes as necessary.

2.9. EXISTING RED HAT GLUSTER STORAGE CONFIGURATIONS

Red Hat Hyperconverged Infrastructure for Virtualization is supported only when deployed as specified in this document. Existing Red Hat Gluster Storage configurations cannot be used in a hyperconverged configuration. If you want to use an existing Red Hat Gluster Storage configuration, refer to the traditional configuration documented in [Configuring Red Hat Virtualization with Red Hat Gluster Storage](#).

2.10. DISASTER RECOVERY

Red Hat strongly recommends configuring a disaster recovery solution. For details on configuring geo-replication as a disaster recovery solution, see *Maintaining Red Hat Hyperconverged Infrastructure for Virtualization*: https://access.redhat.com/documentation/en-us/red_hat_hyperconverged_infrastructure_for_virtualization/1.6/html/maintaining_red_hat_hyperconverged_backup-recovery.

**WARNING**

[Bug 1688239](#) currently prevents IPv6 based geo-replication from working correctly. Do not use IPv6 addresses if you require disaster recovery functionality using geo-replication.

2.10.1. 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.

2.10.2. 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.

2.11. ADDITIONAL REQUIREMENTS FOR SINGLE NODE DEPLOYMENTS

Red Hat Hyperconverged Infrastructure for Virtualization is supported for deployment on a single node provided that all [Support Requirements](#) are met, with the following additions and exceptions.

A single node deployment requires a physical machine with:

- 1 Network Interface Controller
- at least 12 cores
- at least 64GB RAM
- at most 48TB storage

Single node deployments cannot be scaled, and are not highly available.

CHAPTER 3. RECOMMENDATIONS

The configuration described in this section is not required, but may improve the stability or performance of your deployment.

3.1. GENERAL RECOMMENDATIONS

- Take a full backup as soon as deployment is complete, and store the backup in a separate location. Take regular backups thereafter. See [Configuring backup and recovery options](#) for details.
- Avoid running any service that your deployment depends on as a virtual machine in the same RHHI for Virtualization environment. If you must run a required service in the same deployment, carefully plan your deployment to minimize the downtime of the virtual machine running the required service.
- Ensure that hyperconverged hosts have sufficient entropy. Failures can occur when the value in `/proc/sys/kernel/random/entropy_avail` is less than **200**. To increase entropy, install the **rng-tools** package and follow the steps in <https://access.redhat.com/solutions/1395493>.
- Document your environment so that everyone who works with it is aware of its current state and required procedures.

3.2. SECURITY RECOMMENDATIONS

- Do not disable any security features (such as HTTPS, SELinux, and the firewall) on the hosts or virtual machines.
- Register all hosts and Red Hat Enterprise Linux virtual machines to either the Red Hat Content Delivery Network or Red Hat Satellite in order to receive the latest security updates and errata.
- Create individual administrator accounts, instead of allowing many people to use the default admin account, for proper activity tracking.
- Limit access to the hosts and create separate logins. Do not create a single root login for everyone to use. See [Managing Users and Groups](#) in the Red Hat Enterprise Linux 7 System Administrator's Guide.
- Do not create untrusted users on hosts.
- Avoid installing additional packages such as analyzers, compilers, or other components that add unnecessary security risk.

3.3. HOST RECOMMENDATIONS

- Standardize the hosts in the same cluster. This includes having consistent hardware models and firmware versions. Mixing different server hardware within the same cluster can result in inconsistent performance from host to host.
- Configure fencing devices at deployment time. Fencing devices are required for high availability.
- Use separate hardware switches for fencing traffic. If monitoring and fencing go over the same switch, that switch becomes a single point of failure for high availability.

3.4. NETWORKING RECOMMENDATIONS

- Bond network interfaces, especially on production hosts. Bonding improves the overall availability of service, as well as network bandwidth. See [Network Bonding](#) in the Administration Guide.
- For optimal performance and simplified troubleshooting, use VLANs to separate different traffic types and make the best use of 10 GbE or 40 GbE networks.
- If the underlying switches support jumbo frames, set the MTU to the maximum size (for example, **9000**) that the underlying switches support. This setting enables optimal throughput, with higher bandwidth and reduced CPU usage, for most applications. The default MTU is determined by the minimum size supported by the underlying switches. If you have LLDP enabled, you can see the MTU supported by the peer of each host in the NIC's tool tip in the **Setup Host Networks** window.
- 1 GbE networks should only be used for management traffic. Use 10 GbE or 40 GbE for virtual machines and Ethernet-based storage.
- If additional physical interfaces are added to a host for storage use, uncheck **VM network** so that the VLAN is assigned directly to the physical interface.

3.4.1. Recommended practices for configuring host networks

If your network environment is complex, you may need to configure a host network manually before adding the host to Red Hat Virtualization Manager.

Red Hat recommends the following practices for configuring a host network:

- Configure the network with the Web Console. Alternatively, you can use `nmtui` or `nmcli`.
- If a network is not required for a self-hosted engine deployment or for adding a host to the Manager, configure the network in the Administration Portal after adding the host to the Manager. See [Creating a New Logical Network in a Data Center or Cluster](#).
- Use the following naming conventions:
 - VLAN devices: **VLAN_NAME_TYPE_RAW_PLUS_VID_NO_PAD**
 - VLAN interfaces: **physical_device.VLAN_ID** (for example, **eth0.23**, **eth1.128**, **enp3s0.50**)
 - Bond interfaces: **bondnumber** (for example, **bond0**, **bond1**)
 - VLANs on bond interfaces: **bondnumber.VLAN_ID** (for example, **bond0.50**, **bond1.128**)
- Use [network bonding](#). Networking teaming is not supported.
- Use recommended bonding modes:
 - If the **ovirtmgmt** network is not used by virtual machines, the network may use any supported bonding mode.
 - If the **ovirtmgmt** network is used by virtual machines, see [Which bonding modes work when used with a bridge that virtual machine guests or containers connect to?](#)

- Red Hat Virtualization's default bonding mode is **(Mode 4) Dynamic Link Aggregation**. If your switch does not support Link Aggregation Control Protocol (LACP), use **(Mode 1) Active-Backup**. See [Bonding Modes](#) for details.
- Configure a VLAN on a physical NIC as in the following example (although **nmcli** is used, you can use any tool):

```
# nmcli connection add type vlan con-name vlan50 ifname eth0.50 dev eth0 id 50
# nmcli con mod vlan50 +ipv4.dns 8.8.8.8 +ipv4.addresses 123.123.0.1/24 +ipv4.gateway
123.123.0.254
```

- Configure a VLAN on a bond as in the following example (although **nmcli** is used, you can use any tool):

```
# nmcli connection add type bond con-name bond0 ifname bond0 bond.options
"mode=active-backup,miimon=100" ipv4.method disabled ipv6.method ignore
# nmcli connection add type ethernet con-name eth0 ifname eth0 master bond0 slave-type
bond
# nmcli connection add type ethernet con-name eth1 ifname eth1 master bond0 slave-type
bond
# nmcli connection add type vlan con-name vlan50 ifname bond0.50 dev bond0 id 50
# nmcli con mod vlan50 +ipv4.dns 8.8.8.8 +ipv4.addresses 123.123.0.1/24 +ipv4.gateway
123.123.0.254
```

- Do not disable **firewalld**.
- Customize the firewall rules in the Administration Portal after adding the host to the Manager. See [Configuring Host Firewall Rules](#).



IMPORTANT

When creating a management bridge that uses a static IPv6 address, disable network manager control in its interface configuration (ifcfg) file before adding a host. See <https://access.redhat.com/solutions/3981311> for more information.

3.5. SELF-HOSTED ENGINE RECOMMENDATIONS

- Create a separate data center and cluster for the Red Hat Virtualization Manager and other infrastructure-level services, if the environment is large enough to allow it. Although the Manager virtual machine can run on hosts in a regular cluster, separation from production virtual machines helps facilitate backup schedules, performance, availability, and security.
- A storage domain dedicated to the Manager virtual machine is created during self-hosted engine deployment. Do not use this storage domain for any other virtual machines.
- All self-hosted engine nodes should have an equal CPU family so that the Manager virtual machine can safely migrate between them. If you intend to have various families, begin the installation with the lowest one.
- If the Manager virtual machine shuts down or needs to be migrated, there must be enough memory on a self-hosted engine node for the Manager virtual machine to restart on or migrate to it.

PART II. DEPLOY

CHAPTER 4. DEPLOYMENT WORKFLOW

The workflow for deploying Red Hat Hyperconverged Infrastructure for Virtualization (RHVI for Virtualization) is as follows:

1. Verify that your planned deployment meets support requirements: [Chapter 2, Support requirements](#).
2. Install the physical machines that will act as hyperconverged hosts: [Section 4.1, “Install host physical machines”](#).
3. Configure key-based SSH authentication without a password to enable automated configuration of the hosts: [Chapter 5, Configure Public Key based SSH Authentication without a password](#).
4. Configure Red Hat Gluster Storage on the physical hosts using the Web Console: [Chapter 6, Configure Red Hat Gluster Storage for Hosted Engine using the Web Console](#) .
5. Deploy the Hosted Engine using the Web Console: [Chapter 7, Deploy the Hosted Engine using the Web Console](#).
6. Configure the Red Hat Gluster Storage nodes using the Administration Portal: [Log in to the Administration Portal to complete configuration](#).

4.1. INSTALL HOST PHYSICAL MACHINES

Your physical machines need an operating system and access to the appropriate software repositories in order to be used as hyperconverged hosts.

1. Install Red Hat Virtualization Host on each physical machine.
2. Enable the Red Hat Virtualization Host software repository on each physical machine.

4.1.1. Installing Red Hat Virtualization Host

Red Hat Virtualization Host is a minimal operating system designed for setting up a physical machine that acts as a hypervisor in Red Hat Virtualization, or a hyperconverged host in Red Hat Hyperconverged Infrastructure.

Prerequisites

- Ensure that your physical machine meets the requirements outlined in [Physical machines](#).

Procedure

1. Download the Red Hat Virtualization Host ISO image from the Customer Portal:
 - a. Log in to the Customer Portal at <https://access.redhat.com>.
 - b. Click **Downloads** in the menu bar.
 - c. Click **Red Hat Virtualization**. Scroll up and click **Download Latest** to access the product download page.
 - d. Go to *Hypervisor Image for RHV 4.3* and click **Download Now**.


- e. Create a bootable media device. See [Making Media](#) in the *Red Hat Enterprise Linux Installation Guide* for more information.
2. Start the machine on which you are installing Red Hat Virtualization Host, and boot from the prepared installation media.
3. From the boot menu, select **Install RHVH 4.3** and press **Enter**.



NOTE

You can also press the Tab key to edit the kernel parameters. Kernel parameters must be separated by a space, and you can boot the system using the specified kernel parameters by pressing the Enter key. Press the Esc key to clear any changes to the kernel parameters and return to the boot menu.

4. Select a language, and click **Continue**.
5. Select a time zone from the *Date & Time* screen and click **Done**.



IMPORTANT

Red Hat recommends using Coordinated Universal Time (UTC) on all hosts. This helps ensure that data collection and connectivity are not impacted by variation in local time, such as during daylight savings time.

6. Select a keyboard layout from the *Keyboard* screen and click **Done**.
7. Specify the installation location from the *Installation Destination* screen.



IMPORTANT

- Red Hat strongly recommends using the **Automatically configure partitioning** option.
- All disks are selected by default, so deselect disks that you do not want to use as installation locations.
- At-rest encryption is not supported. Do not enable encryption.
- Red Hat recommends increasing the size of **/var/log** to at least 15GB to provide sufficient space for the additional logging requirements of Red Hat Gluster Storage.
Follow the instructions in [Growing a logical volume using the Web Console](#) to increase the size of this partition after installing the operating system.

Click **Done**.

8. Select the Ethernet network from the *Network & Host Name* screen.
 - a. Click **Configure...** → **General** and select the **Automatically connect to this network when it is available** check box.
9. Optionally configure **Language Support**, **Security Policy**, and **Kdump**. See [Installing Using Anaconda](#) in the *Red Hat Enterprise Linux 7 Installation Guide* for more information on each of the sections in the *Installation Summary* screen.

10. Click **Begin Installation**.
11. Set a root password and, optionally, create an additional user while Red Hat Virtualization Host installs.



WARNING

Red Hat strongly recommends not creating untrusted users on Red Hat Virtualization Host, as this can lead to exploitation of local security vulnerabilities.

12. Click **Reboot** to complete the installation.



NOTE

When Red Hat Virtualization Host restarts, **nodectl check** performs a health check on the host and displays the result when you log in on the command line. The message **node status: OK** or **node status: DEGRADED** indicates the health status. Run **nodectl check** to get more information. The service is enabled by default.

4.1.2. Enabling software repositories

1. Log in to the Web Console.
Use the management FQDN and port **9090**, for example, **https://server1.example.com:9090/**.
2. Navigate to **Subscriptions**, click **Register System**, and enter your Customer Portal user name and password.
The **Red Hat Virtualization Host** subscription is automatically attached to the system.
3. Click **Terminal**.
4. Enable the **Red Hat Virtualization Host 7** repository to allow later updates to the Red Hat Virtualization Host:

```
# subscription-manager repos --enable=rhel-7-server-rhvh-4-rpms
```


CHAPTER 5. CONFIGURE PUBLIC KEY BASED SSH AUTHENTICATION WITHOUT A PASSWORD

Configure public key based SSH authentication without a password for the root user on the first hyperconverged host to all hosts, **including itself**. Do this for all storage and management interfaces, and for both IP addresses and FQDNs.

5.1. ADDING KNOWN HOSTS TO THE FIRST HOST

When you use SSH to log in to a host from a system that is not already known to the host, you are prompted to add that system as a known host.

1. Log in to the first hyperconverged host as the root user.
2. Perform the following steps for each host in the cluster, including the first host.
 - a. Use SSH to log in to a host as the root user.

```
[root@server1]# ssh root@server1.example.com
```

- b. Enter **yes** to continue connecting.

```
[root@server1]# ssh root@server2.example.com
The authenticity of host 'server2.example.com (192.51.100.28)' can't be established.
ECDSA key fingerprint is SHA256:Td8KqgVIPXdTlasdfa2xRwn3/asdBasdpnaGM.
Are you sure you want to continue connecting (yes/no)?
```

This automatically adds the host key of the first host to the **known_hosts** file on the target host.

```
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '192.51.100.28' (ECDSA) to the list of known hosts.
```

- c. Enter the password for the root user on the target host to complete the login process.

```
root@server2.example.com's password: *****
Last login: Mon May 27 10:04:49 2019
[root@server2]#
```

- d. Log out of the host.

```
[root@server2]# exit
[root@server1]#
```



NOTE

When you log out of the SSH session from the first host to itself, the user and server in the command line prompt stay the same; it is only the session that changes.

```
[root@server1]# exit
[root@server1]#
```

5.2. GENERATING SSH KEY PAIRS WITHOUT A PASSWORD

Generating a public/private key pair lets you use key-based SSH authentication. Generating a key pair that does not use a password makes it simpler to use Ansible to automate deployment and configuration processes.

Procedure

1. Log in to the first hyperconverged host as the root user.
2. Generate an SSH key that does not use a password.
 - a. Start the key generation process.

```
# ssh-keygen -t rsa
Generating public/private rsa key pair.
```

- b. Enter a location for the key.
The default location, shown in parentheses, is used if no other input is provided.

```
Enter file in which to save the key (/home/username/.ssh/id_rsa): <location>/<keyname>
```

- c. Specify and confirm an empty passphrase by pressing **Enter** twice.

```
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
```

The private key is saved in **<location>/<keyname>**. The public key is saved in **<location>/<keyname>.pub**.

```
Your identification has been saved in <location>/<keyname>.
Your public key has been saved in <location>/<keyname>.pub.
The key fingerprint is SHA256:8BhZageKrLXM99z5f/AM9aPo/KAUd8ZZFPcPFWqK6+M
root@server1.example.com
The key's randomart image is:
+---[ECDSA 256]---+
|  . .  +=|
| ... =  0.0|
| + . * .  0...|
| = . . * . + +..|
|. + . . So o * ..|
| . o . . + = ..|
|   o oo ..= . .|
|   ooo...+ |
|   .E++oo |
+----[SHA256]-----+
```

**WARNING**

Your identification in this output is your private key. Never share your private key. Possession of your private key allows someone else to impersonate you on any system that has your public key.

5.3. COPYING SSH KEYS

To access a host using your private key, that host needs a copy of your public key.

Prerequisites

- Generate a public/private key pair.
- SSH access from the root user on the host to all storage and management interfaces on the same host, using both IP addresses and FQDNs.

Procedure

1. Log in to the first host as the root user.
2. Copy your public key to the host that you want to access.

```
# ssh-copy-id -i <location>/<keyname>.pub <user>@<hostname>
```

Enter the password for **<user>@<hostname>** if prompted.

**WARNING**

Make sure that you use the file that ends in **.pub**. Never share your private key. Possession of your private key allows someone else to impersonate you on any system that has your public key.

CHAPTER 6. CONFIGURE RED HAT GLUSTER STORAGE FOR HOSTED ENGINE USING THE WEB CONSOLE



IMPORTANT

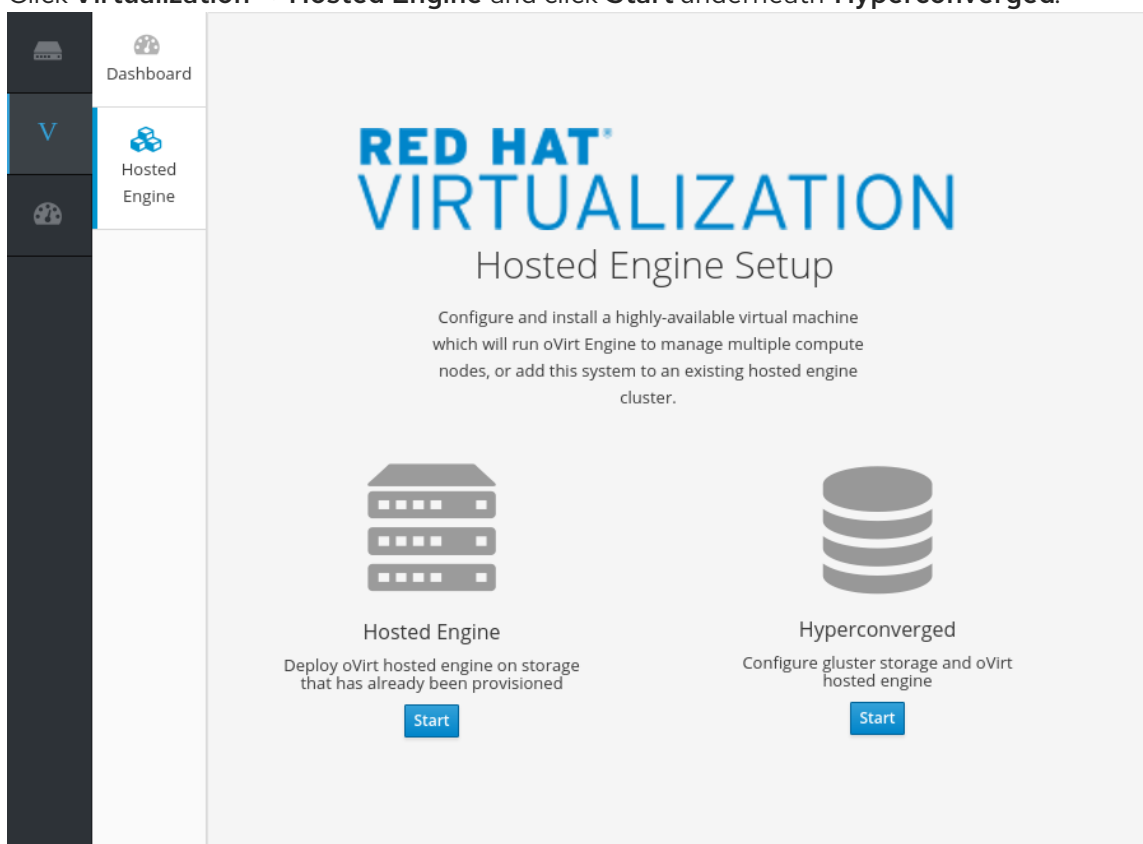
Ensure that disks specified as part of this deployment process do not have any partitions or labels.

1. Log into the Web Console

Browse to the the Web Console management interface of the first hyperconverged host, for example, <https://node1.example.com:9090/>, and log in with the credentials you created in [Section 4.1, "Install host physical machines"](#).


2. Start the deployment wizard

- a. Click **Virtualization** → **Hosted Engine** and click **Start** underneath **Hyperconverged**.



The *Gluster Configuration* window opens.


- b. Click the **Run Gluster Wizard** button.



The *Gluster Deployment* window opens in 3 node mode.

3. Specify hyperconverged hosts

Specify the back-end Additional Hosts on the storage network (not the management network) of the three hyperconverged hosts. The hyperconverged host that can SSH using key pairs should be listed first, as it is the host that will run deployment tasks and the hosted engine.



NOTE

If you plan to create an arbitrated replicated volume, ensure that you specify the host with the arbiter brick as **Host3** on this screen.

Gluster Deployment [X]

Hosts Additional Hosts Volumes Bricks Review

1 — 2 — 3 — 4 — 5

Host1: host1.example.com

Host2: host2.example.com

Host3: host3.example.com

Cancel < Back Next >

Click Next.

4. Specify additional hosts

For multi-node deployments, add the front-end Additional Hosts or IP addresses of the other two hyperconverged hosts to have them automatically added to Red Hat Virtualization Manager when deployment is complete.

Gluster Deployment [X]

Hosts Additional Hosts Volumes Bricks Review

1 — 2 — 3 — 4 — 5

Use same hostnames as in previous step

Host2: FQDN or IP address

Host3: FQDN or IP address

i Provide the address used to add the additional hosts to be managed by Hosted Engine preferably FQDN or IP address. And both FQDN needs to be added in known_hosts file.

Cancel < Back Next >



IMPORTANT

[Bug 1688269](#) means that hosts with IPv6 addresses are not automatically added to Red Hat Virtualization Manager when deployment is complete. You can add them after deployment by following [Add additional hyperconverged hosts to the hosted engine](#).

5. Specify volumes

Specify the volumes to create.

Gluster Deployment
✕

Hosts Additional Hosts Volumes Bricks Review

① — ② — ③ — ④ — ⑤

Name	Volume Type	Arbiter	Brick Dirs
engine	Replicate	<input type="checkbox"/>	/gluster_bricks/engine/engine ✕
data	Replicate	<input checked="" type="checkbox"/>	/gluster_bricks/data/data ✕
vmstore	Replicate	<input checked="" type="checkbox"/>	/gluster_bricks/vmstore/vmsto ✕

[+ Add Volume](#)

ⓘ First volume in the list will be used for hosted-engine deployment

Cancel < Back Next >

Name

Specify the name of the volume to be created.

Volume Type

Specify a **Replicate** volume type. Only replicated volumes are supported for this release.

Arbiter

Specify whether to create the volume with an arbiter brick. If this box is checked, the third disk stores only metadata.

Brick Dirs

The directory that contains this volume's bricks.

The default values are correct for most installations.

6. Specify bricks

Enter details of the bricks to be created. Use the **Select host** drop-down menu to change the host being configured.

Gluster Deployment ✕

Hosts
Additional Hosts
Volumes
Bricks
Review

Raid Information ⓘ

Raid Type RAID 6 ▾

Stripe Size(KB) 256 ▾

Data Disk Count 12 ▾

Brick Configuration

Select Host host1.example.com ▾

LV Name	Device Name	LV Size(GB) ⓘ	Thinp	Mount Point	Enable Dedupe & Expanded Disk Compression	Size(GB) ⓘ
engine	/dev/sdb	100 ▾	<input type="checkbox"/>	/gluster_bricks/engine	<input checked="" type="checkbox"/>	1000 ▾
data	/dev/sdb	500 ▾	<input checked="" type="checkbox"/>	/gluster_bricks/data	<input checked="" type="checkbox"/>	5000 ▾
vmstore	/dev/sdb	500 ▾	<input checked="" type="checkbox"/>	/gluster_bricks/vmstore	<input checked="" type="checkbox"/>	5000 ▾

Configure LV Cache

SSD /dev/sdc

Thinpool device sdb ▾

LV Size(GB) 220 ▾

Cache Mode ⓘ writethrough ▾

ⓘ Dedupe/compression is enabled at the storage device, and will be applicable for all bricks that use the device.

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

Cancel
< Back
Next >

RAID

Specify the RAID configuration to use. This should match the RAID configuration of your host. Supported values are **raid5**, **raid6**, and **jbod**. Setting this option ensures that your storage is correctly tuned for your RAID configuration.

Stripe Size

Specify the RAID stripe size in KB. Do not enter units, only the number. This can be ignored for **jbod** configurations.

Disk Count

Specify the number of data disks in a RAID volume. This can be ignored for **jbod** configurations.

LV Name

The name of the logical volume to be created. This is pre-filled with the name that you specified on the previous page of the wizard.

Device

Specify the raw device you want to use. Red Hat recommends an unpartitioned device.

Size

Specify the size of the logical volume to create in GB. Do not enter units, only the number. This number should be the same for all bricks in a replicated set. Arbiter bricks can be smaller than other bricks in their replication set.

Mount Point

The mount point for the logical volume. This is pre-filled with the brick directory that you specified on the previous page of the wizard.

Thinp

This option is enabled and volumes are thinly provisioned by default, except for the **engine** volume, which must be thickly provisioned.

Enable Dedupe & Compression

Specify whether to provision the volume using VDO for compression and deduplication at deployment time.

Logical Size (GB)

Specify the logical size of the VDO volume. This can be up to 10 times the size of the physical volume, with an absolute maximum logical size of 4 PB.

Configure LV Cache

Optionally, check this checkbox to configure a small, fast SSD device as a logical volume cache for a larger, slower logical volume. Add the device path to the **SSD** field, the size to the **LV Size (GB)** field, and set the **Cache Mode** used by the device.



WARNING

To avoid data loss when using write-back mode, Red Hat recommends using two separate SSD/NVMe devices. Configuring the two devices in a RAID-1 configuration (via software or hardware), significantly reduces the potential of data loss from lost writes.

For further information about lvmcache configuration, see [Red Hat Enterprise Linux 7 LVM Administration](#).

a. **(Optional) If your system has multipath devices, additional configuration is required.**

i. **To use multipath devices**

If you want to use multipath devices in your RHHI for Virtualization deployment, use multipath WWIDs to specify the device. For example, use **/dev/mapper/3600508b1001caab032303683327a6a2e** instead of **/dev/sdb**.

ii. **To disable multipath device use**

If multipath devices exist in your environment, but you do not want to use them for your RHHI for Virtualization deployment, blacklist the devices.

- A. Create a custom multipath configuration file.

```
# touch /etc/multipath/conf.d/99-custom-multipath.conf
```

- B. Add the following content to the file, replacing **<device>** with the name of the device to blacklist:

```
blacklist {
    devnodes "<device>"
}
```

For example, to blacklist the **/dev/sdb** device, add the following:

```
blacklist {
    devnodes "sdb"
}
```

- C. Restart multipathd.

```
# systemctl restart multipathd
```

- D. Verify that your disks no longer have multipath names by using the **lsblk** command. If multipath names are still present, reboot hosts.

7. Review and edit configuration

The screenshot shows the 'Gluster Deployment' web console. At the top, there is a progress bar with five steps: Hosts (1), Additional Hosts (2), Volumes (3), Bricks (4), and Review (5). The 'Review' step is currently selected and highlighted with a blue circle. Below the progress bar, there is a text area displaying the generated Ansible inventory file. The inventory file content is as follows:

```
Generated Ansible inventory : /etc/ansible/hc_wizard_inventory.yml
hc_nodes:
hosts:
host1.example.com:
  gluster_infra_volume_groups:
  - vgroup: gluster_vg_sdb
    pvname: /dev/mapper/vdo_sdb
  gluster_infra_mount_devices:
  - path: /gluster_bricks/engine
    lvname: gluster_lv_engine
    vgroup: gluster_vg_sdb
  - path: /gluster_bricks/data
    lvname: gluster_lv_data
    vgroup: gluster_vg_sdb
```

Below the inventory file, there is a checkbox labeled 'Enable Debug Logging' which is currently unchecked. At the bottom of the console, there are three buttons: 'Cancel', '< Back', and 'Deploy'.

- a. Click **Edit** to begin editing the generated deployment configuration file. Make any changes required and click **Save**.

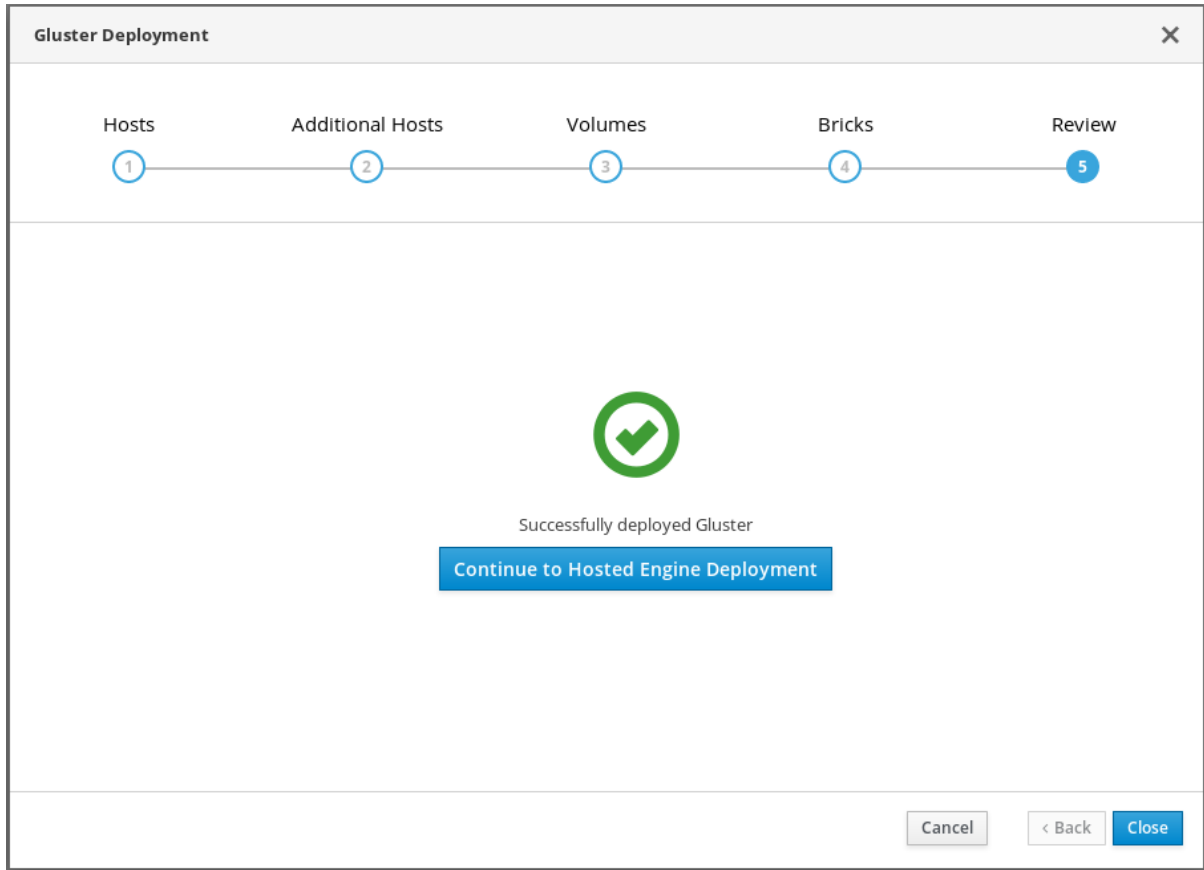
b. **Review the configuration file**

If all configuration details are correct, click **Deploy**.

8. **Wait for deployment to complete**

You can watch the progress of the deployment in the text field.

The window displays **Successfully deployed gluster** when complete.



Click **Continue to Hosted Engine Deployment** and continue the deployment process with the instructions in [Chapter 7, Deploy the Hosted Engine using the Web Console](#).



IMPORTANT

If deployment fails, click **Clean up** to remove any potentially incorrect changes to the system.

When cleanup is complete, click **Redeploy**. This returns you to the *Review and edit configuration* tab so that you can correct any issues in the generated configuration file before reattempting deployment.

CHAPTER 7. DEPLOY THE HOSTED ENGINE USING THE WEB CONSOLE

This section shows you how to deploy the Hosted Engine using the Web Console. Following this process results in Red Hat Virtualization Manager running as a virtual machine on the first physical machine in your deployment. It also configures a Default cluster comprised of the three physical machines, and enables Red Hat Gluster Storage functionality and the **virtual-host** *tuned* performance profile for each machine in the cluster.

Prerequisites

- [Configure Red Hat Gluster Storage for Hosted Engine using the Web Console](#)
- Gather the information you need for Hosted Engine deployment
Have the following information ready before you start the deployment process.
 - IP address for a pingable gateway to the hyperconverged host
 - IP address of the front-end management network
 - Fully-qualified domain name (FQDN) for the Hosted Engine virtual machine
 - MAC address that resolves to the static FQDN and IP address of the Hosted Engine

Procedure

1. Open the Hosted Engine Deployment wizard

If you continued directly from the end of [Configure Red Hat Gluster Storage for Hosted Engine using the Web Console](#), the wizard is already open.

Otherwise:

- a. Click **Virtualization** → **Hosted Engine**.
- b. Click **Start** underneath **Hyperconverged**.
- c. Click **Use existing configuration**.



IMPORTANT

If the previous deployment attempt failed, click **Clean up** instead of **Use existing configuration** to discard the previous attempt and start from scratch.

2. Specify virtual machine details

Hosted Engine Deployment
✕

VM Engine Prepare VM Storage Finish

1 2 3 4 5

VM Settings

Engine VM FQDN

MAC Address

Network Configuration

Bridge Interface

Root Password

Root SSH Access

Number of Virtual CPUs

Memory Size (MiB) 62,047MB available

> Advanced

- a. Enter the following details and click **Validate** for FQDN fields.

Engine VM FQDN

The fully qualified domain name to be used for the Hosted Engine virtual machine, for example, **engine.example.com**.

MAC Address

The MAC address associated with the **Engine VM FQDN**.



IMPORTANT

The pre-populated MAC address must be replaced.

Root password

The root password to be used for the Hosted Engine virtual machine.

- b. Click **Next**. Your FQDNs are validated before the next screen appears.

3. Specify virtualization management details

- a. Enter the password to be used by the **admin** account in the Administration Portal. You can also specify notification behavior here.

Hosted Engine Deployment ✕

VM Engine Prepare VM Storage Finish

① ——— ② ——— ③ ——— ④ ——— ⑤

Engine Credentials

Admin Portal Password

Notification Settings

Server Name

Server Port Number

Sender E-Mail Address

Recipient E-Mail Addresses

b. Click **Next**.

4. Review virtual machine configuration

a. Ensure that the details listed on this tab are correct. Click **Back** to correct any incorrect information.

Hosted Engine Deployment ✕

VM Engine Prepare VM Storage Finish

① ② ③ ④ ⑤

Please review the configuration. Once you click the 'Prepare VM' button, a local virtual machine will be started and used to prepare the management services and their data. This operation may take some time depending on your hardware.

▼ VM

- Engine FQDN: engine.example.com
- MAC Address: 00:xx:xx:xx:xx:xx
- Network Configuration: Static
- VM IP Address: 192.168.0.104
- Gateway Address: 192.168.0.104
- DNS Servers: 192.168.0.254
- Root User SSH Access: yes
- Number of Virtual CPUs: 4
- Memory Size (MiB): 16348
- Root User SSH Public Key: (None)
- Add Lines to /etc/hosts: yes
- Bridge Name: ovirtmgmt

▼ Engine


- SMTP Server Name: localhost
- SMTP Server Port Number: 25
- Sender E-Mail Address: root@localhost

- b. Click **Prepare VM**.
- c. Wait for virtual machine preparation to complete.

Hosted Engine Deployment
✕

VM Engine Prepare VM Storage Finish

① ————— ② ————— ③ ————— ④ ————— ⑤



Execution completed successfully. Please proceed to the next step.

Cancel < Back Next >

If preparation does not occur successfully, see [Viewing Hosted Engine deployment errors](#).

d. Click **Next**.

5. Specify storage for the Hosted Engine virtual machine

- a. Specify the primary host and the location of the **engine** volume.
- b. Ensure that the **Mount Options** field is populated correctly.
 - Ensure that **backup-volfile-servers=<host2-ip-address>:<host3-ip-address>**, with the appropriate IP addresses inserted in place of **<host2-ip-address>** and **<host3-ip-address>**.
 - If your environment uses IPv6 addresses, ensure that you add the **xlator-option=transport.address-family=inet6** option after the **backup-volfile-servers** value, like so:

```
backup-volfile-servers=<host2-ip-address>:<host3-ip-address>,xlator-
option=transport.address-family=inet6
```

Hosted Engine Deployment
✕

VM
Engine
Prepare VM
Storage
Finish

1

2

3

4

5

Please configure the storage domain that will be used to host the disk for the management VM. Please note that the management VM needs to be responsive and reliable enough to be able to manage all resources of your deployment, so highly available storage is preferred.

Storage Settings

Storage Type Gluster ▾

Storage Connection 192.168.0.101:/engine

Mount Options backup-volfile-servers=192.168.0.101;192.

> Advanced

Cancel
< Back
Next >

c. Click **Next**.

6. Finalize Hosted Engine deployment

a. Review your deployment details and verify that they are correct.



NOTE

The responses you provided during configuration are saved to an answer file to help you reinstall the hosted engine if necessary. The answer file is created at **/etc/ovirt-hosted-engine/answers.conf** by default. This file should not be modified manually without assistance from Red Hat Support.

Hosted Engine Deployment ✕

VM Engine Prepare VM Storage Finish

① ——— ② ——— ③ ——— ④ ——— ⑤

Please review the configuration. Once you click the 'Finish Deployment' button, the management VM will be transferred to the configured storage and the configuration of your hosted engine cluster will be finalized. You will be able to use your hosted engine once this step finishes.

▼ Storage

Storage Type: glusterfs

Storage Domain Connection: node1.example.com:/engine

Mount Options: backup-volfile-servers=node2.example.com;node3.example.com

Disk Size (GiB): 58

b. Click **Finish Deployment**.

7. **Wait for deployment to complete**


This takes up to 30 minutes.

The window displays the following when complete.

Hosted Engine Deployment ✕

VM Engine Prepare VM Storage Finish

① ——— ② ——— ③ ——— ④ ——— ⑤



Hosted engine deployment complete!

Close



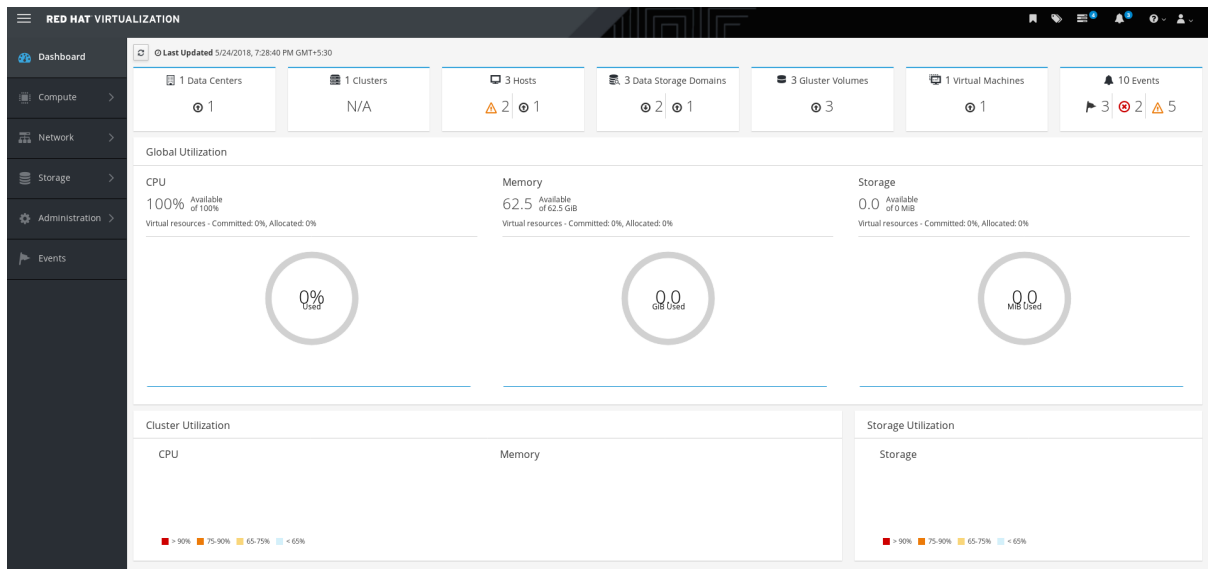
IMPORTANT

If deployment does not complete successfully, see [Viewing Hosted Engine deployment errors](#).

Click **Close**.

8. Verify hosted engine deployment

Browse to the Administration Portal (for example, <http://engine.example.com/ovirt-engine>) and verify that you can log in using the administrative credentials you configured earlier. Click **Dashboard** and look for your hosts, storage domains, and virtual machines.



Next steps

- [Log in to the Administration Portal to complete configuration](#) .

CHAPTER 8. CONFIGURE RED HAT GLUSTER STORAGE AS A RED HAT VIRTUALIZATION STORAGE DOMAIN

8.1. CREATE THE LOGICAL NETWORK FOR GLUSTER TRAFFIC

1. Log in to the Administration Portal

Browse to the Administration Portal (for example, <http://engine.example.com/ovirt-engine>) and log in using the administrative credentials you configured in [Chapter 7, Deploy the Hosted Engine using the Web Console](#).

2. Create a logical network for gluster traffic

- a. Click **Network** → **Networks** and then click **New**. The **New Logical Network** wizard appears.
- b. On the **General** tab of the wizard, provide a **Name** for the new logical network, and uncheck the **VM Network** checkbox.
- c. On the **Cluster** tab of the wizard, uncheck the **Required** checkbox.
- d. Click **OK** to create the new logical network.

3. Enable the new logical network for gluster

- a. Click the **Network** → **Networks** and select the new logical network.
- b. Click the **Clusters** subtab and then click **Manage Network**. The **Manage Network** dialogue appears.
- c. In the **Manage Network** dialogue, check the **Migration Network** and **Gluster Network** checkboxes.
- d. Click **OK** to save.

4. Attach the gluster network to the host

- a. Click **Compute** → **Hosts** and select the host.
- b. Click the **Network Interfaces** subtab and then click **Setup Host Networks**. The **Setup Host Networks** window opens.
- 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.

5. 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**.

8.2. CONFIGURE ADDITIONAL HYPERCONVERGED HOSTS

If your environment uses IPv6 addresses, or if you did not specify additional hyperconverged hosts as part of [Configure Red Hat Gluster Storage for Hosted Engine using the Web Console](#), follow these steps in the Administration Portal for each of the other hyperconverged hosts.

- a. Click **Compute** → **Hosts** and then click **New** to open the *New Host* window.
- b. Provide the **Name**, **Hostname**, and **Password** for the host that you want to manage.
- c. Under **Advanced Parameters**, uncheck the **Automatically configure host firewall** checkbox, as firewall rules are already configured by the deployment process.
- d. In the **Hosted Engine** tab of the *New Host* dialog, set the value of **Choose hosted engine deployment action** to **Deploy**. This ensures that the hosted engine can run on the new host.
- e. Click **OK**.
 1. **Attach the gluster network to all remaining hosts**
- f. Click the name of the newly added host to go to the host page.
- g. Click the **Network Interfaces** subtab and then click **Setup Host Networks**.
- h. Drag and drop the newly created network to the correct interface.
- i. Ensure that the **Verify connectivity** checkbox is checked.
- j. Ensure that the **Save network configuration** checkbox is checked.
- k. Click **OK** to save.
 1. In the **General** subtab for this host, verify that the value of **Hosted Engine HA** is *Active*, with a positive integer as a score.



IMPORTANT

If **Score** is listed as **N/A**, you may have forgotten to select the **deploy** action for **Choose hosted engine deployment action**. Follow the steps in [Reinstalling a hyperconverged host](#) in *Maintaining Red Hat Hyperconverged Infrastructure for Virtualization* to reinstall the host with the **deploy** action.

2. 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**.

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


PART III. VERIFY

CHAPTER 9. VERIFY YOUR DEPLOYMENT

After deployment is complete, verify that your deployment has completed successfully.


1. Browse to the Administration Portal, for example, <http://engine.example.com/ovirt-engine>.

Administration Console Login




2. Log in using the administrative credentials added during hosted engine deployment. When login is successful, the Dashboard appears.

Administration Console Dashboard



3. Verify that your cluster is available.

Administration Console Dashboard - Clusters

 1 Clusters

 1

4. Verify that at least one host is available.
If you provided additional host details during Hosted Engine deployment, 3 hosts are visible here, as shown.

Administration Console Dashboard - Hosts

 3 Hosts

 3

- a. Click **Compute** → **Hosts**.
- b. Verify that all hosts are listed with a **Status** of **Up**.

Administration Console - Hosts

Compute » [Hosts](#)

Host: ✕ ☆ ▾ 🔍

New Edit Remove Management ▾ Installation ▾ Host Console ⋮

↻ ▾ 1 - 3 < >

		Name	Comment	Hostname/IP	Cluster	Data Center	Status	Vi
▲	🔒	rhsdev-grafton2.lab.eng.b		rhsdev-grafton2.lab.en...	Default	Default	Up	
▲	🔒	rhsdev-grafton3.lab.eng.b		rhsdev-grafton3.lab.en...	Default	Default	Up	
▲	🔒	rhsdev-grafton4.lab.eng.b		rhsdev-grafton4.lab.en...	Default	Default	Up	

5. Verify that all storage domains are available.
 - a. Click **Storage** → **Domains**.
 - b. Verify that the **Active** icon is shown in the first column.

Administration Console - Storage Domains

Storage » [Storage Domains](#)

Storage: ✕ ☆ ▾ 🔍

New Domain Import Domain Manage Domain Remove ⋮

↻ ▾ 1 - 5 < >

	Domain Name	Comment	Domain Type	Storage Type	Format	Cross Data Center Status	Total Space	Free Space
▲	data		Data	GlusterFS	V4	Active	4998 GiB	4563 GiB
▲	hosted_storage		Data (Master)	GlusterFS	V4	Active	99 GiB	88 GiB
▲	vmstore		Data	GlusterFS	V4	Active	9998 GiB	9284 GiB

PART IV. NEXT STEPS

CHAPTER 10. POST-DEPLOYMENT CONFIGURATION SUGGESTIONS

Depending on your requirements, you may want to perform some additional configuration on your newly deployed Red Hat Hyperconverged Infrastructure for Virtualization. This section contains suggested next steps for additional configuration.

Details on these processes are available in [Maintaining Red Hat Hyperconverged Infrastructure for Virtualization](#).

10.1. CONFIGURE FENCING FOR HIGH AVAILABILITY

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

See [Configure High Availability using fencing policies](#) for further information.

10.2. CONFIGURE BACKUP AND RECOVERY OPTIONS

Red Hat recommends configuring at least basic disaster recovery capabilities on all production deployments.

See [Configuring backup and recovery options](#) in [Maintaining Red Hat Hyperconverged Infrastructure for Virtualization](#) for more information.

PART V. TROUBLESHOOT

CHAPTER 11. LOG FILE LOCATIONS

During the deployment process, progress information is displayed in the web browser. This information is also stored on the local file system so that the information logged can be archived or reviewed at a later date, for example, if the web browser stops responding or is closed before the information has been reviewed.

The log file for the the Web Console based deployment process (documented in [Chapter 6, Configure Red Hat Gluster Storage for Hosted Engine using the Web Console](#)) is stored in the **/var/log/ansible.log** file by default.

The log files for the Hosted Engine setup portion of the deployment process (documented in [Chapter 7, Deploy the Hosted Engine using the Web Console](#)) are stored in the **/var/log/ovirt-hosted-engine-setup** directory, with file names of the form **ovirt-hosted-engine-setup-<date>.log**.

CHAPTER 12. VIEWING DEPLOYMENT ERRORS

12.1. FAILED TO DEPLOY STORAGE

If an error occurs during [storage deployment](#), the deployment process halts and **Deployment failed** is displayed.

Deploying storage failed

Gluster Deployment [X]

Hosts (1) — Additional Hosts (2) — Volumes (3) — Bricks (4) — Review (5)

Deployment failed [CleanUp] [Redeploy]

```

fatal: [host2.example.com]: UNREACHABLE! => {"changed": false, "msg": "Failed to connect to the host via ssh: ssh: Could not
resolve hostname host2.example.com: Name or service not known", "unreachable": true}

NO MORE HOSTS LEFT *****

NO MORE HOSTS LEFT *****

PLAY RECAP *****
host1.example.com      :ok=0  changed=0  unreachable=1  failed=0  skipped=0  rescued=0  ignored=0
host2.example.com      :ok=0  changed=0  unreachable=1  failed=0  skipped=0  rescued=0  ignored=0
host3.example.com      :ok=0  changed=0  unreachable=1  failed=0  skipped=0  rescued=0  ignored=0

Please check /var/log/cockpit/ovirt-dashboard/gluster-deployment.log for more informations.

```

[Cancel] [< Back] [Close]

- Review the Web Console output for error information.
- Click **Clean up** to remove any potentially incorrect changes to the system.
- Click **Redeploy** and correct any entered values that may have caused errors. If you need help resolving errors, contact Red Hat Support with details.
- Return to [storage deployment](#) to try again.

12.2. FAILED TO PREPARE VIRTUAL MACHINE

If an error occurs while preparing the virtual machine in [Hosted Engine deployment](#), deployment pauses, and you see a screen similar to the following:

Preparing virtual machine failed

Hosted Engine Deployment
✕

VM Engine Prepare VM Storage Finish

① ————— ② ————— ③ ————— ④ ————— ⑤

✕
Deployment failed

```

[ INFO ] changed: [localhost]
[ INFO ] TASK [Check address resolution]
[ INFO ] skipping: [localhost]
[ INFO ] TASK [Parse host address resolution]
[ INFO ] ok: [localhost]
[ INFO ] TASK [Ensure host address resolves locally]
[ INFO ] skipping: [localhost]
[ INFO ] TASK [Get target address from selected interface]
[ INFO ] ok: [localhost]
[ INFO ] TASK [Check the resolved address resolves on the selected interface]
[ INFO ] skipping: [localhost]
[ INFO ] TASK [Check for alias]
[ INFO ] changed: [localhost]
[ INFO ] TASK [Ensure the resolved address resolves only on the selected interface]
[ INFO ] skipping: [localhost]
[ INFO ] TASK [Avoid localhost]
[ INFO ] skipping: [localhost]
[ INFO ] TASK [Get engine FQDN resolution]
[ INFO ] TASK [Check engine FQDN resolution]
[ ERROR ] fatal: [localhost]: FAILED! => {"changed": false, "msg": "Unable to resolve address\n"}

```

Cancel
< Back
Prepare VM

- Review the Web Console output for error information.
- Click **Clean up** to remove any potentially incorrect changes to the system.
- Click **Redeploy** and correct any entered values that may have caused errors. If you need help resolving errors, contact Red Hat Support with details.
- Ensure that the **rhvm-appliance** package is available on all hosts.

```
# yum install rhvm-appliance
```


- Return to [Hosted Engine deployment](#) to try again.
If you closed the deployment wizard while you resolved errors, you can select **Use existing configuration** when you retry the deployment process.

12.3. FAILED TO DEPLOY HOSTED ENGINE

If an error occurs during hosted engine deployment, deployment pauses, and **Deployment failed** is displayed.

Hosted engine deployment failed

Hosted Engine Deployment ✕



✕ Deployment failed

```
[ INFO ] TASK [Obtain SSO token using username/password credentials]
[ INFO ] ok: [localhost]
[ INFO ] TASK [Fetch host facts]
[ INFO ] ok: [localhost]
[ INFO ] TASK [Fetch cluster ID]
[ INFO ] ok: [localhost]
[ INFO ] TASK [Fetch cluster facts]
[ INFO ] ok: [localhost]
[ INFO ] TASK [Fetch Datacenter facts]
[ INFO ] ok: [localhost]
[ INFO ] TASK [Fetch Datacenter ID]
[ INFO ] ok: [localhost]
[ INFO ] TASK [Fetch Datacenter name]
[ INFO ] ok: [localhost]
[ INFO ] TASK [Add NFS storage domain]
[ INFO ] skipping: [localhost]
[ INFO ] TASK [Add glusterfs storage domain]
[ ERROR ] Error: Fault reason is "Operation Failed". Fault detail is "[Failed to fetch Gluster Volume List]". HTTP response code is 400.
[ ERROR ] fatal: [localhost]: FAILED! => {"changed": false, "msg": "Fault reason is \"Operation Failed\". Fault detail is \"[Failed to fetch Gluster Volume List]\". HTTP response code is 400."}
```

Cancel

< Back

Redeploy

- Review the Web Console output for error information.

- Remove the contents of the **engine** volume.

- a. Mount the **engine** volume.

```
# mount -t glusterfs <server1>:/engine /mnt/test
```

- b. Remove the contents of the volume.

```
# rm -rf /mnt/test/*
```

- c. Unmount the **engine** volume.

```
# umount /mnt/test
```

- Click **Clean up** to remove any potentially incorrect changes to the system.
- Click **Redeploy** and correct any entered values that may have caused errors. If you need help resolving errors, contact Red Hat Support with details.
- Return to [Hosted Engine deployment](#) to try again.
If you closed the deployment wizard while you resolved errors, you can select **Use existing configuration** when you retry the deployment process.

PART VI. REFERENCE MATERIAL

APPENDIX A. GLOSSARY OF TERMS

A.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.

A.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.

A.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.

