



Red Hat Hyperconverged Infrastructure for Virtualization 1.8

Managing Red Hat Gluster Storage using the Web Console

Perform common Red Hat Gluster Storage management tasks in the Web Console

Red Hat Hyperconverged Infrastructure for Virtualization 1.8 Managing Red Hat Gluster Storage using the Web Console

Perform common Red Hat Gluster Storage management tasks in the Web Console

Laura Bailey

lbailey@redhat.com

Legal Notice

Copyright © 2022 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

After Red Hat Hyperconverged Infrastructure for Virtualization has been deployed, you can perform many operational and management tasks for Red Hat Gluster Storage using the Web Console. Read this book to understand how to manage storage using the Web Console. This document explains how to perform maintenance tasks specific to Red Hat Hyperconverged Infrastructure for Virtualization.

Table of Contents

MAKING OPEN SOURCE MORE INCLUSIVE	4
CHAPTER 1. UNDERSTANDING THE WEB CONSOLE	5
1.1. UNDERSTANDING THE STORAGE DASHBOARD	5
1.2. UNDERSTANDING THE VOLUME GROUP OVERVIEW PAGE	9
1.3. UNDERSTANDING THE VDO OVERVIEW PAGE	10
1.4. UNDERSTANDING THE DRIVE OVERVIEW PAGE	12
CHAPTER 2. MANAGING PHYSICAL DEVICE DRIVES USING THE WEB CONSOLE	13
2.1. CREATING A PARTITION TABLE USING THE WEB CONSOLE	13
2.2. FORMATTING A DISK PARTITION USING THE WEB CONSOLE	13
CHAPTER 3. MANAGING RAID DEVICES USING THE WEB CONSOLE	15
3.1. CONFIGURING A NEW RAID DEVICE USING THE WEB CONSOLE	15
CHAPTER 4. MANAGING VIRTUAL DATA OPTIMIZERS USING THE WEB CONSOLE	16
4.1. CREATING A VIRTUAL DATA OPTIMIZER USING THE WEB CONSOLE	16
4.2. GROWING THE LOGICAL SIZE OF A VDO DEVICE USING THE WEB CONSOLE	17
4.3. DISABLING COMPRESSION ON A VDO DEVICE USING THE WEB CONSOLE	17
4.4. DISABLING DEDUPLICATION ON A VDO DEVICE USING THE WEB CONSOLE	18
4.5. STOPPING A VDO DEVICE USING THE WEB CONSOLE	18
4.6. DELETING A VDO DEVICE USING THE WEB CONSOLE	18
CHAPTER 5. MANAGING VOLUME GROUPS USING THE WEB CONSOLE	20
5.1. CREATING A VOLUME GROUP USING THE WEB CONSOLE	20
5.2. ADDING PHYSICAL VOLUMES TO A VOLUME GROUP USING THE WEB CONSOLE	20
5.3. RENAMING A VOLUME GROUP USING THE WEB CONSOLE	21
5.4. DELETING A VOLUME GROUP USING THE WEB CONSOLE	21
CHAPTER 6. MANAGING THIN POOLS USING THE WEB CONSOLE	23
6.1. CREATING A THIN POOL USING THE WEB CONSOLE	23
6.2. GROWING A THIN POOL USING THE WEB CONSOLE	23
6.3. DEACTIVATING A THIN POOL USING THE WEB CONSOLE	23
6.4. ACTIVATING A THIN POOL USING THE WEB CONSOLE	24
CHAPTER 7. MANAGING LOGICAL VOLUMES USING THE WEB CONSOLE	25
7.1. ACTIVATING A LOGICAL VOLUME USING THE WEB CONSOLE	25
7.2. CREATING A THINLY PROVISIONED LOGICAL VOLUME USING THE WEB CONSOLE	25
7.3. CREATING A THICKLY PROVISIONED LOGICAL VOLUME USING THE WEB CONSOLE	25
7.4. DEACTIVATING A LOGICAL VOLUME USING THE WEB CONSOLE	26
7.5. DELETING A LOGICAL VOLUME USING THE WEB CONSOLE	26
7.6. GROWING A LOGICAL VOLUME USING THE WEB CONSOLE	27
CHAPTER 8. MANAGING FILE SYSTEMS USING THE WEB CONSOLE	28
8.1. FORMATTING A LOGICAL VOLUME USING THE WEB CONSOLE	28
8.2. CONFIGURING A MOUNT POINT USING THE WEB CONSOLE	28
8.3. MOUNTING A FILE SYSTEM USING THE WEB CONSOLE	29
8.4. UNMOUNTING A FILE SYSTEM USING THE WEB CONSOLE	29
CHAPTER 9. MANAGING GLUSTER VOLUMES USING THE WEB CONSOLE	31
9.1. ACCESSING THE GLUSTER MANAGEMENT DASHBOARD	31
9.2. EXPANDING VOLUME FROM WEB CONSOLE	32
9.3. EXPANDING VOLUME FROM RED HAT VIRTUALIZATION MANAGER	35
9.4. EXPANDING THE HYPERCONVERGED CLUSTER BY ADDING A NEW VOLUME ON NEW NODES USING	

THE WEB CONSOLE	36
9.4.1. Configure additional hyperconverged hosts	39
9.5. CREATING AN ADDITIONAL GLUSTER VOLUME USING THE WEB CONSOLE	40

MAKING OPEN SOURCE MORE INCLUSIVE

Red Hat is committed to replacing problematic language in our code, documentation, and web properties. We are beginning with these four terms: master, slave, blacklist, and whitelist. Because of the enormity of this endeavor, these changes will be implemented gradually over several upcoming releases. For more details, see [our CTO Chris Wright's message](#).

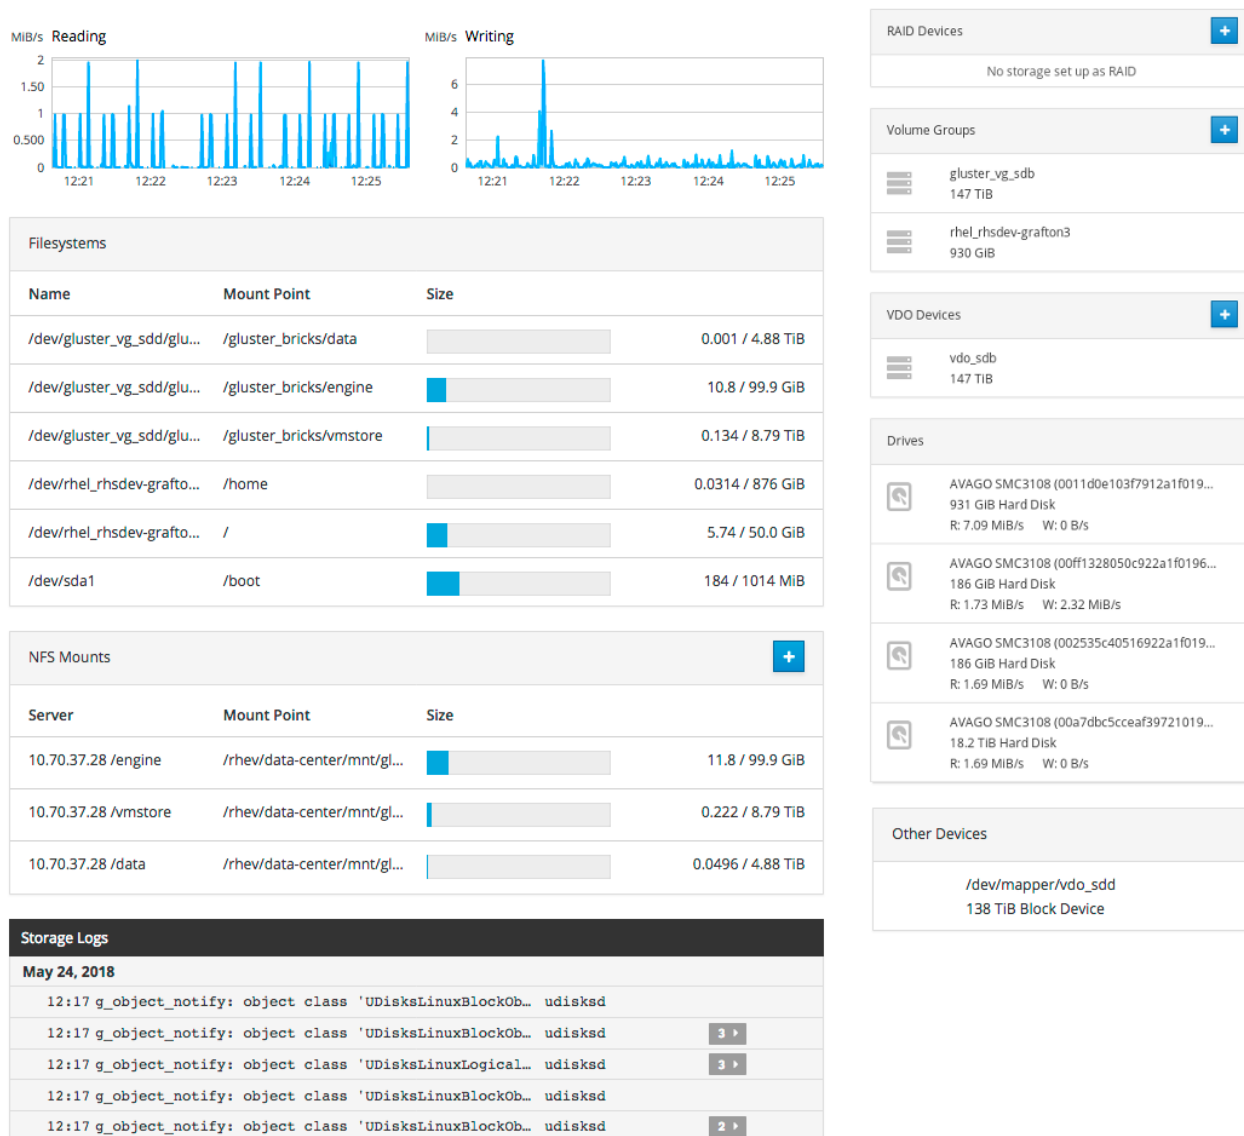
CHAPTER 1. UNDERSTANDING THE WEB CONSOLE

1.1. UNDERSTANDING THE STORAGE DASHBOARD

The Storage Dashboard in the Web Console provides an overview of I/O operations, volumes, and devices available to the server.

Navigate to the Storage Dashboard by logging into the Web Console and clicking the hostname, followed by **Storage**.

Figure 1.1. Storage Dashboard




The Storage Dashboard is divided into a number of sections.

Reading

A graph of the amount of data read in the last five minutes. The scale of the graph adjusts automatically according to server load.


Figure 1.2. Reading graph



Writing

A graph of the amount of data written in the last five minutes. The scale of the graph adjusts automatically according to server load.

Figure 1.3. Writing graph



Filesystems

A list of file systems hosted on this server showing their name, mount point, size, and used versus total storage capacity. Clicking on a file system opens the *Volume Group Overview* page. See [Section 1.2, "Understanding the Volume Group Overview page"](#) for more information.

Figure 1.4. Filesystems

Filesystems			
Name	Mount Point	Size	
/dev/gluster_vg_sdd/glu...	/gluster_bricks/data	<div style="width: 0.001%; background-color: #ccc;"></div>	0.001 / 4.88 TiB
/dev/gluster_vg_sdd/glu...	/gluster_bricks/engine	<div style="width: 10.8%; background-color: #0070c0;"></div>	10.8 / 99.9 GiB
/dev/gluster_vg_sdd/glu...	/gluster_bricks/vmstore	<div style="width: 0.134%; background-color: #ccc;"></div>	0.134 / 8.79 TiB
/dev/rhel_rhsdev-grafto...	/home	<div style="width: 0.0314%; background-color: #ccc;"></div>	0.0314 / 876 GiB
/dev/rhel_rhsdev-grafto...	/	<div style="width: 5.74%; background-color: #0070c0;"></div>	5.74 / 50.0 GiB
/dev/sda1	/boot	<div style="width: 184%; background-color: #0070c0;"></div>	184 / 1014 MiB

NFS Mounts

A list of exported file systems that have been mounted by client systems.

Figure 1.5. Mounts

NFS Mounts +		
Server	Mount Point	Size
10.70.37.28 /engine	/rhev/data-center/mnt/gl...	<div style="width: 10%; background-color: #0070C0; height: 10px;"></div> 11.8 / 99.9 GiB
10.70.37.28 /vmstore	/rhev/data-center/mnt/gl...	<div style="width: 1%; background-color: #0070C0; height: 10px;"></div> 0.222 / 8.79 TiB
10.70.37.28 /data	/rhev/data-center/mnt/gl...	<div style="width: 0%; background-color: #0070C0; height: 10px;"></div> 0.0496 / 4.88 TiB

Storage Logs

A list of recent events and log messages related to storage.

Figure 1.6. Logs

Storage Logs	
May 24, 2018	
12:17 g_object_notify: object class 'UDisksLinuxBlockOb... udisksd	
12:17 g_object_notify: object class 'UDisksLinuxBlockOb... udisksd	3 ▶
12:17 g_object_notify: object class 'UDisksLinuxLogical... udisksd	3 ▶
12:17 g_object_notify: object class 'UDisksLinuxBlockOb... udisksd	
12:17 g_object_notify: object class 'UDisksLinuxBlockOb... udisksd	2 ▶

RAID Devices

A list of RAID devices configured on the server. Clicking a RAID device opens the *RAID Device Summary* page. Clicking + opens the *Create RAID Device* window. See [Section 3.1, "Configuring a new RAID Device using the Web Console"](#) for more information.

Figure 1.7. RAID Devices





RAID Devices +
No storage set up as RAID

Volume Groups

A list of volume groups configured on the server. Clicking a volume group opens the *Volume Group Overview* page. Clicking + opens the *Create Volume Group* window. See the following sections for more information:

- [Section 5.1, "Creating a volume group using the Web Console"](#)
- [Section 1.2, "Understanding the Volume Group Overview page"](#)

Figure 1.8. Volume Groups


Volume Groups		
	gluster_vg_sdb 186 GiB	
	gluster_vg_sdd 138 TiB	
	rhel_rhsdev-grafton2 930 GiB	

VDO Devices

A list of Virtual Data Optimizer devices configured on the server. Clicking a VDO device opens the *VDO Device Summary* page. Clicking + opens the *Create VDO Device* window. See the following sections for more information:

- [Section 4.1, “Creating a Virtual Data Optimizer using the Web Console”](#)
- [Section 1.3, “Understanding the VDO Overview page”](#)





Figure 1.9. VDO Devices

VDO Devices		
	vdo_sdb 147 TiB	

Drives

A list of drives in use by the server. Clicking a drive opens the *Drive Summary* page. See [Section 1.4, “Understanding the Drive Overview page”](#) for more information.

Figure 1.10. Drives

Drives	
	AVAGO SMC3108 (0025e3d21626f... 931 GiB Hard Disk R: 0 B/s W: 0 B/s
	AVAGO SMC3108 (00c6008c1843fe... 186 GiB Hard Disk R: 0 B/s W: 0 B/s
	AVAGO SMC3108 (0007c8301a5ffe... 186 GiB Hard Disk R: 0 B/s W: 0 B/s
	AVAGO SMC3108 (00046eef1c8dfe... 18.2 TiB Hard Disk R: 0 B/s W: 332 KiB/s

Other Devices

A list of other devices attached to the server, that are not yet used.

Figure 1.11. Other Devices

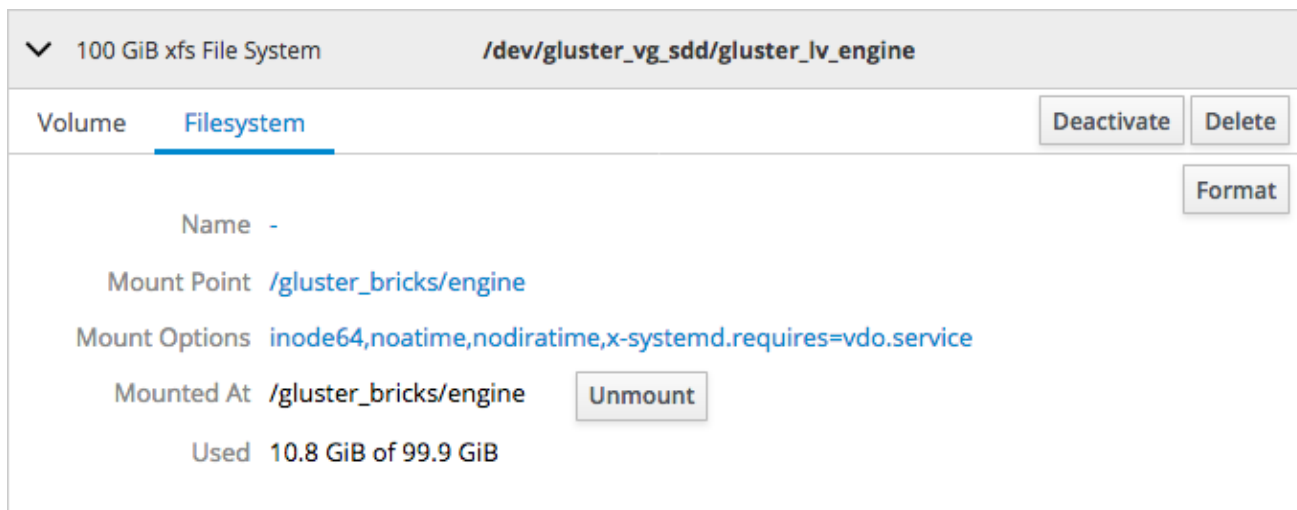
Other Devices
/dev/mapper/vdo_sdd 138 TiB Block Device

1.2. UNDERSTANDING THE VOLUME GROUP OVERVIEW PAGE

The Volume Group Overview page in the Web Console provides an overview of a volume group, the physical volume it resides on, and the logical volumes that are part of that volume group.

Navigate to the Volume Group Summary page for a file system or volume group by logging into the Web Console and clicking the hostname, followed by **Storage**, and clicking any file system or volume group.

Figure 1.12. Volume Group Overview page



The Volume Group Overview page is divided into several sections.

Volume Group summary

A summary of volume group details, showing name, UUID and capacity. You can also rename or delete the volume group from here.

Physical Volumes

A list of physical volumes underlying this volume group. Clicking + opens the *Add Disks* window, which lets you add more disks to the volume group.

Logical Volumes

A list of logical volumes in this volume group, with volume management operations. Clicking a volume shows more details and operations for that volume. See [Chapter 7, Managing logical volumes using the Web Console](#) for more information.

1.3. UNDERSTANDING THE VDO OVERVIEW PAGE

The VDO Overview page provides an overview of Virtual Disk Optimizer devices and their contents.

Navigate to the VDO Overview page for a VDO device by logging into the Web Console and clicking the hostname, followed by **Storage**, and clicking any device listed under **VDO Devices**.

Figure 1.13. VDO Overview page

Storage > vdo_sdb

VDO Device vdo_sdb

Stop
Delete

Device File `/dev/mapper/vdo_sdb`

Backing Device [AVAGO SMC3108 \(00c6008c1843fe6622013bcd1a800403\)](#)

Physical 5.29 GiB data + 25.8 GiB overhead used of 186 GiB (16%)

Logical 18.8 GiB used of 147 TiB (71% saved) Grow

Index Memory 256 MiB

Compression ON

Deduplication ON

Content

▼ 147 TiB Physical volume of gluster_vg_sdb
`/dev/mapper/vdo_sdb`

Physical Volume

Volume Group [gluster_vg_sdb](#)
Format

Free 133 TiB

The VDO Overview page is divided into several sections.

VDO Device summary

A summary of the details of this VDO device, showing the device file location, the drive backing the VDO device, and the physical and logical space available. Several VDO management operations are also available.

Content

A summary of the volumes residing on this VDO device.

The VDO Overview page also provides access to a number of VDO management operations.

See the following sections for more information about each operation.

- [Growing the logical size of a VDO device](#)
- [Disabling compression on a VDO device](#)
- [Disabling deduplication on a VDO device](#)
- [Stopping a VDO device](#)
- [Deleting a VDO device](#)

See [Understanding VDO](#) for more information about VDO devices.

1.4. UNDERSTANDING THE DRIVE OVERVIEW PAGE

The Drive Overview page provides an overview of any drives or disks attached to the server, and any partitions on those drives.

Navigate to the Drive Overview page for any drive by logging into the Web Console and clicking the hostname, followed by **Storage**, and clicking any drive listed under **Drives**.

Figure 1.14. Drive Overview page

[Storage](#) > AVAGO SMC3108 (0025e3d21626fe6622013bcd1a800403)

The screenshot shows the Drive Overview page for an AVAGO SMC3108 drive. The drive details section includes:

- Model: SMC3108
- Firmware Version: 4.29
- Serial Number: 0025e3d21626fe6622013bcd1a800403
- World Wide Name: 0x600304801acd3b012266fe2616d2e325
- Capacity: 931 GiB, 1000 GB, 999653638144 bytes
- Device File: /dev/sda

The Content section shows a table of partitions on the drive:

Content	Device File
1 GiB xfs File System	/dev/sda1
930 GiB Physical volume of rhel_rhsdev-grafton2	/dev/sda2

A "Create partition table" button is visible in the top right corner of the Content section.

The Drive Overview page is divided into several sections.

Drive

A summary of drive details, including the drive model, firmware version, serial number, capacity, and device location.

Content

A summary of partitions on this drive. You can modify the partition table on the drive, or format a partition here.

CHAPTER 2. MANAGING PHYSICAL DEVICE DRIVES USING THE WEB CONSOLE

2.1. CREATING A PARTITION TABLE USING THE WEB CONSOLE

Follow these steps to create a new partition table on a drive using the Web Console.

1. Log in to the Web Console.
2. Click the hostname → **Storage**.
3. Click any drive under **Drives**. The *Drive Overview* page opens.
4. Click **Create partition table**.

Figure 2.1. Drive Content

Content		Create partition table
1 GiB xfs File System	<code>/dev/sda1</code>	
930 GiB Physical volume of rhel_rhdev-grafton2	<code>/dev/sda2</code>	

The *Format device* window opens.

- a. Specify whether to **Erase** existing data completely by overwriting it with zeroes.
- b. Specify the **Partitioning** style to use.
- c. Click **Format**.

2.2. FORMATTING A DISK PARTITION USING THE WEB CONSOLE

Follow these steps to format a partition with a file system using the Web Console.

1. Log in to the Web Console.
2. Click the hostname → **Storage**.
3. Click any drive under **Drives**. The *Drive Overview* page opens.
4. Click the device under **Content**.
5. In the **Filesystem** subtab, click **Format**.

The Filesystem subtab



100 GiB xfs File System /dev/gluster_vg_sdd/gluster_lv_engine

Volume Filesystem Deactivate Delete

Name - Format

Mount Point /gluster_bricks/engine

Mount Options inode64,noatime,nodiratime,x-systemd.requires=vdo.service

Mounted At /gluster_bricks/engine Unmount

Used 10.8 GiB of 99.9 GiB

The *Format Device* window appears.


- Specify whether to **Erase** existing data completely by overwriting it with zeroes.
- Specify the file system **Type** to use.
- Specify a **Name** for the file system.
- Specify whether to use default or customized **Mounting** behavior.
If you selected **Custom**, specify a **Mount Point** and check any **Mount options** you want this file system to use.
- Click **Format**.

CHAPTER 3. MANAGING RAID DEVICES USING THE WEB CONSOLE

3.1. CONFIGURING A NEW RAID DEVICE USING THE WEB CONSOLE

1. Log in to the Web Console.
2. Click the hostname → **Storage**.
3. Click the + button beside **RAID Devices**.

Figure 3.1. RAID Devices



The *Create RAID Device* window opens.

Figure 3.2. Create Raid Device window

- a. Specify a **Name** for your RAID device.
- b. Specify the **RAID Level** to use.



NOTE

Only RAID5 and RAID6 are supported for RHHI for Virtualization.

- c. Leave the **Chunk Size** as the default value.
- d. Check the **Disks** to use in creating the RAID device.
- e. Click **Create**.

CHAPTER 4. MANAGING VIRTUAL DATA OPTIMIZERS USING THE WEB CONSOLE

4.1. CREATING A VIRTUAL DATA OPTIMIZER USING THE WEB CONSOLE

Follow these steps to create a new VDO device.

1. Log in to the Web Console.
2. Click the hostname → **Storage**.
3. Click the + button beside **VDO Devices**. The *Create VDO Device* window opens.

Figure 4.1. VDO Devices




Figure 4.2. Create VDO Device window

Create VDO Device

Name

Disk 186 GiB AVAGO SMC3108 (0007c8301a5ffe6622013bcd1a8004... /dev/sdc

Logical Size 1024 GiB ▾

Index Memory 256 MiB ▾

Options Compression
 Deduplication
 Use 512 Byte emulation

- a. Specify a **Name** for the new VDO device.
- b. Specify the **Disk** to use for the new VDO device.
- c. Specify the **Logical Size** for the new VDO device.

- d. Specify the **Index Memory** size for the new VDO device.
- e. Check any **Options** you want the new VDO device to use.
- f. Click **Create**.

4.2. GROWING THE LOGICAL SIZE OF A VDO DEVICE USING THE WEB CONSOLE

Follow these steps to increase the logical size (the amount of data it can store by using compression and deduplication) of a VDO device.

1. Log in to the Web Console.
2. Click the hostname → **Storage**. The *Storage Dashboard* page opens.
3. Under **VDO Devices**, click the VDO device. The *VDO Device Overview* page opens.
4. Click **Grow**.

Figure 4.3. VDO device summary



The *Grow logical size of device* window opens.

5. Drag the slider or enter a numeric value to set the new logical size of the VDO device, and click **Grow**.

4.3. DISABLING COMPRESSION ON A VDO DEVICE USING THE WEB CONSOLE

Follow these steps to prevent future writes to the VDO device being compressed. This does not affect the compression of existing data on the VDO device.

1. Log in to the Web Console.
2. Click the hostname → **Storage**. The *Storage Dashboard* page opens.
3. Under **VDO Devices**, click the VDO device. The *VDO Device Overview* page opens.
4. Locate the **Compression** toggle.

Figure 4.4. The Compression toggle

5. Click the **Compression** toggle under the device summary so that **OFF** is displayed.

4.4. DISABLING DEDUPLICATION ON A VDO DEVICE USING THE WEB CONSOLE

Follow these steps to prevent future writes to the VDO device being checked for duplication. This does not affect the deduplication of existing data on the VDO device.

1. Log in to the Web Console.
2. Click the hostname → **Storage**. The *Storage Dashboard* page opens.
3. Under **VDO Devices**, click the VDO device. The *VDO Device Overview* page opens.
4. Locate the **Deduplication** toggle.

Figure 4.5. The Deduplication toggle

5. Click the **Deduplication** toggle under the device summary so that **OFF** is displayed.

4.5. STOPPING A VDO DEVICE USING THE WEB CONSOLE


Follow these steps to stop a VDO device for maintenance operations or deletion.

Prerequisites

- Stop any volumes running on the VDO device.

Procedure

1. Log in to the Web Console.
2. Click the hostname → **Storage**. The *Storage Dashboard* page opens.
3. Under **VDO Devices**, click the VDO device. The *VDO Device Overview* page opens.
4. Click **Stop**.

Figure 4.6. The VDO device

4.6. DELETING A VDO DEVICE USING THE WEB CONSOLE

Follow these steps to delete a VDO device.

Prerequisites


Prerequisites

- Stop any volumes running on the VDO device.
- Stop the VDO device.

Procedure

1. Log in to the Web Console.
2. Click the hostname → **Storage**. The *Storage Dashboard* page opens.
3. Under **VDO Devices**, click the VDO device. The *VDO Device Overview* page opens.
4. Click **Delete**.

Figure 4.7. The VDO device

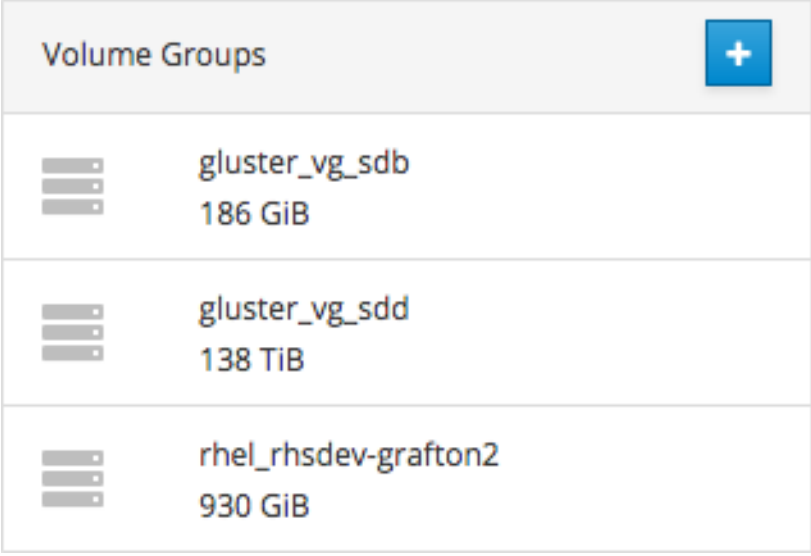





CHAPTER 5. MANAGING VOLUME GROUPS USING THE WEB CONSOLE

5.1. CREATING A VOLUME GROUP USING THE WEB CONSOLE

1. Log in to the Web Console.
2. Click the hostname → **Storage**.
3. Click the + button beside **Volume Groups**.

Figure 5.1. Volume Groups section



Volume Groups		+
	gluster_vg_sdb 186 GiB	
	gluster_vg_sdd 138 TiB	
	rhel_rhsdev-grafton2 930 GiB	

The *Create Volume Group* window opens.


- a. Specify a **Name** for the volume group.
- b. Check the **Disks** to use for the volume group.
- c. Click **Create**.

5.2. ADDING PHYSICAL VOLUMES TO A VOLUME GROUP USING THE WEB CONSOLE

Follow these instructions to expand a volume group using the Web Console.

1. Log in to the Web Console.
2. Click the hostname → **Storage**.
3. Click the volume group. The *Volume Group Overview* page opens.
4. Click the + button beside **Physical Volumes**.

Figure 5.2. Physical Volumes



The *Add Disks* window opens.

- a. Check the **Disks** you want to add to the volume group.
- b. Click **Add**.

The capacity of the volume group is extended automatically.

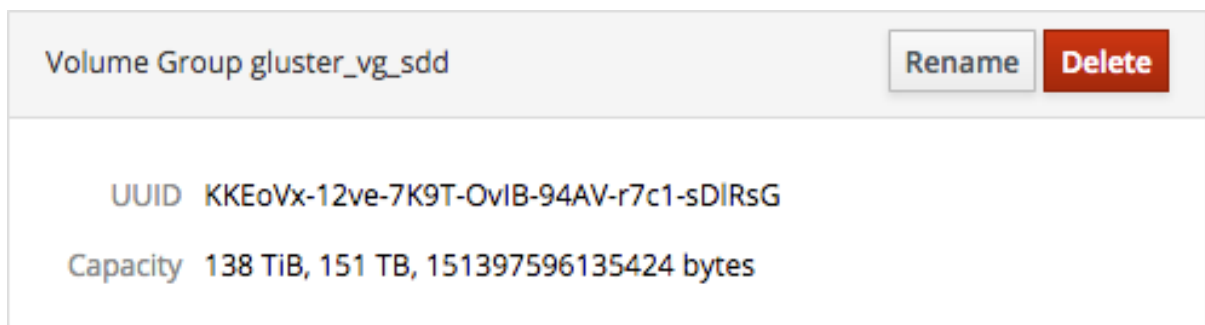
Now that you've added new storage to the volume group, you may want to [grow your thin pool](#) or [grow your logical volume](#).

5.3. RENAMING A VOLUME GROUP USING THE WEB CONSOLE

Follow these instructions to rename a volume group using the Web Console.

1. Log in to the Web Console.
2. Click the hostname → **Storage**.
3. Click the volume group. The *Volume Group Overview* page opens.
4. Click **Rename**.

Figure 5.3. Volume group summary



The *Rename Volume Group* window opens. .. Specify a new **Name**. .. Click **Rename**.

The volume group is renamed, and the change is propagated to all logical volumes in the volume group. This does not affect the naming of thin pools.


5.4. DELETING A VOLUME GROUP USING THE WEB CONSOLE

Follow these instructions to delete a volume group using the Web Console.

1. Log in to the Web Console.

2. Click the hostname → **Storage**.
3. Click the volume group. The *Volume Group Overview* page opens.
4. In the Volume Group summary, click **Delete**.

Figure 5.4. Volume group summary



The confirmation window opens.

5. Click **Delete** to confirm deletion.

CHAPTER 6. MANAGING THIN POOLS USING THE WEB CONSOLE

6.1. CREATING A THIN POOL USING THE WEB CONSOLE

Follow these instructions to create a logical thin pool using the Web Console.

1. Log in to the Web Console.
2. Click the hostname → **Storage**.
3. Click the volume group. The *Volume Group Overview* page opens.
4. Click + **Create new Logical Volume** The *Create Logical Volume* window opens.
 - a. Specify a **Name** for your thin pool.
 - b. Set **Purpose** to **Pool for thinly provisioned volumes**.
 - c. Specify a **Size** for your thin pool.
 - d. Click **Create**.

Your new thin pool appears in the list of logical volumes in this volume group.

6.2. GROWING A THIN POOL USING THE WEB CONSOLE

Follow these instructions to increase the size of a logical thin pool using the Web Console.

1. Log in to the Web Console.
2. Click the hostname → **Storage**.
3. Click the volume group. The *Volume Group Overview* page opens.
4. Click the thin pool.
5. On the **Pool** subtab, click **Grow**. The *Grow Logical Volume* window opens.
 - a. Specify the new **Size** of the thin pool.
 - b. Click **Grow**.

6.3. DEACTIVATING A THIN POOL USING THE WEB CONSOLE

Follow these instructions to deactivate a logical thin pool using the Web Console. This deactivates all thinly provisioned logical volumes in the pool.

1. Log in to the Web Console.
2. Click the hostname → **Storage**.
3. Click the volume group. The *Volume Group Overview* page opens.
4. Click the thin pool.

5. Click **Deactivate**.

The thin pool is deactivated.

6.4. ACTIVATING A THIN POOL USING THE WEB CONSOLE

Follow these instructions to activate a logical thin pool using the Web Console.

1. Log in to the Web Console.
2. Click the hostname → **Storage**.
3. Click the volume group. The *Volume Group Overview* page opens.
4. Click the thin pool.
5. Click **Activate**.

The thin pool is activated. This does not activate thin provisioned logical volumes in the pool.

CHAPTER 7. MANAGING LOGICAL VOLUMES USING THE WEB CONSOLE

7.1. ACTIVATING A LOGICAL VOLUME USING THE WEB CONSOLE

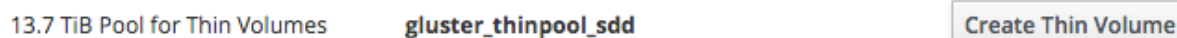
Follow these instructions to activate a logical volume using the Web Console.

1. Log in to the Web Console.
2. Click the hostname → **Storage**.
3. Click the volume group. The *Volume Group Overview* page opens.
4. Click the logical volume.
5. Click **Activate**.

7.2. CREATING A THINLY PROVISIONED LOGICAL VOLUME USING THE WEB CONSOLE

1. Log in to the Web Console.
2. Click the hostname → **Storage**.
3. Click the volume group. The *Volume Group Overview* page opens.
4. Click **Create Thin Volume** beside the thin pool that should host the volume.

Figure 7.1. A thin pool



13.7 TiB Pool for Thin Volumes **gluster_thinpool_sdd** Create Thin Volume

The *Create Thin Volume* window opens.

- a. Specify a **Name** for the new volume.
- b. Specify a **Size** for the new volume.
- c. Click **Create**.

The new volume appears in the list of logical volumes.

7.3. CREATING A THICKLY PROVISIONED LOGICAL VOLUME USING THE WEB CONSOLE

Follow these instructions to create a logical thin pool using the Web Console.

1. Log in to the Web Console.
2. Click the hostname → **Storage**.
3. Click the volume group. The *Volume Group Overview* page opens.

4. Click + **Create new Logical Volume** The *Create Logical Volume* window opens.

Figure 7.2. The Create Logical Volume window

- a. Specify a **Name** for your logical volume.
- b. Set **Purpose** to **Block device for file systems**.
- c. Specify a **Size** for your logical volume.
- d. Click **Create**.

Your new logical volume appears in the list of logical volumes in this volume group.

7.4. DEACTIVATING A LOGICAL VOLUME USING THE WEB CONSOLE

Follow these instructions to deactivate a logical volume using the Web Console.

1. Log in to the Web Console.
2. Click the hostname → **Storage**.
3. Click the volume group. The *Volume Group Overview* page opens.
4. Click the logical volume.
5. Click **Deactivate**.

Figure 7.3. The logical volume summary

7.5. DELETING A LOGICAL VOLUME USING THE WEB CONSOLE

Follow these instructions to delete a thinly- or thickly-provisioned logical volume.

1. Log in to the Web Console.
2. Click the hostname → **Storage**.
3. Click the volume group. The *Volume Group Overview* page opens.
4. Click the logical volume.
5. Click **Delete** in the logical volume summary.
6. Click **Delete** to confirm deletion.

7.6. GROWING A LOGICAL VOLUME USING THE WEB CONSOLE

Follow these instructions to increase the size of a logical volume using the Web Console.


1. Log in to the Web Console.
2. Click the hostname → **Storage**.
3. Click the volume group. The *Volume Group Overview* page opens.
4. Click the logical volume.
5. On the **Volume** subtab, click **Grow**.

Figure 7.4. Logical Volume section expanded



The *Grow Logical Volume* window opens.

Figure 7.5. The Grow Logical Volume window



- a. Specify the new **Size** of the logical volume.
- b. Click **Grow**.

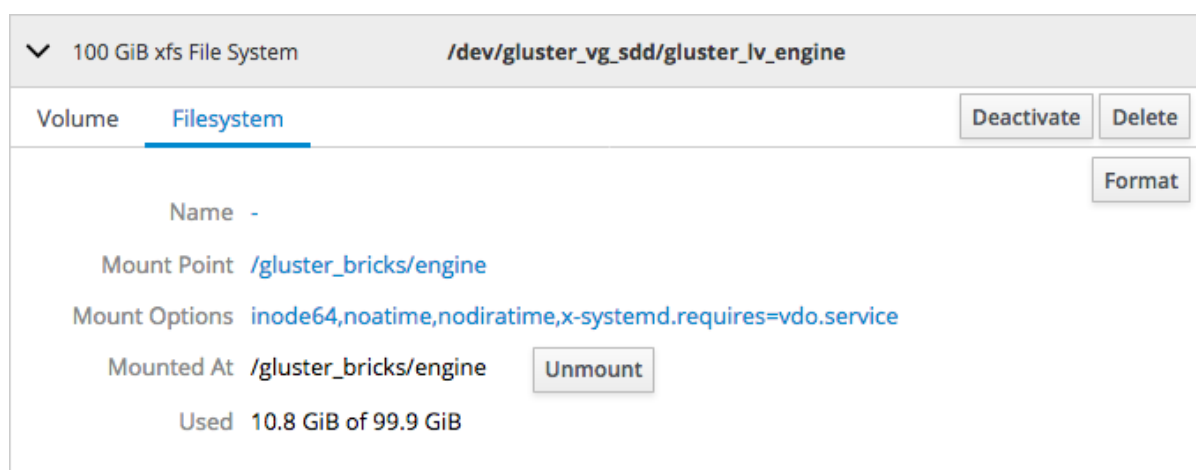
CHAPTER 8. MANAGING FILE SYSTEMS USING THE WEB CONSOLE

8.1. FORMATTING A LOGICAL VOLUME USING THE WEB CONSOLE

Follow these steps to format a logical volume with a file system using the Web Console.

1. Log in to the Web Console.
2. Click the hostname → **Storage**.
3. Click any volume group under **Volume Group**. The *Volume Group Overview* page opens.
4. Click the logical volume.
5. On the **Filesystem** subtab, click **Format**.

Figure 8.1. The logical volume summary



NOTE

If this logical volume does not have an existing file system, this tab is labelled **Unrecognised Data**.

- a. Specify whether to **Erase** existing data completely by overwriting it with zeroes.
- b. Specify the file system **Type** to use.
- c. Specify a **Name** for the file system.
- d. Specify whether to use default or customized **Mounting** behavior.
If you selected **Custom**, specify a **Mount Point** and check any **Mount options** you want this file system to use.
- e. Click **Format**.

8.2. CONFIGURING A MOUNT POINT USING THE WEB CONSOLE

1. Log in to the Web Console.

2. Click the hostname → **Storage**.
3. Click any volume group under **Volume Group**. The *Volume Group Overview* page opens.
4. Click the logical volume.
5. On the **Filesystem** subtab, click the value of **Mount Point**.

Figure 8.2. The logical volume summary



For this volume the mount point is **/gluster_bricks/engine**

The *Filesystem Mounting* window opens.

- a. Specify a behavior to use for **Mounting**.
If you select **Custom**, specify a **Mount Point** and check any **Mount Options** you want to use.
- b. Click **Apply**.

8.3. MOUNTING A FILE SYSTEM USING THE WEB CONSOLE

1. Log in to the Web Console.
2. Click the hostname → **Storage**.
3. Click any volume group under **Volume Group**. The *Volume Group Overview* page opens.
4. Click the logical volume.
5. On the **Filesystem** subtab, click **Mount**.

8.4. UNMOUNTING A FILE SYSTEM USING THE WEB CONSOLE

1. Log in to the Web Console.
2. Click the hostname → **Storage**.
3. Click any volume group under **Volume Group**. The *Volume Group Overview* page opens.
4. Click the logical volume.

5. On the **Filesystem** subtab, click **Unmount**.

Figure 8.3. The logical volume summary

The screenshot shows the logical volume summary for a 100 GiB xfs File System. The path is /dev/gluster_vg_sdd/gluster_lv_engine. The interface has two tabs: 'Volume' and 'Filesystem', with 'Filesystem' selected. There are buttons for 'Deactivate', 'Delete', and 'Format'. The summary includes the following information:

- Name -
- Mount Point /gluster_bricks/engine
- Mount Options inode64,noatime,nodiratime,x-systemd.requires=vdo.service
- Mounted At /gluster_bricks/engine with an 'Unmount' button.
- Used 10.8 GiB of 99.9 GiB

CHAPTER 9. MANAGING GLUSTER VOLUMES USING THE WEB CONSOLE

9.1. ACCESSING THE GLUSTER MANAGEMENT DASHBOARD

The *Gluster Management* dashboard lets you view information about the currently configured Gluster volumes in your hyperconverged cluster.

To access the *Gluster Management* dashboard:

1. Click **Virtualization** → **Hosted Engine** to reach the *Hosted Engine* dashboard.

The Hosted Engine dashboard in the Web Console

The screenshot shows the Hosted Engine dashboard. At the top, it indicates 'Hosted Engine is up!' with a green checkmark icon and 'Hosted Engine is running on rhsdev-grafton2.lab.eng.blr.redhat.com'. Below this, there is a section titled 'Status of this host (rhsdev-grafton2)' showing the host name 'rhsdev-grafton2' with a green checkmark. There are three buttons: 'Put this host into local maintenance', 'Remove this host from maintenance', and 'Put this cluster into global maintenance'. Below that is a section titled 'Hosts in this cluster' with a 'Manage Gluster' button. The hosts listed are:

Host Name	Agent stopped	Local Maintenance	VM Status	State
rhsdev-grafton2.l...	false	false	up	up
rhsdev-grafton3.l...	false	false	down_unexpected	down_unexpected
rhsdev-grafton4.l...	false	false	down	down

2. Click **Manage Gluster** to reach the *Gluster Management* dashboard.

The Gluster Management dashboard in the Web Console

Dashboard

Hosted Engine

Gluster Management

Hosts

Name	Peer Status
rhsdev-grafton2 ⓘ	CONNECTED
rhsdev-grafton2 ⓘ	CONNECTED
rhsdev-grafton2 ⓘ	CONNECTED

[Expand Cluster](#)

Volumes

Name	Cluster	Volume Type	Volume Status	Bricks	
engine	Default	REPLICATE	ONLINE	↕3 ↘0	ⓘ
plaindis	Default	DISTRIBUTE	ONLINE	↕1 ↘0	ⓘ
sharded-dis	Default	DISTRIBUTE	OFFLINE	↕0 ↘0	ⓘ
data	Default	REPLICATE	ONLINE	↕3 ↘0	ⓘ
vmstore	Default	REPLICATE	ONLINE	↕3 ↘0	ⓘ

[Create Volume](#)

9.2. EXPANDING VOLUME FROM WEB CONSOLE

Follow these instructions to use the Web Console to expand your volume.

Prerequisites

- Verify that your scaling plans are supported: [Requirements for scaling](#).

Procedure

1. Log in to the Web Console.
2. Click **Virtualization** → **Hosted Engine** and then click **Manage Gluster**.

Gluster Management

Hosts ⓘ [Expand Cluster](#)

Name	Peer status	UUID
10.70.36.79	CONNECTED	cc2fea6c-5e0d-4800-b10c-5692de7bfc79
rhsqa-grafton2-nic2.lab.eng.blr.redhat.com	CONNECTED	cd44a114-8804-437d-ab02-b159ebc7ce35
rhsqa-grafton3-nic2.lab.eng.blr.redhat.com	CONNECTED	fb0c96fe-5a41-4a95-988c-fee3c2de2d36

Volumes ⓘ [Create Volume](#)

Name	Volume Type	Status	More Info	
> engine	REPLICATE	🟢 ONLINE	ⓘ	Expand Volume
> data	REPLICATE (ARBITER)	🟢 ONLINE	ⓘ	Expand Volume
> vmstore	REPLICATE	🟢 ONLINE	ⓘ	Expand Volume

3. Click **Expand volume** button beside the volume you want to expand. The *expand volume* data page opens.

Expand Volume data [X]

Hosts Volumes Bricks Review

1 — 2 — 3 — 4

Host1:

Host2:

Host3:

i Ansible will login to gluster hosts as root user using passwordless ssh connections. Make sure, passwordless ssh is configured for all gluster hosts from the first host.

[Cancel] [< Back] [Next >]

4. On the *Hosts* tab, verify the Host details and click **Next**.
5. On the *Volumes* tab, specify the details of the brick path to be configured for the new disk.

Expand Volume data [X]

Hosts Volumes Bricks Review

1 — 2 — 3 — 4

Name	Volume Type	Arbiter	Brick Dirs
<input type="text" value="data"/>	<input type="text" value="Replicate"/>	<input checked="" type="checkbox"/>	<input type="text" value="/gluster_bricks/newdata/newdata"/>

[Cancel] [< Back] [Next >]

6. On the *Bricks* tab, specify the details of the disks to be used to expand the Gluster volume.

Expand Volume data
✕

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: 10.70.36.79

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

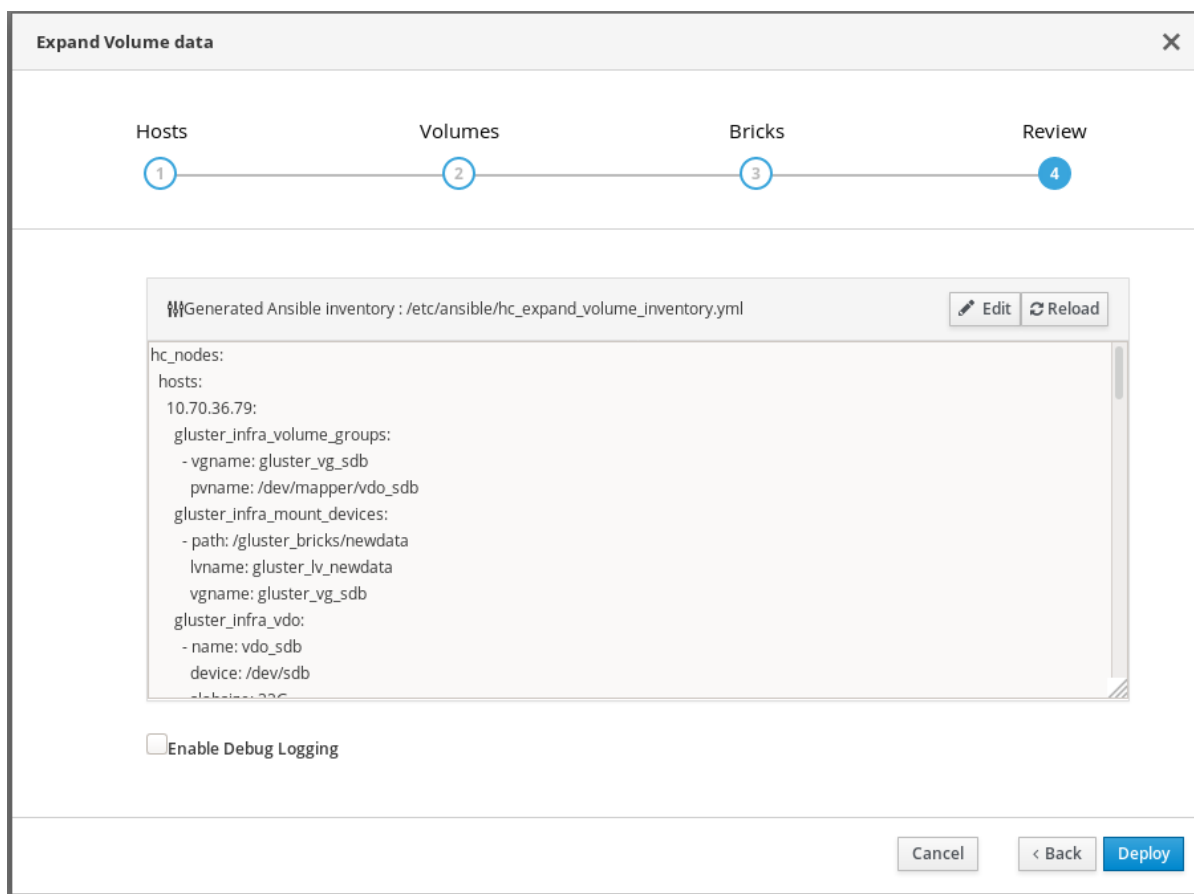
Configure LV Cache

ⓘ 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 >

7. On the *Review* tab, check the generated file for any problems. Here, **Enable debug logging**, runs ansible-playbook in verbose mode, and provides more logs to add information. When you are satisfied, click **Deploy**.



9.3. 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](#).
- Install three physical machines to serve as the new hyperconverged nodes. Follow the instructions in [Installing hyperconverged hosts](#).
- 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 pairs instead of passwords for SSH authentication](#) to configure key based 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.

9.4. 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](#).
- Install three physical machines to serve as the new hyperconverged nodes. Follow the instructions in [Installing hyperconverged hosts](#).
- 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 pairs instead of passwords for SSH authentication](#) to configure key based 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**.

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

- 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

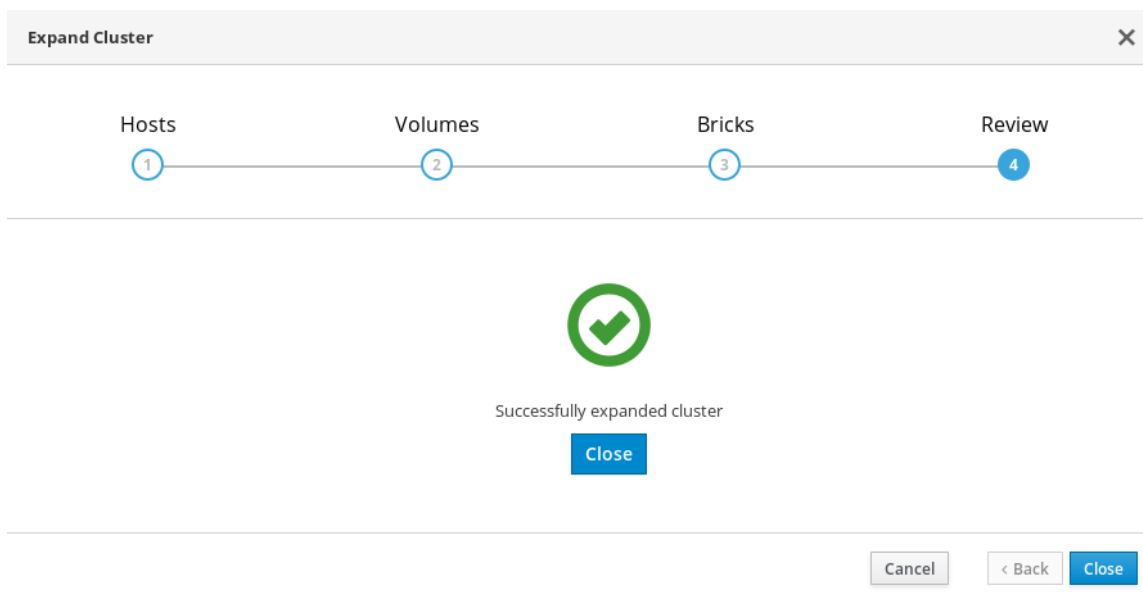
Generated Ansible inventory : /etc/ansible/hc_wizard_inventory.yml

```

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_thinpool:
- 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.



9.4.1. 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.

1. Click **Compute** → **Hosts** and then click **New** to open the *New Host* window.
2. Provide the **Name**, **Hostname**, and **Password** for the host that you want to manage.
3. Under **Advanced Parameters**, uncheck the **Automatically configure host firewall** checkbox, as firewall rules are already configured by the deployment process.
4. 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.
5. Click **OK**.
6. **Attach the gluster network to all remaining hosts**
 - a. Click the name of the newly added host to go to the host page.
 - 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** checkbox is checked.
 - e. Ensure that the **Save network configuration** checkbox is checked.
 - f. Click **OK** to save.
7. 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.

8. 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 4.4 *Self-Hosted Engine Guide* for further details:

https://access.redhat.com/documentation/en-us/red_hat_virtualization/4.4/html/self-hosted_engine_guide/chap-installing_additional_hosts_to_a_self-hosted_environment

9.5. CREATING AN ADDITIONAL GLUSTER VOLUME USING THE WEB CONSOLE

Follow these instructions to use the Web Console to create a new Red Hat Gluster Storage volume using raw disks that are available on hyperconverged hosts in your cluster.

Prerequisites

- Verify that the raw disk drives you plan to use for the new volume are visible under the *Drives* section of the *Storage Dashboard*, and do not have any file systems listed on their *Drive Overview* page.

Procedure

1. Log in to the Web Console.
2. Click **Virtualization** → **Hosted Engine** and then click **Manage Gluster**.
3. Click **Create Volume**. The *Create Volume* window opens.
 - a. On the *Hosts* tab, select three different hyperconverged hosts with unused disks and click **Next**.

Create Volume
✕

Hosts Volumes Bricks Review

1 ————— 2 ————— 3 ————— 4

Host1

Host2

Host3

i Ansible will login to gluster hosts as root user using passwordless ssh connections. Make sure, passwordless ssh is configured for all gluster hosts from the first host.

Cancel < Back Next >

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

Create Volume
✕

Hosts Volumes Bricks Review

1 ————— 2 ————— 3 ————— 4

Name	Volume Type	Arbiter	Brick Dirs
<input type="text" value="vol2"/>	<input type="text" value="Replicate"/>	<input type="checkbox"/>	<input type="text" value="/gluster_bricks/vol2/vol2"/> 🗑

[+ Add Volume](#)

Cancel < Back Next >

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

Create Volume
✕

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: host1.example.com

LV Name	Device Name	LV Size(GB)	Thinp	Mount Point	Enable Dedupe & Compression
vol2	/dev/sdb	500	<input checked="" type="checkbox"/>	/gluster_bricks/vol2	<input type="checkbox"/>

Configure LV Cache

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

Cancel
< Back
Next >

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

Create Volume
✕

Hosts
Volumes
Bricks
Review

1

2

3

4

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

```

            hc_nodes:
            hosts:
            host1.example.com:
            gluster_infra_volume_groups:
            - vgname: gluster_vg_sdb
            pvname: /dev/sdb
            gluster_infra_mount_devices:
            - path: /gluster_bricks/vol2
            lvname: gluster_lv_vol2
            vgname: gluster_vg_sdb
            gluster_infra_thinpoools:
            - vgname: gluster_vg_sdb
            thinpoolname: gluster_thinpool_gluster_vg_sdb
            
```

Enable Debug Logging

Cancel
< Back
Deploy

Successfully created volume is displayed when deployment completes successfully.

