



Red Hat Gluster Storage 3.2

Installation Guide

Installing Red Hat Gluster Storage 3.2

Red Hat Gluster Storage 3.2 Installation Guide

Installing Red Hat Gluster Storage 3.2

Bhavana Mohanraj
Red Hat Engineering Content Services
bmohanra@redhat.com

Anjana Suparna Sriram
Red Hat Engineering Content Services
asriram@redhat.com

Divya Muntimadugu
Red Hat Engineering Content Services
divya@redhat.com

Laura Bailey
Red Hat Customer Content Services
lbailey@redhat.com

Legal Notice

Copyright © 2015-2017 Red Hat, Inc.

This document is licensed by Red Hat under the [Creative Commons Attribution-ShareAlike 3.0 Unported License](#). If you distribute this document, or a modified version of it, you must provide attribution to Red Hat, Inc. and provide a link to the original. If the document is modified, all Red Hat trademarks must be removed.

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

Abstract

This guide describes the prerequisites and provides step-by-step instructions to install Red Hat Gluster Storage using different methods.

Table of Contents

CHAPTER 1. PLANNING RED HAT GLUSTER STORAGE INSTALLATION	4
1.1. ABOUT RED HAT GLUSTER STORAGE	4
1.2. PREREQUISITES	4
1.3. HARDWARE COMPATIBILITY	6
1.4. PORT INFORMATION	6
1.5. SUPPORTED VERSIONS OF RED HAT GLUSTER STORAGE	8
CHAPTER 2. INSTALLING RED HAT GLUSTER STORAGE	10
2.1. OBTAINING RED HAT GLUSTER STORAGE	10
2.2. INSTALLING FROM AN ISO IMAGE	11
2.3. INSTALLING RED HAT GLUSTER STORAGE SERVER ON RED HAT ENTERPRISE LINUX (LAYERED INSTALL)	18
2.4. INSTALLING FROM A PXE SERVER	20
2.5. INSTALLING FROM RED HAT SATELLITE SERVER	21
2.6. SUBSCRIBING TO THE RED HAT GLUSTER STORAGE SERVER CHANNELS	23
2.7. MANAGING THE GLUSTERD SERVICE	26
2.8. INSTALLING ANSIBLE TO SUPPORT GDEPLOY	27
CHAPTER 3. DEPLOYING SAMBA ON RED HAT GLUSTER STORAGE	28
3.1. PREREQUISITES	28
3.2. INSTALLING SAMBA USING ISO	29
3.3. INSTALLING SAMBA USING YUM	30
CHAPTER 4. DEPLOYING NFS-GANESHA ON RED HAT GLUSTER STORAGE	31
4.1. PREREQUISITES	31
4.2. INSTALLING NFS-GANESHA DURING AN ISO INSTALLATION	32
4.3. INSTALLING NFS-GANESHA USING YUM	32
CHAPTER 5. DEPLOYING GSTATUS ON RED HAT GLUSTER STORAGE	33
5.1. INSTALLING GSTATUS DURING AN ISO INSTALLATION	33
5.2. INSTALLING USING YUM OR THE RED HAT SATELLITE SERVER OR RED HAT NETWORK	33
CHAPTER 6. DEPLOYING CONTAINERIZED RED HAT GLUSTER STORAGE WITH ATOMIC HOST	34
6.1. SUPPORTED DEPLOYMENT CONFIGURATION	34
6.2. SUPPORTED VERSIONS	34
6.3. DOWNLOADING THE RED HAT GLUSTER STORAGE IMAGE FROM THE RED HAT REGISTRY	34
CHAPTER 7. UPDATING RED HAT GLUSTER STORAGE FROM 3.2.X TO 3.2.Y	36
7.1. UPDATING RED HAT GLUSTER STORAGE IN THE OFFLINE MODE	36
7.2. UPDATING THE NFS SERVER	39
7.3. IN-SERVICE SOFTWARE UPDATE FROM RED HAT GLUSTER STORAGE	45
CHAPTER 8. UPGRADING TO RED HAT GLUSTER STORAGE 3.2	49
8.1. OFFLINE UPGRADE FROM RED HAT GLUSTER STORAGE 3.1 TO 3.2	49
8.2. IN-SERVICE SOFTWARE UPGRADE FROM RED HAT GLUSTER STORAGE 3.1 TO RED HAT GLUSTER STORAGE 3.2	53
CHAPTER 9. UPGRADING FROM RED HAT GLUSTER STORAGE 3.1 TO RED HAT GLUSTER STORAGE 3.2 IN A RED HAT ENTERPRISE VIRTUALIZATION-RED HAT GLUSTER STORAGE ENVIRONMENT	67
9.1. UPGRADING USING AN ISO	67
9.2. UPGRADING USING YUM	69
CHAPTER 10. ENABLING SELINUX	71
CHAPTER 11. USING THE GLUSTER COMMAND LINE INTERFACE	73

APPENDIX A. REVISION HISTORY **75**

CHAPTER 1. PLANNING RED HAT GLUSTER STORAGE INSTALLATION

This chapter outlines the minimum hardware and software installation requirements for a successful installation, configuration, and operation of a Red Hat Gluster Storage Server environment.

1.1. ABOUT RED HAT GLUSTER STORAGE

Red Hat Gluster Storage is a software-only, scale-out storage that provides flexible and affordable unstructured data storage for the enterprise. Red Hat Gluster Storage 3.2 provides new opportunities to unify data storage and infrastructure, increase performance, and improve availability and manageability in order to meet a broader set of an organization's storage challenges and requirements.

GlusterFS, a key building block of Red Hat Gluster Storage, is based on a stackable user space design and can deliver exceptional performance for diverse workloads. GlusterFS aggregates various storage servers over network interconnects into one large parallel network file system. The POSIX compatible GlusterFS servers, which use XFS file system format to store data on disks, can be accessed using industry standard access protocols including NFS and CIFS.

Red Hat Gluster Storage can be deployed in the private cloud or datacenter using Red Hat Gluster Storage Server for On-Premise. Red Hat Gluster Storage can be installed on commodity servers and storage hardware resulting in a powerful, massively scalable, and highly available NAS environment. Additionally, Red Hat Gluster Storage can be deployed in the public cloud using Red Hat Gluster Storage Server for Public Cloud, for example, within the Amazon Web Services (AWS) cloud. It delivers all the features and functionality possible in a private cloud or datacenter to the public cloud by providing massively scalable and highly available NAS in the cloud.

Red Hat Gluster Storage Server for On-Premise

Red Hat Gluster Storage Server for On-Premise enables enterprises to treat physical storage as a virtualized, scalable, and centrally managed pool of storage by using commodity server and storage hardware.

Red Hat Gluster Storage Server for Public Cloud

Red Hat Gluster Storage Server for Public Cloud packages GlusterFS as an Amazon Machine Image (AMI) for deploying scalable NAS in the AWS public cloud. This powerful storage server provides a highly available, scalable, virtualized, and centrally managed pool of storage for Amazon users.

1.2. PREREQUISITES

Ensure that your environment meets the following requirements.

File System Requirements

XFS - Format the back-end file system using XFS for glusterFS bricks. XFS can journal metadata, resulting in faster crash recovery. The XFS file system can also be defragmented and expanded while mounted and active.



NOTE

Red Hat assists existing Gluster Storage Software Appliance customers using `ext3` or `ext4` to upgrade to a supported version of Red Hat Gluster Storage using the XFS back-end file system.

Logical Volume Manager

Format glusterFS bricks using XFS on the Logical Volume Manager to prepare for the installation.

Network Time Configuration

- Synchronize time across all Red Hat Gluster Storage servers using the Network Time Protocol (NTP) daemon.

1.2.1. Network Time Protocol Setup

Use a remote server over the Network Time Protocol (NTP) to synchronize the system clock. Set the `ntpd` daemon to automatically synchronize the time during the boot process as follows:

1. Edit the NTP configuration file `/etc/ntp.conf` using a text editor such as `vim` or `nano`.

```
# nano /etc/ntp.conf
```

2. Add or edit the list of public NTP servers in the `ntp.conf` file as follows:

```
server 0.rhel.pool.ntp.org
server 1.rhel.pool.ntp.org
server 2.rhel.pool.ntp.org
```

The Red Hat Enterprise Linux 6 version of this file already contains the required information. Edit the contents of this file if customization is required.

3. Optionally, increase the initial synchronization speed by appending the `iburst` directive to each line:

```
server 0.rhel.pool.ntp.org iburst
server 1.rhel.pool.ntp.org iburst
server 2.rhel.pool.ntp.org iburst
```

4. After the list of servers is complete, set the required permissions in the same file. Ensure that only `localhost` has unrestricted access:

```
restrict default kod nomodify notrap nopeer noquery
restrict -6 default kod nomodify notrap nopeer noquery
restrict 127.0.0.1
restrict -6 ::1
```

5. Save all changes, exit the editor, and restart the NTP daemon:

```
# service ntpd restart
```

6. Ensure that the `ntpd` daemon starts at boot time:

```
# chkconfig ntpd on
```

Use the `ntpdate` command for a one-time synchronization of NTP. For more information about this feature, see the *Red Hat Enterprise Linux Deployment Guide*

1.3. HARDWARE COMPATIBILITY

Hardware specifications change almost daily, it is recommended that all systems be checked for compatibility. The most recent list of supported hardware can be found in the *Red Hat Gluster Storage Server Compatible Physical, Virtual Server and Client OS Platforms List*, available online at <https://access.redhat.com/knowledge/articles/66206>. You must ensure that your environments meets the hardware compatibility outlined in this article. Hardware specifications change rapidly and full compatibility is not guaranteed.

Hardware compatibility is a particularly important concern if you have an older or custom-built system.

1.4. PORT INFORMATION

Red Hat Gluster Storage Server uses the listed ports. Ensure that firewall settings do not prevent access to these ports.

Table 1.1. TCP Port Numbers

Port Number	Usage
22	For sshd used by geo-replication.
111	For rpc port mapper.
139	For netbios service.
445	For CIFS protocol.
965	For NFS's Lock Manager (NLM).
2049	For glusterFS's NFS exports (nfsd process).
24007	For glusterd (for management).
24009 - 24108	For client communication with Red Hat Gluster Storage 2.0.
38465	For NFS mount protocol.
38466	For NFS mount protocol.
38468	For NFS's Lock Manager (NLM).
38469	For NFS's ACL support.
39543	For oVirt (Red Hat Gluster Storage-Console).

Port Number	Usage
49152 - 49251	For client communication with Red Hat Gluster Storage 2.1 and for brick processes depending on the availability of the ports. The total number of ports required to be open depends on the total number of bricks exported on the machine.
54321	For VDSM (Red Hat Gluster Storage Console).
55863	For oVirt (Red Hat Gluster Storage-Console).

Table 1.2. TCP Port Numbers used for Object Storage (Swift)

Port Number	Usage
443	For HTTPS request.
6010	For Object Server.
6011	For Container Server.
6012	For Account Server.
8080	For Proxy Server.

Table 1.3. TCP Port Numbers for Nagios Monitoring

Port Number	Usage
80	For HTTP protocol (required only if Nagios server is running on a Red Hat Gluster Storage node).
443	For HTTPS protocol (required only for Nagios server).
5667	For NSCA service (required only if Nagios server is running on a Red Hat Gluster Storage node).
5666	For NRPE service (required in all Red Hat Gluster Storage nodes).

Table 1.4. UDP Port Numbers

Port Number	Usage
111	For RPC Bind.

Port Number	Usage
963	For NFS's Lock Manager (NLM).

1.5. SUPPORTED VERSIONS OF RED HAT GLUSTER STORAGE

Several versions of Red Hat Gluster Storage (RHGS) are currently supported. In order to upgrade or use Red Hat Gluster Storage in conjunction with other software, you may need to know the support details of each version. This table provides a summary.

Table 1.5. Version Details

RHGS version	glusterfs and glusterfs-fuse	RHGS op-version	SMB	NFS	gDeploy
3.0	3.6.0.28-1	30000	SMB 1, 2.0	NFSv3	-
3.0 Update 1	3.6.0.29-1	30000	SMB 1, 2.0	NFSv3	gdeploy-1.0-12
3.0 Update 2	3.6.0.29-3	30000	SMB 1, 2.0	NFSv3	gdeploy-1.0-12
3.0 Update 3	3.6.0.42-1	30000	SMB 1, 2.0	NFSv3	gdeploy-1.0-12
3.0 Update 4	3.6.0.53-1	30004	SMB 1, 2.0, 2.1, 3.0	NFSv3, NFSv4	gdeploy-1.0-12
3.1	3.7.1-11	30703	SMB 1, 2.0, 2.1, 3.0	NFSv3, NFSv4	gdeploy-1.0-12
3.1 Update 1	3.7.1-16	30703	SMB 1, 2.0, 2.1, 3.0	NFSv3, NFSv4	gdeploy-1.0-12
3.1 Update 2	3.7.5-19	30707	SMB 1, 2.0, 2.1, 3.0	NFSv3, NFSv4	gdeploy-1.0-12
3.1 Update 3	3.7.9-12	30712	SMB 1, 2.0, 2.1, 3.0, 3.1.1	NFSv3, NFSv4	gdeploy-2.0-16
3.1 Update 3 (Async Security Fix)	3.7.9-12	30712	SMB 1, 2.0, 2.1, 3.0, 3.1.1	NFSv3, NFSv4	gdeploy-2.0.1-8
3.2	3.8.4-18.6	31001	SMB 1, 2.0, 2.1, 3.0, 3.1.1	NFSv3, NFSv4	gdeploy-2.0.1-13
3.2 (Async Bug Fix)	3.8.4-18.7	31001	SMB 1, 2.0, 2.1, 3.0, 3.1.1	NFSv3, NFSv4	gdeploy-2.0.1-13

Also, note that the repository that you install packages from may affect the package version installed.

This in turn may affect the compatibility of your systems. In general, Red Hat recommends installing the client from the Red Hat Gluster Storage repository (`rhel-x86_64-server-rhsc1ient-6`).

CHAPTER 2. INSTALLING RED HAT GLUSTER STORAGE

Red Hat Gluster Storage can be installed in a data center using Red Hat Gluster Storage Server On-Premise.

This chapter describes the three different methods for installing Red Hat Gluster Storage Server: using an ISO image, using a PXE server, or using the Red Hat Satellite Server.

For information on launching Red Hat Gluster Storage Server for Public Cloud, see the Red Hat Gluster Storage Administration Guide.



IMPORTANT

- Technology preview packages will also be installed with this installation of Red Hat Gluster Storage Server. For more information about the list of technology preview features, see chapter *Technology Previews* in the *Red Hat Gluster Storage 3.2 Release Notes*.
- When you clone a virtual machine that has Red Hat Gluster Storage Server installed, you need to remove the `/var/lib/glusterd/glusterd.info` file (if present) before you clone. If you do not remove this file, all cloned machines will have the same UUID. The file will be automatically recreated with a UUID on initial start-up of the `glusterd` daemon on the cloned virtual machines.

2.1. OBTAINING RED HAT GLUSTER STORAGE

This chapter details the steps to obtain the Red Hat Gluster Storage software.

2.1.1. Obtaining Red Hat Gluster Storage Server for On-Premise

Visit the **Software & Download Center** in the Red Hat Customer Service Portal (<https://access.redhat.com/downloads>) to obtain the Red Hat Gluster Storage Server for On-Premise installation *ISO image files*. Use a valid Red Hat Subscription to download the full installation files, obtain a free evaluation installation, or follow the links in this page to purchase a new Red Hat Subscription.

To download the Red Hat Gluster Storage Server installation files using a Red Hat Subscription or a Red Hat Evaluation Subscription:

1. Visit the Red Hat Customer Service Portal at <https://access.redhat.com/login> and enter your user name and password to log in.
2. Click **Downloads** to visit the **Software & Download Center**.
3. In the Red Hat Gluster Storage Server area, click **Download Software** to download the latest version of the software.

2.1.2. Obtaining Red Hat Gluster Storage Server for Public Cloud

Red Hat Gluster Storage Server for Public Cloud is pre-integrated, pre-verified, and ready to run the Amazon Machine Image (AMI). This AMI provides a fully POSIX-compatible, highly available, scale-out NAS and object storage solution for the Amazon Web Services (AWS) public cloud infrastructure.

For more information about obtaining access to AMI, see <https://access.redhat.com/knowledge/articles/145693>.

2.2. INSTALLING FROM AN ISO IMAGE

2.2.1. Installing Red Hat Gluster Storage 3.2 on Red Hat Enterprise Linux 6.7 and later

To install Red Hat Gluster Storage Server from the ISO image:

1. Download an ISO image file for Red Hat Gluster Storage Server as described in [Section 2.1, “Obtaining Red Hat Gluster Storage”](#)

The installation process launches automatically when you boot the system using the ISO image file.

Press **Enter** to begin the installation process.



NOTE

For some hypervisors, while installing Red Hat Gluster Storage on a virtual machine, you must select the **Install System with basic video driver** option.

2. The **Configure TCP/IP** screen displays.

To configure your computer to support TCP/IP, accept the default values for Internet Protocol Version 4 (IPv4) and Internet Protocol Version 6 (IPv6) and click **OK**. Alternatively, you can manually configure network settings for both Internet Protocol Version 4 (IPv4) and Internet Protocol Version 6 (IPv6).



IMPORTANT

NLM Locking protocol implementation in Red Hat Gluster Storage does not support clients over IPv6.




Figure 2.1. Configure TCP/IP

3. The **Welcome** screen displays.

Click **Next**.

4. The **Language Selection** screen displays. Select the preferred language for the installation and the system default and click **Next**.
5. The **Keyboard Configuration** screen displays. Select the preferred keyboard layout for the installation and the system default and click **Next**.

6. The **Storage Devices** screen displays. Select **Basic Storage Devices**.

Click **Next**.

7. The **Hostname** configuration screen displays.

Enter the hostname for the computer. You can also configure network interfaces if required. Click **Next**.

8. The **Time Zone Configuration** screen displays. Set your time zone by selecting the city closest to your computer's physical location.

9. The **Set Root Password** screen displays.

The root account's credentials will be used to install packages, upgrade RPMs, and perform most system maintenance. As such, setting up a root account and password is one of the most important steps in the installation process.



NOTE

The root user (also known as the superuser) has complete access to the entire system. For this reason, you should only log in as the root user to perform system maintenance or administration.

The **Set Root Password** screen prompts you to set a root password for your system. You cannot proceed to the next stage of the installation process without entering a root password.

Enter the root password into the **Root Password** field. The characters you enter will be masked for security reasons. Then, type the same password into the **Confirm** field to ensure the password is set correctly. After you set the root password, click **Next**.

10. The **Partitioning Type** screen displays.

Partitioning allows you to divide your hard drive into isolated sections that each behave as their own hard drive. Partitioning is particularly useful if you run multiple operating systems. If you are unsure how to partition your system, see *An Introduction to Disk Partitions* in *Red Hat Enterprise Linux 6 Installation Guide* for more information.

In this screen you can choose to create the default partition layout in one of four different ways, or choose to partition storage devices manually to create a custom layout.

If you are not comfortable partitioning your system, choose one of the first four options. These options allow you to perform an automated installation without having to partition your storage devices yourself. Depending on the option you choose, you can still control what data, if any, is removed from the system. Your options are:

- Use All Space
- Replace Existing Linux System(s)
- Shrink Current System
- Use Free Space
- Create Custom Layout

Choose the preferred partitioning method by clicking the radio button to the left of its description in the dialog box.

Click **Next** once you have made your selection. For more information on disk partitioning, see *Disk Partitioning Setup* in the *Red Hat Enterprise Linux 6 Installation Guide*



IMPORTANT

- It is recommended to create a separate `/var` partition that is large enough (50GB - 100GB) for log files, geo-replication related miscellaneous files, and other files.
- If a user does not select **Create Custom Layout**, all the connected/detected disks will be used in the Volume Group for the `/` and `/home` filesystems.

11. The **Boot Loader** screen displays with the default settings.

Click **Next**.

12. The **Minimal Selection** screen displays.

Click **Next** to retain the default selections and proceed with the installation.

- o To customize your package set further, select the **Customize now** option and click **Next**. This will take you to the **Customizing the Software Selection** screen.

Click **Next** to retain the default selections and proceed with the installation.

- o For Red Hat Gluster Storage 3.0.4 or later, if you require the Samba packages, ensure you select the **RH-Gluster-Samba-Server** component, in the **Customizing the Software Selection** screen. If you require samba active directory integration with gluster, ensure you select **RH-Gluster-AD-Integration** component. For NFS-Ganesha select **RH-Gluster-NFS-Ganesha**.

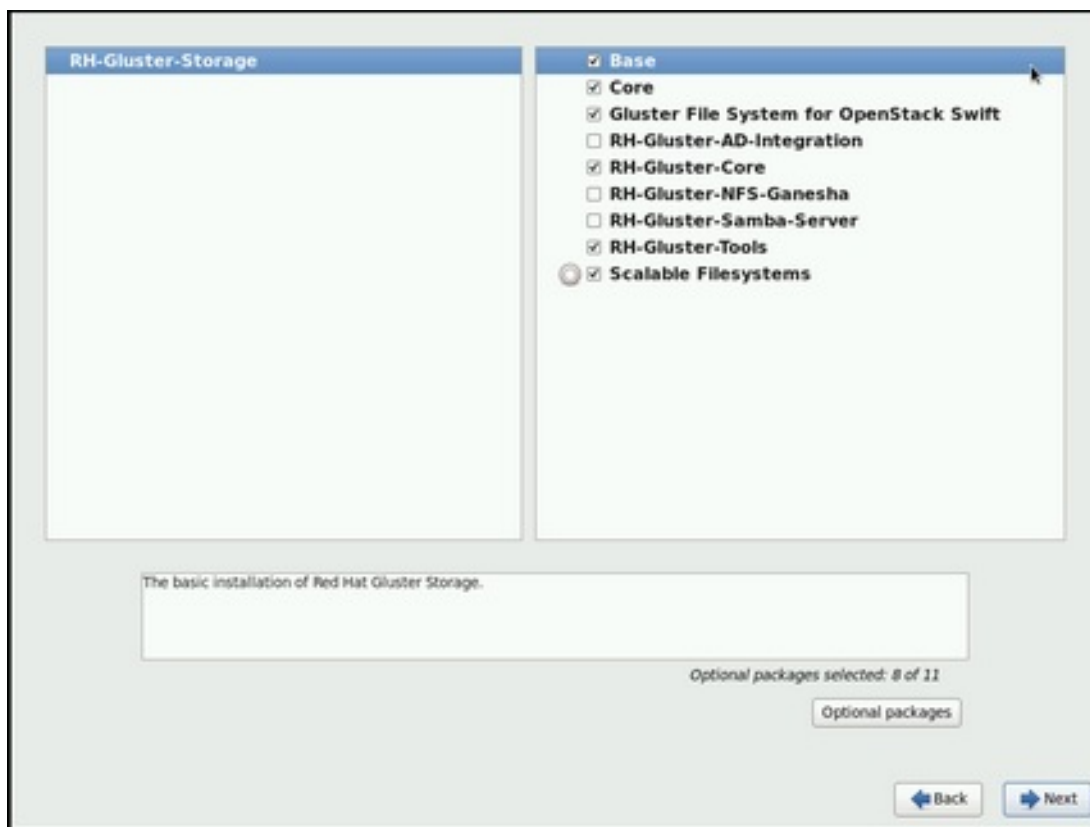


Figure 2.2. Customize Packages

13. The **Package Installation** screen displays.

Red Hat Gluster Storage Server reports the progress on the screen as it installs the selected packages in the system.

14. On successful completion, the **Installation Complete** screen displays.
15. Click **Reboot** to reboot the system and complete the installation of Red Hat Gluster Storage Server.

Ensure that you remove any installation media if it is not automatically ejected upon reboot.

Congratulations! Your Red Hat Gluster Storage Server installation is now complete.

For more information about installing Samba, see [Chapter 3, Deploying Samba on Red Hat Gluster Storage](#)

For more information about installing NFS-Ganesha, see [Chapter 4, Deploying NFS-Ganesha on Red Hat Gluster Storage](#).

2.2.2. Installing Red Hat Gluster Storage 3.2 on Red Hat Enterprise Linux 7.3

To install Red Hat Gluster Storage 3.2 on Red Hat Enterprise Linux 7.3:

1. Download an ISO image file for Red Hat Gluster Storage Server as described in [Section 2.1, “Obtaining Red Hat Gluster Storage”](#)
2. In the **Welcome to Red Hat Gluster Storage 3.2** screen, select the language that will be used for the rest of the installation and click **Continue**. This selection will also become the default for the installed system, unless changed later.



NOTE

One language is pre-selected by default on top of the list. If network access is configured at this point (for example, if you booted from a network server instead of local media), the pre-selected language will be determined based on automatic location detection using the GeoIP module.

3. The **Installation Summary** screen is the central location for setting up an installation.

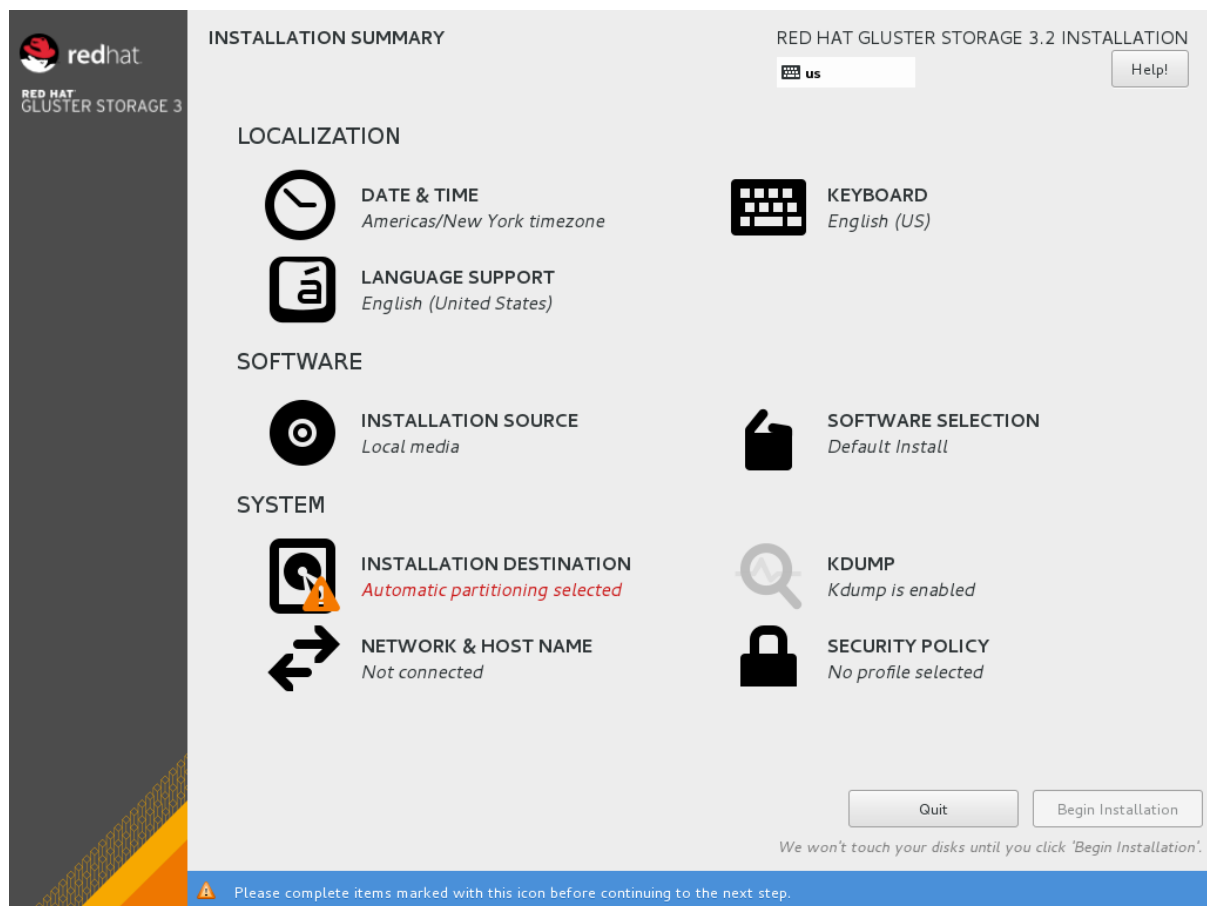


Figure 2.3. Installation Summary

Instead of directing you through consecutive screens, the Red Hat Gluster Storage 3.2 installation program on Red Hat Enterprise Linux 7.3 allows you to configure the installation in the order you choose.

Select a menu item to configure a section of the installation. When you have completed configuring a section, or if you would like to complete that section later, click the **Done** button located in the upper left corner of the screen.

Only sections marked with a warning symbol are mandatory. A note at the bottom of the screen warns you that these sections must be completed before the installation can begin. The remaining sections are optional. Beneath each section's title, the current configuration is summarized. Using this you can determine whether you need to visit the section to configure it further.

The following list provides a brief information of each of the menu item on the **Installation Summary** screen:

- **Date & Time**
To configure time zone, date, and optional settings for network time, select **Date & Time**.
- **Language Support**
To install support for additional locales and language dialects, select **Language Support**.
- **Keyboard Configuration**
To add multiple keyboard layouts to your system, select **Keyboard**.
- **Installation Source**
To specify a file or a location to install Red Hat Enterprise Linux from, select **Installation Source**. On this screen, you can choose between locally available installation media, such as a DVD or an ISO file, or a network location.
- **Network & Hostname**
To configure essential networking features for your system, select **Network & Hostname**.



IMPORTANT

When the Red Hat Gluster Storage 3.2 on Red Hat Enterprise Linux 7.3 installation finishes and the system boots for the first time, any network interfaces which you configured during the installation will be activated. However, the installation does not prompt you to configure network interfaces on some common installation paths - for example, when you install Red Hat Gluster Storage 3.2 on Red Hat Enterprise Linux 7.3 from a DVD to a local hard drive.

When you install Red Hat Gluster Storage 3.2 on Red Hat Enterprise Linux 7.3 from a local installation source to a local storage device, be sure to configure at least one network interface manually if you require network access when the system boots for the first time. You will also need to set the connection to connect automatically after boot when editing the configuration.

- **Software Selection**
To specify which packages will be installed, select **Software Selection**. If you require the following optional Add-Ons, then select the required Add-Ons and click **Done**:
 - RH-Gluster-AD-Integration
 - RH-Gluster-NFS-Ganesha

- RH-Gluster-Samba-Server

- **Installation Destination**

To select the disks and partition the storage space on which you will install Red Hat Gluster Storage, select **Installation Destination**. For more information on Installation Destination, see *Red Hat Enterprise Linux 7 Installation Guide*

- **Kdump**

Kdump is a kernel crash dumping mechanism which, in the event of a system crash, captures information that can be invaluable in determining the cause of the crash. Use this option to select whether or not to use Kdump on the system

4. After making the necessary configurations, click **Begin Installation** on the **Installation Summary** screen.



WARNING

Up to this point in the installation process, no lasting changes have been made on your computer. When you click **Begin Installation**, the installation program will allocate space on your hard drive and start to transfer Red Hat Gluster Storage into this space. Depending on the partitioning option that you chose, this process might include erasing data that already exists on your computer.

To revise any of the choices that you made up to this point, return to the relevant section of the **Installation Summary** screen. To cancel installation completely, click **Quit** or switch off your computer.

If you have finished customizing the installation and are certain that you want to proceed, click **Begin Installation**.

After you click **Begin Installation**, allow the installation process to complete. If the process is interrupted, for example, by you switching off or resetting the computer, or by a power outage, you will probably not be able to use your computer until you restart and complete the Red Hat Gluster Storage installation process

5. Once you click **Begin Installation**, the progress screen appears. Red Hat Gluster Storage reports the installation progress on the screen as it writes the selected packages to your system. Following is a brief description of the options on this screen:

- **Root Password**

The Root Password menu item is used to set the password for the root account. The root account is used to perform critical system management and administration tasks. The password can be configured either while the packages are being installed or afterwards, but you will not be able to complete the installation process until it has been configured.

- **User Creation**

Creating a user account is optional and can be done after installation, but it is recommended to do it on this screen. A user account is used for normal work and to access

the system. Best practice suggests that you always access the system via a user account and not the root account.

6. After the installation is completed, click **Reboot** to reboot your system and begin using Red Hat Gluster Storage.

2.3. INSTALLING RED HAT GLUSTER STORAGE SERVER ON RED HAT ENTERPRISE LINUX (LAYERED INSTALL)

Layered install involves installing Red Hat Gluster Storage over Red Hat Enterprise Linux



IMPORTANT

- Layered Installation is not available to all customers. Contact your Red Hat representative for more details on whether you can use it.
- It is recommended to create a separate `/var` partition that is large enough (50GB - 100GB) for log files, geo-replication related miscellaneous files, and other files.

1. Perform a base install of Red Hat Enterprise Linux Server

Red Hat Gluster Storage requires a kernel version of 2.6.32-431.17.1 or higher. This is available in Red Hat Enterprise Linux 6.5 and higher, and in Red Hat Enterprise Linux 7.

2. Register the System with Subscription Manager

Run the following command and enter your Red Hat Network user name and password to register the system with the Red Hat Network:

```
# subscription-manager register
```

3. Identify Available Entitlement Pools

Run the following commands to find entitlement pools containing the repositories required to install Red Hat Gluster Storage:

```
# subscription-manager list --available
```

4. Attach Entitlement Pools to the System

Use the pool identifiers located in the previous step to attach the **Red Hat Enterprise Linux Server** and **Red Hat Gluster Storage** entitlements to the system. Run the following command to attach the entitlements:

```
# subscription-manager attach --pool=[POOLID]
```

For example:

```
# subscription-manager attach --  
pool=8a85f9814999f69101499c05aa706e47
```

5. Enable the Required Channels

For Red Hat Gluster Storage 3.2 on Red Hat Enterprise Linux 6.7 and later

1. Run the following commands to enable the repositories required to install Red Hat Gluster Storage:

```
# subscription-manager repos --enable=rhel-6-server-rpms
# subscription-manager repos --enable=rhel-scalefs-for-rhel-6-server-rpms
# subscription-manager repos --enable=rhs-3-for-rhel-6-server-rpms
```

2. For Red Hat Gluster Storage 3.0.4 and later, if you require Samba, then enable the following repository:

```
# subscription-manager repos --enable=rh-gluster-3-samba-for-rhel-6-server-rpms
```

For more information about installing Samba, see [Chapter 3, Deploying Samba on Red Hat Gluster Storage](#)

3. If you require NFS-Ganesha then enable the following repository:

```
# subscription-manager repos --enable=rh-gluster-3-nfs-for-rhel-6-server-rpms
```

For more information about installing NFS-Ganesha, see [Chapter 4, Deploying NFS-Ganesha on Red Hat Gluster Storage](#).

For Red Hat Gluster Storage 3.2 on Red Hat Enterprise Linux 7.x

1. Run the following commands to enable the repositories required to install Red Hat Gluster Storage

```
# subscription-manager repos --enable=rhel-7-server-rpms
# subscription-manager repos --enable=rh-gluster-3-for-rhel-7-server-rpms
```

2. For Red Hat Gluster Storage 3.0.4 and later, if you require Samba, then enable the following repository:

```
# subscription-manager repos --enable=rh-gluster-3-samba-for-rhel-7-server-rpms
```

For more information about installing Samba, see [Chapter 3, Deploying Samba on Red Hat Gluster Storage](#)

3. If you require NFS-Ganesha then enable the following repository:

```
# subscription-manager repos --enable=rh-gluster-3-nfs-for-rhel-7-server-rpms
```

For more information about installing NFS-Ganesha, see [Chapter 4, Deploying NFS-Ganesha on Red Hat Gluster Storage](#).

6. Verify if the Channels are Enabled

Run the following command to verify if the channels are enabled:

```
# yum repolist
```

7. Kernel Version Requirement

Red Hat Gluster Storage requires the kernel-2.6.32-431.17.1.el6 version or higher to be used on the system. Verify the installed and running kernel versions by running the following command:


```
# rpm -q kernel
kernel-2.6.32-431.el6.x86_64
kernel-2.6.32-431.17.1.el6.x86_64
```

```
# uname -r
2.6.32-431.17.1.el6.x86_64
```

8. Update all packages

Ensure that all packages are up to date by running the following command.

```
# yum update
```



IMPORTANT

If any kernel packages are updated, reboot the system with the following command.

```
# shutdown -r now
```

9. Install Red Hat Gluster Storage

Run the following command to install Red Hat Gluster Storage:

```
# yum install redhat-storage-server
```

1. For Red Hat Gluster Storage 3.0.4 and later, if you require Samba, then execute the following command to install Samba:

```
# yum groupinstall RH-Gluster-Samba-Server
```

2. If you require Samba Active Directory integration with gluster, execute the following command:

```
# yum groupinstall RH-Gluster-AD-Integration
```

3. To install NFS-Ganesha, see [Chapter 4, Deploying NFS-Ganesha on Red Hat Gluster Storage](#)

10. Reboot

Reboot the system.

2.4. INSTALLING FROM A PXE SERVER

To boot your computer using a PXE server, you need a properly configured server and a network interface in your computer that supports PXE.

Configure the computer to boot from the network interface. This option is in the BIOS, and may be labeled **Network Boot** or **Boot Services**. Once you properly configure PXE booting, the computer can boot the Red Hat Gluster Storage Server installation system without any other media.

To boot a computer from a PXE server:

1. Ensure that the network cable is attached. The link indicator light on the network socket should be lit, even if the computer is not switched on.
2. Switch on the computer.
3. A menu screen appears. Press the number key that corresponds to the preferred option.

If your computer does not boot from the netboot server, ensure that the BIOS is configured so that the computer boots first from the correct network interface. Some BIOS systems specify the network interface as a possible boot device, but do not support the PXE standard. See your hardware documentation for more information.

2.5. INSTALLING FROM RED HAT SATELLITE SERVER

Ensure that the firewall settings are configured so that the required ports are open. For a list of port numbers, see [Section 1.4, “Port Information”](#).

Creating the Activation Key

For more information on how to create an activation key, see *Activation Keys* in the *Red Hat Network Satellite Reference Guide*.

- In the **Details** tab of the **Activation Keys** screen, select **Red Hat Enterprise Linux Server (v. 6 for 64-bit x86_64)** from the **Base Channels** drop-down list.

The screenshot shows the Red Hat Satellite web interface. At the top, there's a navigation bar with 'Overview', 'Systems', 'Emails', 'Channels', 'Audit', 'Configuration', 'Schedules', 'Users', 'Admin', and 'Help'. A red banner at the top right says 'NO SYSTEMS SELECTED' with 'MANAGE' and 'CLEAR' buttons. Below this, a blue banner says 'Activation key RHS3_ActivationKey has been created.' with a 'delete key' button. The main content area is titled 'RHS3_ActivationKey' and has tabs for 'Details', 'Child Channels', 'Packages', 'Groups', and 'Activated Systems'. The 'Details' tab is active, showing 'Activation Key Details'. A note says 'Systems registered with this activation key will inherit the settings listed below.' The form fields are:

- Description:** RHS3_ActivationKey (with a tip: 'Use this to describe what kind of settings this key will reflect on systems that use it. If left blank, this field will be filled in "None"').
- Key:** 1- [a6adcaaf4d1d15b7c75fd77] (with a tip: 'Leave blank for automatic key generation. Note that the prefix is an indication of the Red Hat Satellite organization the key is associated with.')
- Usage:** (with a tip: 'Leave blank for unlimited use.')
- Base Channels:** Red Hat Enterprise Linux Server (v. 6 for 64-bit x86_64) (with a tip: 'Choose "Red Hat Satellite Default" to allow systems to register to the default Red Hat provided channel that corresponds to their installed version of Red Hat Enterprise Linux. You may also choose particular Red Hat provided channels or custom base channels here, but please note if a system using this key is not compatible with the selected channel, it will fall back to its Red Hat default channel.')
- Add-On Entitlements:** Monitoring, Provisioning, Virtualization, Virtualization Platform (all unchecked).
- Configuration File Deployment:** (unchecked, with a tip: 'Configuration File Deployment is currently disabled. Select "Provisioning" under "Add-On Entitlements" to enable configuration management facilities for this activation key.')
- Universal Default:** (unchecked, with a tip: 'Only one universal default activation key may be set for this organization. By setting this key as universal default, you will remove universal default status from the current universal default key if it exists. If this key is set as universal default, then newly-registered systems to your organization will inherit the properties of this key.')

 At the bottom right, there is an 'Update Activation Key' button.

Figure 2.4. Base Channels

- In the **Child Channels** tab of the **Activation Keys** screen, select the following child channels:

RHEL Server Scalable File System (v. 6 for x86_64)
Red Hat Gluster Storage Server 3 (RHEL 6 for x86_64)

For Red Hat Gluster Storage 3.0.4 or later, if you require the Samba package, then select the following child channel:

Red Hat Gluster 3 Samba (RHEL 6 for x86_64)

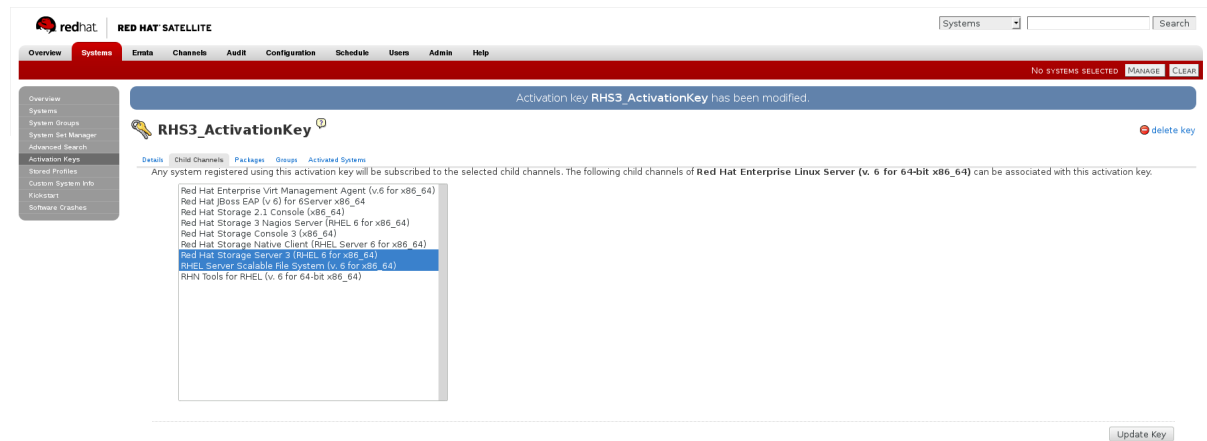


Figure 2.5. Child Channels

- In the **Packages** tab of the **Activation Keys** screen, enter the following package name:

redhat-storage-server

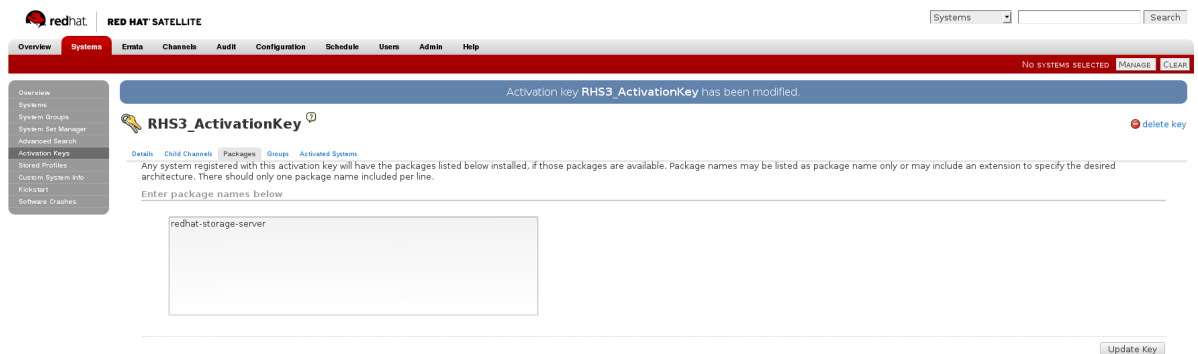


Figure 2.6. Package

- For Red Hat Gluster Storage 3.0.4 or later, if you require the Samba package, then enter the following package name:

samba

Creating the Kickstart Profile

For more information on creating a kickstart profile, see *Kickstart* in the *Red Hat Network Satellite Reference Guide*.

- When creating a kickstart profile, the following **Base Channel** and **Tree** must be selected.

Base Channel: Red Hat Enterprise Linux Server (v.6 for 64-bit x86_64)

Tree: ks-rhel-x86_64-server-6-6.5

- Do not associate any child channels with the kickstart profile.
- Associate the previously created activation key with the kickstart profile.



IMPORTANT

- By default, the kickstart profile chooses **md5** as the hash algorithm for user passwords.

You must change this algorithm to **sha512** by providing the following settings in the **auth** field of the **Kickstart Details, Advanced Options** page of the kickstart profile:

```
--enablshadow --passalgo=sha512
```

- After creating the kickstart profile, you must change the root password in the **Kickstart Details, Advanced Options** page of the kickstart profile and add a root password based on the prepared sha512 hash algorithm.

Installing Red Hat Gluster Storage Server using the Kickstart Profile

For more information on installing Red Hat Gluster Storage Server using a kickstart profile, see *Kickstart* in *Red Hat Network Satellite Reference Guide*

2.6. SUBSCRIBING TO THE RED HAT GLUSTER STORAGE SERVER CHANNELS

After you have successfully installed Red Hat Gluster Storage, you must subscribe to the required channels:



NOTE

If you used Red Hat Satellite or Layered Installation method to install Red Hat Gluster Storage 3.0, then you can skip the steps mentioned in this chapter. For more information regarding layered installation, see [Section 2.3, “Installing Red Hat Gluster Storage Server on Red Hat Enterprise Linux \(Layered Install\)”](#)

Using Subscription Manager

1. Register the System with Subscription Manager

Run the following command and enter your Red Hat Network user name and password to register the system with Subscription Manager:

```
# subscription-manager register
```

2. Identify Available Entitlement Pools

Run the following commands to find entitlement pools containing the repositories required to install Red Hat Gluster Storage:

```
# subscription-manager list --available | grep -A8 "Red Hat Enterprise Linux Server"
# subscription-manager list --available | grep -A8 "Red Hat Storage"
```

3. Attach Entitlement Pools to the System

Use the pool identifiers located in the previous step to attach the **Red Hat Enterprise Linux Server** and **Red Hat Gluster Storage** entitlements to the system. Run the following command to attach the entitlements:

```
# subscription-manager attach --pool=[POOLID]
```

For example:

```
# subscription-manager attach --  
pool=8a85f9814999f69101499c05aa706e47
```

4. Enable the Required Channels for Red Hat Gluster Storage on Red Hat Enterprise Linux

For Red Hat Gluster Storage 3.2 on Red Hat Enterprise Linux 6.7 and later

1. Run the following commands to enable the repositories required to install Red Hat Gluster Storage:

```
# subscription-manager repos --enable=rhel-6-server-rpms  
# subscription-manager repos --enable=rhel-scalefs-for-rhel-6-  
server-rpms  
# subscription-manager repos --enable=rhs-3-for-rhel-6-server-  
rpms
```

2. For Red Hat Gluster Storage 3.0.4 or later, if you require Samba, then enable the following repository:

```
# subscription-manager repos --enable=rh-gluster-3-samba-for-  
rhel-6-server-rpms
```

3. For Red Hat Gluster Storage 3.2, if NFS-Ganesha is required, then enable the following repository:

```
# subscription-manager repos  
--enable=rh-gluster-3-nfs-for-rhel-6-server-rpms
```

For Red Hat Gluster Storage 3.2 on Red Hat Enterprise Linux 7.3

1. Run the following commands to enable the repositories required to install Red Hat Gluster Storage

```
# subscription-manager repos --enable=rhel-7-server-rpms  
# subscription-manager repos --enable=rh-gluster-3-for-rhel-7-  
server-rpms
```

2. For Red Hat Gluster Storage 3.0.4 or later, if you require Samba, then enable the following repository:

```
# subscription-manager repos --enable=rh-gluster-3-samba-for-  
rhel-7-server-rpms
```

3. For Red Hat Gluster Storage 3.2, if NFS-Ganesha is required, then enable the following repository:

```
# subscription-manager repos
--enable=rh-gluster-3-nfs-for-rhel-7-server-rpms
```

4. For Red Hat Gluster Storage 3.2, if you require CTDB, then enable the following repository:

```
# subscription-manager repos --enable=rh-gluster-3-samba-for-
rhel-7-server-rpms
```

5. Verify if the Channels are Enabled

Run the following command to verify if the channels are enabled:

```
# yum repolist
```

Using Red Hat Satellite Server

1. Configure the Client System to Access Red Hat Satellite

Configure the client system to access Red Hat Satellite. Refer section *Registering Clients with Red Hat Satellite Server* in *Red Hat Satellite 5.6 Client Configuration Guide*

2. Register to the Red Hat Satellite Server

Run the following command to register the system to the Red Hat Satellite Server:

```
# rhn_register
```

3. Register to the Standard Base Channel

In the select operating system release page, select **All available updates** and follow the prompts to register the system to the standard base channel for RHEL6 - `rhel-x86_64-server-6`. The standard base channel for RHEL 7 is `RHEL7-rhel-x86_64-server-7`

4. Subscribe to the Required Red Hat Gluster Storage Server Channels

For Red Hat Gluster Storage 3.2 on Red Hat Enterprise Linux 6.7 and later

1. Run the following command to subscribe the system to the required Red Hat Gluster Storage server channel:

```
# rhn-channel --add --channel rhel-x86_64-server-6-rhs-3 --
channel rhel-x86_64-server-sfs-6
```

2. For Red Hat Gluster Storage 3.0.4 or later, if you require Samba, then execute the following command to enable the required channel:

```
# rhn-channel --add --channel rhel-x86_64-server-6-rh-gluster-3-
samba
```

3. For Red Hat Gluster Storage 3.2, for NFS-Ganesha enable the following channel:

```
# rhn-channel --add --channel
rhel-x86_64-server-6-rh-gluster-3-nfs
```

For Red Hat Gluster Storage 3.2 on Red Hat Enterprise Linux 7.3

1. Run the following command to subscribe the system to the required Red Hat Gluster Storage server channels for RHEL 7:

```
# rhn-channel --add --channel rhel-x86_64-server-7-rh-gluster-3
```

2. For Red Hat Gluster Storage 3.0.4 or later, if you require Samba, then execute the following command to enable the required channel:

```
# rhn-channel --add --channel rhel-x86_64-server-7-rh-gluster-3-samba
```

3. For Red Hat Gluster Storage 3.2, for NFS-Ganesha enable the following channel:

```
# rhn-channel --add --channel rhel-x86_64-server-7-rh-gluster-3-nfs
```

4. For Red Hat Gluster Storage 3.2, if CTDB is required, then enable the following channel:

```
# rhn-channel --add --channel rhel-x86_64-server-7-rh-gluster-3-samba
```

5. Verify if the System is Registered Successfully

Run the following command to verify if the system is registered successfully:

```
# rhn-channel --list
rhel-x86_64-server-7
rhel-x86_64-server-7-rh-gluster-3
```

2.7. MANAGING THE GLUSTERD SERVICE

After installing Red Hat Gluster Storage, the `glusterd` service automatically starts on all the servers in the trusted storage pool. The service can be manually started and stopped using the `glusterd` service commands.

Use Red Hat Gluster Storage to dynamically change the configuration of glusterFS volumes without restarting servers or remounting volumes on clients. The glusterFS daemon `glusterd` also offers elastic volume management.

Use the `gluster` CLI commands to decouple logical storage volumes from physical hardware. This allows the user to grow, shrink, and migrate storage volumes without any application downtime. As storage is added to the cluster, the volumes are distributed across the cluster. This distribution ensures that the cluster is always available despite changes to the underlying hardware.

2.7.1. Manually Starting and Stopping glusterd

Use the following instructions to manually start and stop the `glusterd` service.

- Manually start `glusterd` as follows:

```
# /etc/init.d/glusterd start
```

or

```
# service glusterd start
```

- Manually stop **glusterd** as follows:

```
# /etc/init.d/glusterd stop
```

or

```
# service glusterd stop
```

2.8. INSTALLING ANSIBLE TO SUPPORT GDEPLOY



NOTE

Consult with your IT department to confirm your organization's supported download instructions for Ansible.

gDeploy depends on Ansible to execute the playbooks and modules. You must install Ansible 1.9 or higher to use gDeploy.

CHAPTER 3. DEPLOYING SAMBA ON RED HAT GLUSTER STORAGE

Red Hat Gluster Storage provides a more recent version of Samba than the one shipped with Red Hat Enterprise Linux 6.6. This allows Red Hat Gluster Storage to take advantage of the latest features and enhancements. It includes a plug-in for directly accessing Red Hat Gluster Storage server.

3.1. PREREQUISITES

To install Samba on Red Hat Gluster Storage you require access to the installation media either through an ISO or a properly configured software repository. The Red Hat Gluster Storage server requirements are:

1. You must install Red Hat Gluster Storage Server 3.0.4 on the target server.



WARNING

- For layered installation of Red Hat Gluster Storage, ensure to have only the default Red Hat Enterprise Linux server installation, without the Samba or CTDB packages installed from Red Hat Enterprise Linux.
- The Samba version 3 is being deprecated from Red Hat Gluster Storage 3.0 Update 4. Further updates will not be provided for samba-3.x. It is recommended that you upgrade to Samba-4.x, which is provided in a separate channel or repository, for all updates including the security updates.
- CTDB version 2.5 is not supported from Red Hat Gluster Storage 3.1 Update 2. To use CTDB in Red Hat Gluster Storage 3.1.2 and later, you must upgrade the system to CTDB 4.x, which is provided in the Samba channel of Red Hat Gluster Storage.
- Downgrade of Samba from Samba 4.x to Samba 3.x is not supported.
- Ensure that Samba is upgraded on all the nodes simultaneously, as running different versions of Samba in the same cluster will lead to data corruption.

2. Enable the channel where the Samba packages are available:

For Red Hat Gluster Storage 3.2 on Red Hat Enterprise Linux 6.x

1. If you have registered your machine using Red Hat Subscription Manager, enable the repository by running the following command:

```
# subscription-manager repos --enable=rh-gluster-3-samba-for-rhel-6-server-rpms
```


2. If you have registered your machine using Satellite server, enable the channel by running the following command:

```
# rhn-channel --add --channel rhel-x86_64-server-6-rh-gluster-3-samba
```

For Red Hat Gluster Storage 3.2 on Red Hat Enterprise Linux 7.x

1. If you have registered your machine using Red Hat Subscription Manager, enable the repository by running the following command:

```
# subscription-manager repos --enable=rh-gluster-3-samba-for-rhel-7-server-rpms
```

2. If you have registered your machine using Satellite server, enable the channel by running the following command:

```
# rhn-channel --add --channel rhel-x86_64-server-7-rh-gluster-3-samba
```

3.2. INSTALLING SAMBA USING ISO

During the installation of Red Hat Gluster Storage, ensure you select the **RH-Gluster-Samba-Server** component, in the **Customizing the Software Selection** screen:

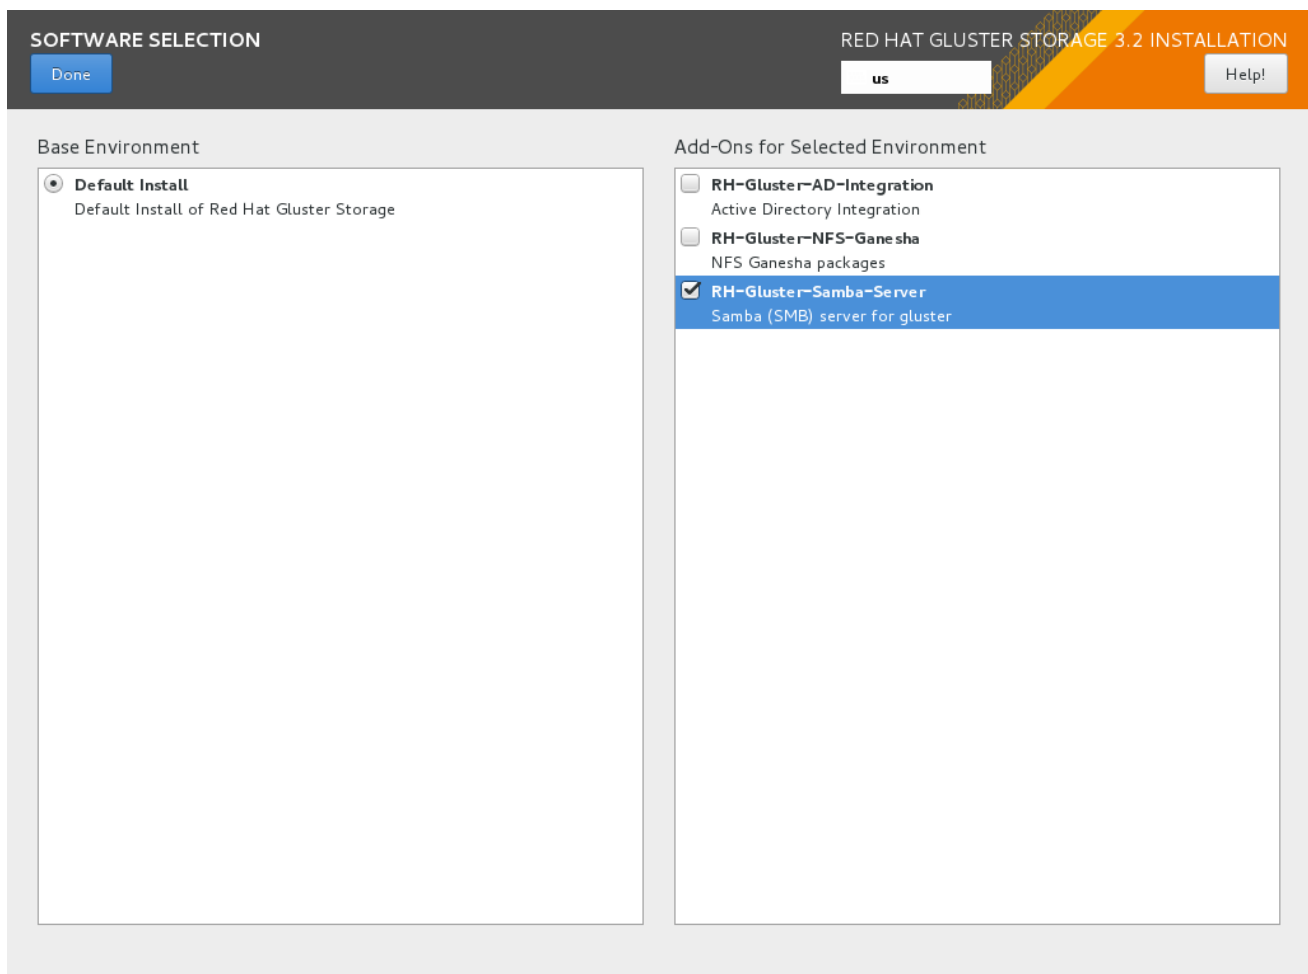


Figure 3.1. Customize Packages

For more information about installing Red Hat Gluster Storage using an ISO, see [Section 2.2, “Installing from an ISO Image”](#)

3.3. INSTALLING SAMBA USING YUM

To install Samba using yum, install the Samba group using the following command:

```
# yum groupinstall RH-Gluster-Samba-Server
```

If you require Samba Active Directory integration with gluster, execute the following command:

```
# yum groupinstall RH-Gluster-AD-Integration
```

For more information about installing Red Hat Gluster Storage using yum, see [Section 2.3, “Installing Red Hat Gluster Storage Server on Red Hat Enterprise Linux \(Layered Install\)”](#).

Samba can also be installed by following these steps:

1. To install the basic Samba packages, execute the following command:

```
# yum install samba
```

2. If you require the `smbclient` on the server, then execute the following command:

```
# yum install samba-client
```

3. If you require an Active directory setup, then execute the following commands:

```
# yum install samba-winbind  
# yum install samba-winbind-clients  
# yum install samba-winbind-krb5-locator
```

4. Verify if the following packages are installed.

```
samba-libs  
samba-winbind-krb5-locator  
samba-winbind-modules  
samba-vfs-glusterfs  
samba-winbind  
samba-client  
samba-common  
samba-winbind-clients  
samba
```

CHAPTER 4. DEPLOYING NFS-GANESHA ON RED HAT GLUSTER STORAGE

NFS-Ganesha can be installed using any of the following methods:

- Installing NFS-Ganesha using yum
- Installing NFS-Ganesha during an ISO Installation

4.1. PREREQUISITES

To install NFS-Ganesha on Red Hat Gluster Storage you require access to the installation media either through an ISO or a properly configured software repository.



NOTE

Ensure to install the required packages that are mentioned in the following sections on all the nodes in the Trusted Storage Pool.

For Red Hat Gluster Storage 3.2 on Red Hat Enterprise Linux 6.x

Enable the channel where the NFS-Ganesha packages are available:

1. If you have registered your machine using Red Hat Subscription Manager, enable the repository by running the following command:

```
# subscription-manager repos
--enable=rh-gluster-3-nfs-for-rhel-6-server-rpms
```

2. If you have registered your machine using Satellite server, enable the channel by running the following command:

```
# rhn-channel --add --channel
rhel-x86_64-server-6-rh-gluster-3-nfs
```

3. To subscribe to the HA channel, execute the following command:

```
# rhn-channel --add --channel rhel-x86_64-server-ha-6
```

4. To add the HA repository, execute the following command:

```
# subscription-manager repos --enable=rhel-6-server-rpms --
enable=rhel-ha-for-rhel-6-server-rpms
```

For Red Hat Gluster Storage 3.2 on Red Hat Enterprise Linux 7.x

Enable the channel where the NFS-Ganesha packages are available:

1. If you have registered your machine using Red Hat Subscription Manager, enable the repository by running the following command:

```
# subscription-manager repos
--enable=rh-gluster-3-nfs-for-rhel-7-server-rpms
```

2. If you have registered your machine using Satellite server, enable the channel by running the following command:

```
# rhn-channel --add --channel  
rhel-x86_64-server-7-rh-gluster-3-nfs
```

3. To subscribe to the HA channel, execute the following command:

```
# rhn-channel --add --channel rhel-x86_64-server-ha-7
```

4. To add the HA repository, execute the following command:

```
# subscription-manager repos --enable=rhel-7-server-rpms --  
enable=rhel-ha-for-rhel-7-server-rpms
```

4.2. INSTALLING NFS-GANESHA DURING AN ISO INSTALLATION

For more information about installing Red Hat Gluster Storage using an ISO image, see *Installing from an ISO Image*.

1. While installing Red Hat Storage using an ISO, in the Customizing the Software Selection screen, select **RH-Gluster-NFS-Ganesha** and click **Next**.
2. Proceed with the remaining installation steps for installing Red Hat Gluster Storage. For more information on how to install Red Hat Storage using an ISO, see *Installing from an ISO Image*.

4.3. INSTALLING NFS-GANESHA USING YUM

For more information about installing Red Hat Gluster Storage using yum , see *Installing Red Hat Gluster Storage Server on Red Hat Enterprise Linux*.

Execute the following steps to install NFS-Ganesha:

1. The glusterfs-ganesha package can be installed using the following command:

```
# yum install glusterfs-ganesha
```

NFS-Ganesha is installed along with the above package. nfs-ganesha-gluster and HA packages are also installed.

CHAPTER 5. DEPLOYING GSTATUS ON RED HAT GLUSTER STORAGE

Gstatus can be installed using any of the following methods:

- Installing gstatus during an ISO Installation
- Installing using yum or the Red Hat Satellite Server or Red Hat Network

5.1. INSTALLING GSTATUS DURING AN ISO INSTALLATION

For more information about installing Red Hat Gluster Storage using an ISO image, see *Installing from an ISO Image*.

1. While installing Red Hat Storage using an ISO, in the **Customizing the Software Selection** screen, select **Red Hat Storage Tools Group** and click **Optional Packages**.
2. From the list of packages, select **gstatus** and click **Close**.
3. Proceed with the remaining installation steps for installing Red Hat Storage. For more information on how to install Red Hat Storage using an ISO, see *Installing from an ISO Image*

5.2. INSTALLING USING YUM OR THE RED HAT SATELLITE SERVER OR RED HAT NETWORK

The gstatus package can be installed using the following command:

```
# yum install gstatus
```



NOTE

If you are installing using the Red Hat Network or Satellite, ensure that your system is subscribed to the required channels.

To verify the installation execute the following command:

```
# yum list gstatus
```

```
Installed Packages  
gstatus.x86_640.65-1@rhs-3-for-rhel-6-server-rpms
```

CHAPTER 6. DEPLOYING CONTAINERIZED RED HAT GLUSTER STORAGE WITH ATOMIC HOST

Red Hat Gluster Storage environment can be set up in a container on a Red Hat Enterprise Linux Atomic Host 7.2.

This chapter provides step-by-step instructions for users to deploy containerized Red Hat Gluster Storage on a Atomic Host.

6.1. SUPPORTED DEPLOYMENT CONFIGURATION

You must install one Red Hat Gluster Storage Container per Atomic Host and configure four bricks per Atomic Host.

6.2. SUPPORTED VERSIONS

The following table lists the supported versions of Docker and Atomic Host with Red Hat Gluster Storage Server.

Table 6.1. Supported Versions

Red Hat Gluster Storage Server	Docker	Red Hat Enterprise Linux Atomic Host
3.2	1.13	7.3

For information on installing and upgrading to Red Hat Enterprise Linux Atomic Host 7.2, see <https://access.redhat.com/documentation/en/red-hat-enterprise-linux-atomic-host/version-7/installation-and-configuration-guide/>.

6.3. DOWNLOADING THE RED HAT GLUSTER STORAGE IMAGE FROM THE RED HAT REGISTRY

A Docker registry provides a place to store and share docker containers that are saved as images that can be shared with other people. With the docker package available with Red Hat Enterprise Linux Atomic Host, you can pull images from the Red Hat Customer Portal. You see what images are available to pull from the Red Hat Customer Portal (using docker pull) by searching the Red Hat Container Images Search Page at <https://access.redhat.com/search/browse/container-images>.

1. You can verify the docker service is running with the following command:

```
# systemctl status docker
docker.service - Docker Application Container Engine
Loaded: loaded (/usr/lib/systemd/system/docker.service; enabled)
Active: active (running) since Fri 2017-01-22 12:03:34 CEST; 3h
19min ago
Docs: http://docs.docker.com
Main PID: 843 (docker)
CGroup: /system.slice/docker.service
        └─843 /usr/bin/docker -d --selinux-enabled
```

2. Download the Red Hat Gluster Storage Container images from the Red Hat Docker Registry using the following command:

```
# docker pull registry.access.redhat.com/rhgs3/rhgs-server-rhel7
```

3. Verify that you have a Red Hat Gluster Storage image on a Red Hat Enterprise Linux Atomic Host 7.2 by checking the contents of the `/etc/redhat-release` file. First, on the host system:

```
# cat /etc/redhat-release  
Red Hat Enterprise Linux Atomic Host 7.3
```

4. And next, you can check the Red Hat Gluster Storage Container Image, again from the host system (no need to log into the container image):

```
# docker run rhgs3/rhgs-server-rhel7 cat /etc/redhat-storage-release  
Red Hat Gluster Storage Server 3.1 Update 2 (Container)
```

This means that you have successfully downloaded a Red Hat Gluster Storage Image on a Red Hat Enterprise Linux 7.3 Atomic Host and you can use that image.

For information on configuring Containerized Red Hat Gluster Storage, see *Red Hat Gluster Storage Administration Guide*.

CHAPTER 7. UPDATING RED HAT GLUSTER STORAGE FROM 3.2.X TO 3.2.Y

A software update is a minor release that includes bug fixes for features, software enhancements etc. Red Hat strongly recommends you update your Red Hat Gluster Storage software regularly with the latest security patches and upgrades. Associate your system with a content server to update existing content or to install new content. This ensures that your system is up-to-date with security updates and upgrade.

To keep your Red Hat Gluster Storage system up-to-date, associate the system with the RHN or your locally-managed content service. This ensures your system automatically stays up-to-date with security patches and bug fixes.

Be aware of the following when updating your Red Hat Gluster Storage 3.2 installation:

- Updating from Red Hat Enterprise Linux 6 based Red Hat Gluster Storage to Red Hat Enterprise Linux 7 based Red Hat Gluster Storage is not supported.
- Asynchronous errata update releases of Red Hat Gluster Storage include all fixes that were released asynchronously since the last release as a cumulative update.
- When there are large number of snapshots, ensure to deactivate the snapshots before performing an update. The snapshots can be activated after the update is complete. For more information, see *Chapter 4.1 Starting and Stopping the glusterd service* in the *Red Hat Gluster Storage 3 Administration Guide*.

7.1. UPDATING RED HAT GLUSTER STORAGE IN THE OFFLINE MODE



WARNING

Before you update, be aware of changed requirements that exist after Red Hat Gluster Storage 3.1.3. If you want to access a volume being provided by a Red Hat Gluster Storage 3.1.3 or higher server, your client must also be using Red Hat Gluster Storage 3.1.3 or higher. Accessing volumes from other client versions can result in data becoming unavailable and problems with directory operations. This requirement exists because Red Hat Gluster Storage 3.1.3 contained a number of changes that affect how the Distributed Hash Table works in order to improve directory consistency and remove the effects seen in [BZ#1115367](#) and [BZ#1118762](#).

Before you update, be aware:

- Offline updates result in server downtime, as volumes are offline during upgrade.
- You must update all Red Hat Gluster Storage servers before updating any clients.
- This process assumes that you are updating to a thinly provisioned volume.

Updating Red Hat Gluster Storage 3.2 in the offline mode

1. Make a complete backup using a reliable backup solution. This Knowledge Base solution covers one possible approach: <https://access.redhat.com/solutions/1484053>.

If you use an alternative backup solution:

- Ensure that you have sufficient space available for a complete backup.
 - Copy the `.glusterfs` directory before copying any data files.
 - Ensure that no new files are created on Red Hat Gluster Storage file systems during the backup.
 - Ensure that all extended attributes, ACLs, owners, groups, and symbolic and hard links are backed up.
 - Check that the backup restores correctly before you continue with the migration.
2. Delete the existing Logical Volume (LV) and recreate a new thinly provisioned LV. For more information, see https://access.redhat.com/site/documentation/en-US/Red_Hat_Enterprise_Linux/6/html/Logical_Volume_Manager_Administration/thinprovisioning.html.
 3. Restore backed up content to the newly created thinly provisioned LV.
 4. When you are certain that your backup works, stop all volumes.

```
# gluster volume stop volname
```

5. Run the following commands to stop gluster services and update Red Hat Gluster Storage in the offline mode:

On Red Hat Enterprise Linux 6:

```
# service glusterd stop
# pkill glusterfs
# pkill glusterfsd
# yum update
```

On Red Hat Enterprise Linux 7:

```
# systemctl stop glusterd
# pkill glusterfs
# pkill glusterfsd
# yum update
```

Wait for the update to complete.

6. Start glusterd.

On Red Hat Enterprise Linux 6:

```
# service glusterd start
```

On Red Hat Enterprise Linux 7:

```
# systemctl start glusterd
```

7. When all nodes have been updated, run the following command to update the *op-version* of the cluster. This helps to prevent any compatibility issues within the cluster.

```
# gluster volume set all cluster.op-version 31001
```

**NOTE**

31001 is the *cluster.op-version* value for Red Hat Gluster Storage 3.2.0. Refer to [Section 1.5, “Supported Versions of Red Hat Gluster Storage”](#) for the correct *cluster.op-version* value for other versions.

8. Start your volumes with the following command:

```
# gluster volume start volname
```

9. If you had a meta volume configured prior to this upgrade, and you did not reboot as part of the upgrade process, mount the meta volume:

```
# mount /var/run/gluster/shared_storage/
```

If this command does not work, review the content of `/etc/fstab` and ensure that the entry for the shared storage is configured correctly and re-run the `mount` command. The line for the meta volume in the `/etc/fstab` file should look like the following:

```
hostname:/gluster_shared_storage /var/run/gluster/shared_storage/  
glusterfs defaults 0 0
```

10. If using NFS to access volumes, enable gluster-NFS using the following command:

```
# gluster volume set volname nfs.disable off
```

For example:

```
# gluster volume set testvol nfs.disable off  
volume set: success
```

11. If you use geo-replication, restart geo-replication sessions when upgrade is complete.

IMPORTANT

In Red Hat Gluster Storage 3.1 and higher, a meta volume is recommended when geo-replication is configured. However, when upgrading geo-replicated Red Hat Gluster Storage from version 3.0.x to 3.1.y, the older geo-replicated configuration that did not use shared volumes was persisted to the upgraded installation. Red Hat recommends reconfiguring geo-replication following upgrade to Red Hat Gluster Storage 3.2 to ensure that shared volumes are used and a meta volume is configured.

To enable shared volumes, set the `cluster.enable-shared-storage` parameter to `enable` from the master node:

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

Then configure geo-replication to use shared volumes as a meta volume by setting `use_meta_volume` to `true`.

```
# gluster volume geo-replication MASTER_VOL
SLAVE_HOST::SLAVE_VOL config use_meta_volume true
```

For further information see the following sections in the *Red Hat Gluster Storage 3.2 Administration Guide*:

- https://access.redhat.com/documentation/en-US/Red_Hat_Storage/3.2/html-single/Administration_Guide/index.html#chap-Managing_Red_Hat_Storage_Volumes-Shared_Volume
- https://access.redhat.com/documentation/en-US/Red_Hat_Storage/3.2/html-single/Administration_Guide/index.html#chap-Managing_Geo-replication-Meta_Volume

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

NOTE

As a result of [BZ#1347625](#), you may need to use the `force` parameter to successfully restart in some circumstances.

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

7.2. UPDATING THE NFS SERVER

Depending on the environment, the NFS server can be updated in the following ways:

- Updating Gluster NFS

- Updating NFS-Ganesha in the Offline Mode
- Migrating from Gluster NFS to NFS Ganesha in Offline mode

More detailed information about each is provided in the following sections.

7.2.1. Updating Gluster NFS

To update gluster NFS, refer [Section 7.1, “Updating Red Hat Gluster Storage in the Offline Mode”](#)

If you have a CTDB setup, then refer [Section 8.2.4.1, “In-Service Software Upgrade for a CTDB Setup”](#)

7.2.2. Updating NFS-Ganesha in the Offline Mode



NOTE

NFS-Ganesha does not support in-service updates. This means all running services and I/O operations must be stopped before starting the update process.

Execute the following steps to update the NFS-Ganesha service from Red Hat Gluster Storage 3.1.x to Red Hat Gluster Storage 3.2:

1. Back up all the volume export files under `/etc/ganesha/exports` and `ganesha.conf` under `/etc/ganesha`, in a backup directory on all the nodes:

From Red Hat Gluster Storage 3.1.x to Red Hat Gluster Storage 3.2

For example:

```
# cp /etc/ganesha/exports/export.v1.conf backup/
# cp /etc/ganesha/exports/export.v2.conf backup/
# cp /etc/ganesha/exports/export.v3.conf backup/
# cp /etc/ganesha/exports/export.v4.conf backup/
# cp /etc/ganesha/exports/export.v5.conf backup/
# cp /etc/ganesha/ganesha.conf backup/
# cp /etc/ganesha/ganesha-ha.conf backup/
```

From Red Hat Gluster Storage 3.2 to Red Hat Gluster Storage 3.2 Async

For example:

```
# cp /var/run/gluster/shared_storage/nfs-ganesha/exports/export.v1.conf backup/
# cp /var/run/gluster/shared_storage/nfs-ganesha/exports/export.v2.conf backup/
# cp /var/run/gluster/shared_storage/nfs-ganesha/exports/export.v3.conf backup/
# cp /var/run/gluster/shared_storage/nfs-ganesha/exports/export.v4.conf backup/
# cp /var/run/gluster/shared_storage/nfs-ganesha/exports/export.v5.conf backup/
# cp /var/run/gluster/shared_storage/nfs-ganesha/ganesha.conf backup/
# cp /var/run/gluster/shared_storage/nfs-ganesha/ganesha-ha.conf backup/
```

-
2. Disable nfs-ganesha on the cluster by executing the following command:

```
# gluster nfs-ganesha disable
```

For example:

```
# gluster nfs-ganesha disable
This will take a few minutes to complete. Please wait ..
nfs-ganesha : success
```

3. Disable the shared volume in cluster by executing the following command:

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

For example:

```
# gluster volume set all cluster.enable-shared-storage disable
Disabling cluster.enable-shared-storage will delete the shared
storage volumec(gluster_shared_storage), which is used by snapshot
scheduler, geo-replication and NFS-Ganesha.
Do you still want to continue? (y/n) y
volume set: success
```

4. Stop the glusterd service and kill any running gluster process on all the nodes.

On Red Hat Enterprise Linux 7:

```
# systemctl stop glusterd
# pkill glusterfs
# pkill glusterfsd
```

On Red Hat Enterprise Linux 6:

```
# service glusterd stop
# pkill glusterfs
# pkill glusterfsd
```

5. Ensure all gluster processes are stopped and if there are any gluster processes still running, terminate the process using kill, on all the nodes by executing the following command:

```
# pgrep gluster
```

6. Stop the pcsd service on all nodes of the cluster.

On Red Hat Enterprise Linux 7:

```
# systemctl stop pcsd
```

On Red Hat Enterprise Linux 6:

```
# service pcsd stop
```

7. Update the packages on all the nodes by executing the following command:

```
# yum update
```

This updates the required packages and any dependencies of those packages.



IMPORTANT

- From Red Hat Gluster Storage 3.2, NFS-Ganesha packages must be installed on all the nodes of the trusted storage pool.
- Verify on all the nodes that the required packages are updated, the nodes are fully functional and are using the correct versions. If anything does not seem correct, then do not proceed until the situation is resolved. Contact the Red Hat Global Support Services for assistance if needed.

8. Start the `glusterd` and `pcsd` service on all the nodes by executing the following commands.

On Red Hat Enterprise Linux 7:

```
# systemctl start glusterd
# systemctl start pcsd
```

On Red Hat Enterprise Linux 6:

```
# service glusterd start
# service pcsd start
```

9. When all nodes have been upgraded, run the following command to update the `op-version` of the cluster. This helps to prevent any compatibility issues within the cluster.

```
# gluster volume set all cluster.op-version 31001
```

10. 1. Copy the volume's export information from backup copy of `ganesha.conf` to the newly renamed `ganesha.conf` under `/etc/ganesha`.

Export entries will look like as below in backup copy of `ganesha.conf` :

```
%include "/etc/ganesha/exports/export.v1.conf"
%include "/etc/ganesha/exports/export.v2.conf"
%include "/etc/ganesha/exports/export.v3.conf"
%include "/etc/ganesha/exports/export.v4.conf"
%include "/etc/ganesha/exports/export.v5.conf"
```

2. Copy the backup volume export files from backup directory to `/etc/ganesha/exports`

```
# cp export.* /etc/ganesha/exports/
```

11. Enable the firewall settings for the new services and ports. Information on how to enable the services is available in the *Red Hat Gluster Storage Administration Guide*

12. Enable the shared volume in the cluster:

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

For example:

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

13. Ensure that the shared storage volume mount exists on the server after node reboot/shutdown. If it does not, then mount the shared storage volume manually using the following command:

```
# mount -t glusterfs <local_node's_hostname>:gluster_shared_storage
/var/run/gluster/shared_storage
```

14. Once the shared volume is created, create a folder named “nfs-ganesha” inside `/var/run/gluster/shared_storage`:

```
# cd /var/run/gluster/shared_storage/
# mkdir nfs-ganesha
```

15. Copy the `ganesha.conf`, `ganesha-ha.conf`, and the `exports` folder from `/etc/ganesha` to `/var/run/gluster/shared_storage/nfs-ganesha`

```
# cd /etc/ganesha/
# cp ganesha.conf ganesha-ha.conf
/var/run/gluster/shared_storage/nfs-ganesha/
# cp -r exports/ /var/run/gluster/shared_storage/nfs-ganesha/
```

16. If there are any export entries in the `ganesha.conf` file, then update the path in the file using the following command:

```
# sed -i 's/\/etc\/ganesha/\/var\/run\/gluster\/shared_storage\/nfs-ganesha/' /var/run/gluster/shared_storage/nfs-ganesha/ganesha.conf
```

17. Execute the following command to cleanup any already existing cluster related configuration:

```
/usr/libexec/ganesha/ganesha-ha.sh --cleanup
/var/run/gluster/shared_storage/nfs-ganesha
```

18. If you have upgraded to Red Hat Enterprise Linux 7.4, enable the `ganesha_use_fusefs` and the `gluster_use_execmem` boolean before enabling NFS-Ganesha by executing the following commands:

```
# setsebool -P ganesha_use_fusefs on
# setsebool -P gluster_use_execmem on
```

19. Enable `nfs-ganesha` on the cluster:

```
# gluster nfs-ganesha enable
```

For example:

```
# gluster nfs-ganesha enable
Enabling NFS-Ganesha requires Gluster-NFS to be disabled across the
trusted pool. Do you still want to continue?
(y/n) y
This will take a few minutes to complete. Please wait ..
nfs-ganesha : success
```



IMPORTANT

Verify that all the nodes are functional. If anything does not seem correct, then do not proceed until the situation is resolved. Contact Red Hat Global Support Services for assistance if required.

7.2.3. Migrating from Gluster NFS to NFS Ganesha in Offline mode

The following steps have to be performed on each node of the replica pair to migrate from Gluster NFS to NFS Ganesha

1. To ensure that CTDB does not start automatically after a reboot run the following command on each node of the CTDB cluster:

```
# chkconfig ctdb off
```

2. Stop the CTDB service on the Red Hat Gluster Storage node using the following command on each node of the CTDB cluster:

```
# service ctdb stop
```

3. To verify if the CTDB and NFS services are stopped, execute the following command:

```
ps axf | grep -E '(ctdb|nfs)[d]'
```

4. Stop the gluster services on the storage server using the following commands:

```
# service glusterd stop
# pkill glusterfs
# pkill glusterfsd
```

5. Delete the CTDB volume by executing the following command:

```
# gluster vol delete <ctdb_vol_name>
```

6. Update the server using the following command:

```
# yum update
```

7. Reboot the server

8. Start the glusterd service using the following command:

```
# service glusterd start
```


On Red Hat Enterprise Linux 7, execute the following command:

```
# systemctl start glusterd
```

9. When all nodes have been upgraded, run the following command to update the op-version of the cluster. This helps to prevent any compatibility issues within the cluster.

```
# gluster volume set all cluster.op-version 31001
```

10. To install nfs-ganesha packages, refer [Chapter 4, Deploying NFS-Ganesha on Red Hat Gluster Storage](#)
11. To configure nfs-ganesha cluster, refer section *NFS-Ganesha* in the *Red Hat Gluster Storage Administration Guide*.

7.3. IN-SERVICE SOFTWARE UPDATE FROM RED HAT GLUSTER STORAGE



WARNING

Before you update, be aware of changed requirements that exist after Red Hat Gluster Storage 3.1.3. If you want to access a volume being provided by a Red Hat Gluster Storage 3.1.3 or higher server, your client must also be using Red Hat Gluster Storage 3.1.3 or higher. Accessing volumes from other client versions can result in data becoming unavailable and problems with directory operations. This requirement exists because Red Hat Gluster Storage 3.1.3 contained a number of changes that affect how the Distributed Hash Table works in order to improve directory consistency and remove the effects seen in [BZ#1115367](#) and [BZ#1118762](#).



IMPORTANT

In Red Hat Enterprise Linux 7 based Red Hat Gluster Storage, updating to 3.1 or higher reloads firewall rules. All runtime-only changes made before the reload are lost.

Before you update, be aware:

- You must update all Red Hat Gluster Storage servers before updating any clients.
- If geo-replication is in use, slave nodes must be updated before master nodes.
- NFS-Ganesha does not support in-service updates. All the running services and I/O operations must be stopped before starting the update process. For more information see, [Section 7.2.2, “Updating NFS-Ganesha in the Offline Mode”](#).
- Dispersed volumes (volumes that use erasure coding) do not support in-service updates and cannot be updated in a non-disruptive manner.

- The SMB and CTDB services do not support in-service updates. The procedure outlined in this section involves service interruptions to the SMB and CTDB services.
- If updating Samba, ensure that Samba is upgraded on all nodes simultaneously, as running different versions of Samba in the same cluster results in data corruption.
- Your system must be registered to Red Hat Network. For more information refer to [Section 2.6, “Subscribing to the Red Hat Gluster Storage Server Channels”](#)
- Do not perform any volume operations while the cluster is being updated.

To update your system to Red Hat Gluster Storage 3.2.x, follow these steps. The following steps must be performed on each node of the replica pair.

Updating Red Hat Gluster Storage 3.2 in in-service mode

1. If you have a replicated configuration, perform these steps on all nodes of a replica set.

If you have a distributed-replicated setup, perform these steps on one replica set at a time, for all replica sets.

2. Stop any geo-replication sessions.

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

3. Verify that there are no pending self-heals:

```
# gluster volume heal volname info
```

Wait for any self-heal operations to complete before continuing.

4. Stop the gluster services on the storage server using the following commands:

```
# service glusterd stop
# pkill glusterfs
# pkill glusterfsd
```

5. If you use Samba:

1. Enable the required repository.

On Red Hat Enterprise Linux 6.7 or later:

```
# subscription-manager repos --enable=rh-gluster-3-samba-for-
rhel-6-server-rpms
```

On Red Hat Enterprise Linux 7:

```
# subscription-manager repos --enable=rh-gluster-3-samba-for-
rhel-7-server-rpms
```

2. Stop the CTDB and SMB services across all nodes in the Samba cluster using the following command. Stopping the CTDB service also stops the SMB service.

```
# service ctdb stop
```

This ensures different versions of Samba do not run in the same Samba cluster.

3. Verify that the CTDB and SMB services are stopped by running the following command:

```
ps axf | grep -E '(ctdb|smb|winbind|nmb)[d]'
```

6. Update the server using the following command:

```
# yum update
```

Wait for the update to complete.

7. If a kernel update was included as part of the update process in the previous step, reboot the server.
8. If a reboot of the server was not required, then start the gluster services on the storage server using the following command:

```
# service glusterd start
```

Additionally, if you use Samba:

1. Mount `/gluster/lock` before starting CTDB by executing the following command:

```
# mount -a
```

2. If the CTDB and SMB services were stopped earlier, then start the services by executing the following command.

```
# service ctdb start
```

3. To verify if the CTDB and SMB services have started, execute the following command:

```
ps axf | grep -E '(ctdb|smb|winbind|nmb)[d]'
```

9. To verify that you have updated to the latest version of the Red Hat Gluster Storage server execute the following command and compare output with the desired version in [Section 1.5, “Supported Versions of Red Hat Gluster Storage”](#).

```
# gluster --version
```

10. Ensure that all bricks are online. To check the status, execute the following command:

```
# gluster volume status
```

11. Start self-heal on the volume.


```
# gluster volume heal volname
```

12. Ensure self-heal is complete on the replica using the following command:

```
# gluster volume heal volname info
```

13. When all nodes in the volume have been upgraded, run the following command to update the *op-version* of the cluster. This helps to prevent any compatibility issues within the cluster.

```
# gluster volume set all cluster.op-version 31001
```



NOTE

31001 is the *cluster.op-version* value for Red Hat Gluster Storage 3.2.0. Refer to [Section 1.5, “Supported Versions of Red Hat Gluster Storage”](#) for the correct *cluster.op-version* value for other versions.

14. If you had a meta volume configured prior to this upgrade, and you did not reboot as part of the upgrade process, mount the meta volume:

```
# mount /var/run/gluster/shared_storage/
```

If this command does not work, review the content of `/etc/fstab` and ensure that the entry for the shared storage is configured correctly and re-run the `mount` command. The line for the meta volume in the `/etc/fstab` file should look like the following:

```
hostname:/gluster_shared_storage /var/run/gluster/shared_storage/  
glusterfs defaults 0 0
```

15. If you use geo-replication, restart geo-replication sessions when upgrade is complete.

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



NOTE

As a result of [BZ#1347625](#), you may need to use the *force* parameter to successfully restart in some circumstances.

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

CHAPTER 8. UPGRADING TO RED HAT GLUSTER STORAGE 3.2

This chapter describes the procedure to upgrade to Red Hat Gluster Storage 3.2 from Red Hat Gluster Storage 3.1.

Upgrade support limitations

- Upgrading from Red Hat Enterprise Linux 6 based Red Hat Gluster Storage to Red Hat Enterprise Linux 7 based Red Hat Gluster Storage is not supported.
- Servers must be upgraded prior to upgrading clients.

8.1. OFFLINE UPGRADE FROM RED HAT GLUSTER STORAGE 3.1 TO 3.2



WARNING

Before you upgrade, be aware of changed requirements that exist after Red Hat Gluster Storage 3.1.3. If you want to access a volume being provided by a Red Hat Gluster Storage 3.1.3 or higher server, your client must also be using Red Hat Gluster Storage 3.1.3 or higher. Accessing volumes from other client versions can result in data becoming unavailable and problems with directory operations. This requirement exists because Red Hat Gluster Storage 3.1.3 contained a number of changes that affect how the Distributed Hash Table works in order to improve directory consistency and remove the effects seen in [BZ#1115367](#) and [BZ#1118762](#).



IMPORTANT

In Red Hat Enterprise Linux 7 based Red Hat Gluster Storage, updating to 3.1 or higher reloads firewall rules. All runtime-only changes made before the reload are lost.

8.1.1. Upgrading from Red Hat Gluster Storage 3.1 to Red Hat Gluster Storage 3.2 for Systems Subscribed to Red Hat Network

Procedure 8.1. Before you upgrade

1. Unmount gluster volumes from all clients. On a client, use the following command to unmount a volume from a mount point.

```
# umount mount-point
```

2. On a gluster server, stop all volumes by running the following command for each volume.

```
# gluster volume stop volname
```

3. Unmount the data partition(s) from the servers using the following command.

```
# umount mount-point
```

4. Verify that all volumes are stopped.

```
# gluster volume info
```

5. Stop the `glusterd` services on all servers using the following command:

```
# service glusterd stop
```

Procedure 8.2. Upgrade using yum

1. Execute the following command to kill all gluster processes:

```
# pkill gluster
```

2. Verify that your system is not on the legacy Red Hat Network Classic update system.

```
# migrate-rhs-classic-to-rhsm --status
```

If you are still on Red Hat Network Classic, run the following command to migrate to Red Hat Subscription Manager.

```
# migrate-rhs-classic-to-rhsm --rhn-to-rhsm
```

Then verify that your status has changed.

```
# migrate-rhs-classic-to-rhsm --status
```

3. If you use Samba:

1. For Red Hat Enterprise Linux 6.7 or higher, enable the following repository:

```
# subscription-manager repos --enable=rh-gluster-3-samba-for-rhel-6-server-rpms
```

For Red Hat Enterprise Linux 7, enable the following repository:

```
# subscription-manager repos --enable=rh-gluster-3-samba-for-rhel-7-server-rpms
```

2. Ensure that Samba is upgraded on all the nodes simultaneously, as running different versions of Samba in the same cluster will lead to data corruption.

Stop the CTDB and SMB services and verify that they are stopped.

```
# service ctdb stop
```

```
# ps axf | grep -E '(ctdb|smb|winbind|nmb)[d]'
```

- Upgrade the server to Red Hat Gluster Storage 3.2.

```
# yum update
```

Wait for the update to complete.

- Reboot the server to ensure that kernel updates are applied.
- When all nodes have been upgraded, run the following command to update the *op-version* of the cluster. This helps to prevent any compatibility issues within the cluster.

```
# gluster volume set all cluster.op-version 31001
```



NOTE

31001 is the *cluster.op-version* value for Red Hat Gluster Storage 3.2.0. Refer to [Section 1.5, “Supported Versions of Red Hat Gluster Storage”](#) for the correct *cluster.op-version* value for other versions.

8.1.2. Upgrading from Red Hat Gluster Storage 3.1 to Red Hat Gluster Storage 3.2 for Systems Subscribed to Red Hat Satellite Server

- Unmount all volumes. On a client, use the following command to unmount a volume from the mount point.

```
# umount mount-name
```

- On a gluster server, stop all volumes using the following command:

```
# gluster volume stop volname
```

- Unmount the data partition(s) on the servers using the following command:

```
umount mount-point
```

- Create an Activation Key at the Red Hat Satellite Server, and associate it with the following channels. For more information, refer to [Section 2.5, “Installing from Red Hat Satellite Server”](#)
 - For Red Hat Enterprise Linux 6.7 or higher:

```
Base Channel: Red Hat Enterprise Linux Server (v.6 for 64-bit x86_64)
```

```
Child channels:
```

```
RHEL Server Scalable File System (v. 6 for x86_64)
```

```
Red Hat Gluster Storage Server 3 (RHEL 6 for x86_64)
```

If you use Samba, add the following channel:

```
Red Hat Gluster 3 Samba (RHEL 6 for x86_64)
```

- o For Red Hat Enterprise Linux 7:

```
Base Channel: Red Hat Enterprise Linux Server (v.7 for 64-bit
x86_64)
```

Child channels:

```
RHEL Server Scalable File System (v. 7 for x86_64)
Red Hat Gluster Storage Server 3 (RHEL 7 for x86_64)
```

If you use Samba, add the following channel:

```
Red Hat Gluster 3 Samba (RHEL 6 for x86_64)
```

5. Unregister your system from Red Hat Satellite by following these steps:

1. Log in to the Red Hat Satellite server.
 2. Click on the **Systems** tab in the top navigation bar and then the name of the old or duplicated system in the **System List**.
 3. Click the **delete system** link in the top-right corner of the page.
 4. To confirm the system profile deletion by clicking the **Delete System** button.
6. Run the following command on your Red Hat Gluster Storage server, using your credentials and the Activation Key you prepared earlier. This re-registers the system to the Red Hat Gluster Storage 3.2 channels on the Red Hat Satellite Server.

```
# rhnreg_ks --username username --password password --force --
activationkey Activation Key ID
```

7. Verify that the channel subscriptions have been updated.

On Red Hat Enterprise Linux 6.7 and higher, look for the following channels, as well as the **rhel-x86_64-server-6-rh-gluster-3-samba** channel if you use Samba.

```
# rhn-channel --list
rhel-x86_64-server-6
rhel-x86_64-server-6-rhs-3
rhel-x86_64-server-sfs-6
```

On Red Hat Enterprise Linux 7, look for the following channels, as well as the **rhel-x86_64-server-7-rh-gluster-3-samba** channel if you use Samba.

```
# rhn-channel --list
rhel-x86_64-server-7
rhel-x86_64-server-7-rhs-3
rhel-x86_64-server-sfs-7
```

8. Upgrade to Red Hat Gluster Storage 3.2.

```
# yum update
```


9. Reboot the server and run volume and data integrity checks.
10. When all nodes have been upgraded, run the following command to update the *op-version* of the cluster. This helps to prevent any compatibility issues within the cluster.

```
# gluster volume set all cluster.op-version 31001
```



NOTE

31001 is the *cluster.op-version* value for Red Hat Gluster Storage 3.2.0. Refer to [Section 1.5, “Supported Versions of Red Hat Gluster Storage”](#) for the correct *cluster.op-version* value for other versions.

8.2. IN-SERVICE SOFTWARE UPGRADE FROM RED HAT GLUSTER STORAGE 3.1 TO RED HAT GLUSTER STORAGE 3.2



IMPORTANT

Update all Red Hat Gluster Storage servers before updating clients.

In-service software upgrade refers to the ability to progressively update a Red Hat Gluster Storage Server cluster with a new version of the software without taking the volumes hosted on the cluster offline. In most cases normal I/O operations on the volume can continue even when the cluster is being updated under most circumstances. This method of updating the storage cluster is only supported for replicated and distributed-replicated volumes.

8.2.1. Pre-upgrade Tasks

Ensure you perform the following steps based on the set-up before proceeding with the in-service software upgrade process.

8.2.1.1. Upgrade Requirements for Red Hat Gluster Storage 3.2

The following are the upgrade requirements to upgrade to Red Hat Gluster Storage 3.2 from the preceding update:

- In-service software upgrade is supported only for nodes with replicate and distributed replicate volumes.
- If you want to use snapshots for your existing environment, each brick must be an independent thinly provisioned logical volume (LV). If you do not plan to use snapshots, thickly provisioned volumes remain supported.
- A Logical Volume that contains a brick must not be used for any other purpose.
- Only linear LVM is supported with Red Hat Gluster Storage 3.2. For more information, see https://access.redhat.com/site/documentation/en-US/Red_Hat_Enterprise_Linux/4/html-single/Cluster_Logical_Volume_Manager/#lv_overview
- When server-side quorum is enabled, ensure that bringing one node down does not violate server-side quorum. Add dummy peers to ensure the server-side quorum will not be violated until the completion of rolling upgrade using the following command:

```
# gluster peer probe dummynode
```



NOTE

If you have a geo-replication session, then to add a node follow the steps mentioned in the section *Starting Geo-replication for a New Brick or New Node* in the *Red Hat Gluster Storage Administration Guide*

For example, when the server-side quorum percentage is set to the default value (>50%), for a plain replicate volume with two nodes and one brick on each machine, a dummy node that does not contain any bricks must be added to the trusted storage pool to provide high availability of the volume using the command mentioned above.

In a three node cluster, if the server-side quorum percentage is set to 77%, bringing down one node would violate the server-side quorum. In this scenario, you have to add two dummy nodes to meet server-side quorum.

- For replica 2 volumes, disable client-side quorum. This is not recommended for replica 3 volumes, as it increases the risk of split brain conditions developing.

```
# gluster volume reset <vol-name> cluster.quorum-type
```

- Stop any geo-replication sessions running between the master and slave.

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

- Ensure that there are no pending self-heals before proceeding with in-service software upgrade using the following command:

```
# gluster volume heal volname info
```

- Ensure the Red Hat Gluster Storage server is registered to the required channels.

On Red Hat Enterprise Linux 6:

```
rhel-x86_64-server-6
rhel-x86_64-server-6-rhs-3
rhel-x86_64-server-sfs-6
```

On Red Hat Enterprise Linux 7:

```
rhel-x86_64-server-7
rhel-x86_64-server-7-rhs-3
rhel-x86_64-server-sfs-7
```

To subscribe to the channels, run the following command:

```
# rhn-channel --add --channel=<channel>
```

8.2.1.2. Restrictions for In-Service Software Upgrade

The following lists some of the restrictions for in-service software upgrade:

- Do not perform in-service software upgrade when the I/O or load is high on the Red Hat Gluster Storage server.
- Do not perform any volume operations on the Red Hat Gluster Storage server.
- Do not change hardware configurations.
- Do not run mixed versions of Red Hat Gluster Storage for an extended period of time. For example, do not have a mixed environment of Red Hat Gluster Storage 3.1 and Red Hat Gluster Storage 3.2 for a prolonged time.
- Do not combine different upgrade methods.
- It is not recommended to use in-service software upgrade for migrating to thinly provisioned volumes, but to use offline upgrade scenario instead. For more information see, [Section 8.1, “Offline Upgrade from Red Hat Gluster Storage 3.1 to 3.2”](#)

8.2.1.3. Configuring repo for Upgrading using ISO

To configure the repo to upgrade using ISO, execute the following steps:



NOTE

Upgrading Red Hat Gluster Storage using ISO can be performed only from the immediately preceding release. This means that upgrading to Red Hat Gluster Storage 3.2 using ISO can only be done from Red Hat Gluster Storage 3.1 Update 3. For a complete list of supported Red Hat Gluster Storage releases, see [Section 1.5, “Supported Versions of Red Hat Gluster Storage”](#).

1. Mount the ISO image file under any directory using the following command:

```
mount -o loop <ISO image file> <mount-point>
```

For example:

```
mount -o loop RHGS-3.2-20170209.0-RHS-x86_64-dvd1.iso /mnt
```

2. Set the repo options in a file in the following location:

```
/etc/yum.repos.d/<file_name.repo>
```

3. Add the following information to the repo file:

```
[local]
name=local
baseurl=file:///mnt
enabled=1
gpgcheck=0
```

8.2.1.4. Preparing and Monitoring the Upgrade Activity

Before proceeding with the in-service software upgrade, prepare and monitor the following processes:

- Check the peer and volume status to ensure that all peers are connected and there are no active volume tasks.

```
# gluster peer status
```

```
# gluster volume status
```

- Check the rebalance status using the following command:

```
# gluster volume rebalance r2 status
Node   Rebalanced-files size      scanned  failures  skipped
status run time in secs
-----
10.70.43.198      0      0Bytes    99       0         0
completed      1.00
10.70.43.148     49     196Bytes  100      0         0
completed      3.00
```

- Ensure that there are no pending self-heals before proceeding with in-service software upgrade using the following command:

```
# gluster volume heal volname info
```

The following example shows no pending self-heals.

```
# gluster volume heal drvol info
Gathering list of entries to be healed on volume drvol has been
successful

Brick 10.70.37.51:/rhs/brick1/dir1
Number of entries: 0

Brick 10.70.37.78:/rhs/brick1/dir1
Number of entries: 0

Brick 10.70.37.51:/rhs/brick2/dir2
Number of entries: 0

Brick 10.70.37.78:/rhs/brick2/dir2
Number of entries: 0
```

8.2.2. Service Impact of In-Service Upgrade

In-service software upgrade will impact the following services. Ensure you take the required precautionary measures.

SWIFT

ReST requests that are in transition will fail during in-service software upgrade. Hence it is recommended to stop all swift services before in-service software upgrade using the following commands:

```
# service openstack-swift-proxy stop
# service openstack-swift-account stop
# service openstack-swift-container stop
# service openstack-swift-object stop
```

NFS

When you NFS mount a volume, any new or outstanding file operations on that file system will hang uninterruptedly during in-service software upgrade until the server is upgraded.

Samba / CTDB

Ongoing I/O on Samba shares will fail as the Samba shares will be temporarily unavailable during the in-service software upgrade, hence it is recommended to stop the Samba service using the following command:

```
# service ctdb stop ;Stopping CTDB will also stop the SMB service.
```

Distribute Volume

In-service software upgrade is not supported for distributed volume. If you have a distributed volume in the cluster, stop that volume for the duration of the upgrade.

```
# gluster volume stop <VOLNAME>
```

Virtual Machine Store

The virtual machine images are likely to be modified constantly. The virtual machine listed in the output of the volume heal command does not imply that the self-heal of the virtual machine is incomplete. It could mean that the modifications on the virtual machine are happening constantly.

Hence, if you are using a gluster volume for storing virtual machine images (Red Hat Enterprise Linux, Red Hat Enterprise Virtualization and Red Hat OpenStack), then it is recommended to power-off all virtual machine instances before in-service software upgrade.

8.2.3. In-Service Software Upgrade

The following steps have to be performed on each node of the replica pair:

1. Back up the following configuration directory and files in a location that is not on the operating system partition.

```
/var/lib/glusterd, /etc/swift, /etc/samba, /etc/ctdb, /etc/glusterfs,
/var/lib/samba, /var/lib/ctdb
```

```
# cp -a /var/lib/glusterd /backup-disk/
# cp -a /etc/swift /backup-disk/
# cp -a /etc/samba /backup-disk/
# cp -a /etc/ctdb /backup-disk/
# cp -a /etc/glusterfs /backup-disk/
# cp -a /var/lib/samba /backup-disk/
# cp -a /var/lib/ctdb /backup-disk/
```

**NOTE**

- If you have a CTDB environment, see [Section 8.2.4.1, “In-Service Software Upgrade for a CTDB Setup”](#).

2. Ensure that there are no pending self-heals before proceeding with in-service software upgrade using the following command:

```
# gluster volume heal volname info
```

3. Stop the gluster services on the storage server using the following commands:

```
# service glusterd stop
# pkill glusterfs
# pkill glusterfsd
```

4. Verify that your system is not using the legacy Red Hat Classic update software.

```
# migrate-rhs-classic-to-rhsm --status
```

If your system uses this legacy software, migrate to Red Hat Subscription Manager and verify that your status has changed when migration is complete.

```
# migrate-rhs-classic-to-rhsm --rhn-to-rhsm
```

```
# migrate-rhs-classic-to-rhsm --status
```

5. Update the server using the following command:

```
# yum update
```

6. If the volumes are thickly provisioned, and you plan to use snapshots, perform the following steps to migrate to thinly provisioned volumes:

**NOTE**

Migrating from thickly provisioned volume to thinly provisioned volume during in-service software upgrade takes a significant amount of time based on the data you have in the bricks. If you do not plan to use snapshots, you can skip this step. However, if you plan to use snapshots on your existing environment, the offline method to upgrade is recommended. For more information regarding offline upgrade, see [Section 8.1, “Offline Upgrade from Red Hat Gluster Storage 3.1 to 3.2”](#)

Contact a Red Hat Support representative before migrating from thickly provisioned volumes to thinly provisioned volumes using in-service software upgrade.

1. Unmount all the bricks associated with the volume by executing the following command:

```
# umount mount_point
```

2. Remove the LVM associated with the brick by executing the following command:

```
# lvremove logical_volume_name
```

For example:

```
# lvremove /dev/RHS_vg/brick1
```

3. Remove the volume group by executing the following command:

```
# vgrename -ff volume_group_name
```

For example:

```
vgremove -ff RHS_vg
```

4. Remove the physical volume by executing the following command:

```
# pvremove -ff physical_volume
```

5. If the physical volume (PV) not created then create the PV for a RAID 6 volume by executing the following command, else proceed with the next step:

```
# pvcreate --dataalignment 2560K /dev/vdb
```

For more information, see the *Red Hat Gluster Storage 3.2 Administration Guide* https://access.redhat.com/documentation/en-us/red_hat_gluster_storage/3.2/html-single/administration_guide/#Formatting_and_Mounting_Bricks.

6. Create a single volume group from the PV by executing the following command:

```
# vgcreate volume_group_name disk
```

For example:

```
vgcreate RHS_vg /dev/vdb
```

7. Create a thinpool using the following command:

```
# lvcreate -L size --poolmetadatasize md size --chunksize chunk size -T pool device
```

For example:

```
lvcreate -L 2T --poolmetadatasize 16G --chunksize 256 -T /dev/RHS_vg/thin_pool
```

8. Create a thin volume from the pool by executing the following command:

```
# lvcreate -V size -T pool device -n thinp
```

For example:

```
lvcreate -V 1.5T -T /dev/RHS_vg/thin_pool -n thin_vol
```

9. Create filesystem in the new volume by executing the following command:

```
mkfs.xfs -i size=512 thin_pool device
```

For example:

```
mkfs.xfs -i size=512 /dev/RHS_vg/thin_vol
```

The back-end is now converted to a thinly provisioned volume.

10. Mount the thinly provisioned volume to the brick directory and setup the extended attributes on the bricks. For example:

```
# setfattr -n trusted.glusterfs.volume-id \ -v 0x$(grep volume-id  
/var/lib/glusterd/vols/volname/info \ | cut -d= -f2 | sed 's/-  
//g') $brick
```

7. To ensure Red Hat Gluster Storage Server node is healthy after reboot and so that it can then be joined back to the cluster, it is recommended that you disable glusterd during boot using the following command:

```
# chkconfig glusterd off
```

8. Reboot the server.

9. Perform the following operations to change the Automatic File Replication extended attributes so that the heal process happens from a brick in the replica subvolume to the thin provisioned brick.

1. Create a FUSE mount point from any server to edit the extended attributes. Using the NFS and CIFS mount points, you will not be able to edit the extended attributes.

Note that /mnt/r2 is the FUSE mount path.

2. Create a new directory on the mount point and ensure that a directory with such a name is not already present.

```
# mkdir /mnt/r2/name-of-nonexistent-dir
```

3. Delete the directory and set the extended attributes.

```
# rmdir /mnt/r2/name-of-nonexistent-dir
```

```
# setfattr -n trusted.non-existent-key -v abc /mnt/r2  
# setfattr -x trusted.non-existent-key /mnt/r2
```


4. Ensure that the extended attributes of the brick in the replica subvolume(In this example, **brick: /dev/RHS_vg/brick2** , extended attribute: **trusted.afr.r2-client-1**), is not set to zero.

```
# getfattr -d -m. -e hex /dev/RHS_vg/brick2 # file:
/dev/RHS_vg/brick2
security.selinux=0x756e636f6e66696e65645f753a6f626a6563745f723a66
696c655f743a733000
trusted.afr.r2-client-0=0x000000000000000000000000
trusted.afr.r2-client-1=0x0000000000000000300000002
trusted.gfid=0x00000000000000000000000000000001
trusted.glusterfs.dht=0x0000000100000000000000007fffffff
trusted.glusterfs.volume-id=0xde822e25ebd049ea83bfaa3c4be2b440
```

10. Start the **glusterd** service using the following command:

```
# service glusterd start
```

11. To automatically start the **glusterd** daemon every time the system boots, run the following command:

```
# chkconfig glusterd on
```

12. To verify if you have upgraded to the latest version of the Red Hat Gluster Storage server execute the following command:

```
# gluster --version
```

13. Ensure that all the bricks are online. To check the status execute the following command:

```
# gluster volume status
```

For example:

```
# gluster volume status
Status of volume: r2

Gluster process                                     Port
Online  Pid
-----
-----
Brick 10.70.43.198:/brick/r2_0                        49152  Y
32259
Brick 10.70.42.237:/brick/r2_1                        49152  Y
25266
Brick 10.70.43.148:/brick/r2_2                        49154  Y
2857
Brick 10.70.43.198:/brick/r2_3                        49153  Y
32270
NFS Server on localhost                             2049   Y
25280
Self-heal Daemon on localhost                       N/A    Y
25284
```

```

NFS Server on 10.70.43.148                2049    Y
2871
Self-heal Daemon on 10.70.43.148         N/A     Y
2875
NFS Server on 10.70.43.198                2049    Y
32284
Self-heal Daemon on 10.70.43.198         N/A     Y
32288

```

```
Task Status of Volume r2
```

```
-----
```

```
-----
There are no active volume tasks
```

14. Start self-heal on the volume.

```
# gluster volume heal vol-name
```

15. Ensure self-heal is complete on the replica using the following command:

```
# gluster volume heal volname info
```

The following example shows self heal completion:

```

# gluster volume heal drvol info
Gathering list of entries to be healed on volume drvol has been
successful

Brick 10.70.37.51:/rhs/brick1/dir1
Number of entries: 0

Brick 10.70.37.78:/rhs/brick1/dir1
Number of entries: 0

Brick 10.70.37.51:/rhs/brick2/dir2
Number of entries: 0

Brick 10.70.37.78:/rhs/brick2/dir2
Number of entries: 0

```

16. Repeat the above steps on the other node of the replica pair.



NOTE

In the case of a distributed-replicate setup, repeat the above steps on all the replica pairs.

17. When all nodes have been upgraded, run the following command to update the *op-version* of the cluster. This helps to prevent any compatibility issues within the cluster.

```
# gluster volume set all cluster.op-version 31001
```

**NOTE**

31001 is the *cluster.op-version* value for Red Hat Gluster Storage 3.2.0. Refer to [Section 1.5, “Supported Versions of Red Hat Gluster Storage”](#) for the correct *cluster.op-version* value for other versions.

**NOTE**

If you want to enable snapshots, see the *Red Hat Gluster Storage 3.2 Administration Guide*: https://access.redhat.com/documentation/en-us/red_hat_gluster_storage/3.2/html-single/administration_guide/#Troubleshooting1.

18. If the client-side quorum was disabled before upgrade, then upgrade it by executing the following command:

```
# gluster volume set volname cluster.quorum-type auto
```

19. If a dummy node was created earlier, then detach it by executing the following command:

```
# gluster peer detach <dummy_node name>
```

20. If the geo-replication session between master and slave was disabled before upgrade, then configure the meta volume and restart the session:

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

```
# gluster volume geo-replication Volume1 example.com::slave-vol
config use_meta_volume true
```

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

8.2.4. Special Consideration for In-Service Software Upgrade

The following sections describe the in-service software upgrade steps for a CTDB setup.

8.2.4.1. In-Service Software Upgrade for a CTDB Setup

Before you upgrade the CTDB packages, ensure you upgrade the Red Hat Gluster Storage server by following these steps. The following steps have to be performed on each node of the replica pair.

1. To ensure that the CTDB does not start automatically after a reboot run the following command on each node of the CTDB cluster:

```
# chkconfig ctdb off
```

2. Stop the CTDB service on the Red Hat Gluster Storage node using the following command on each node of the CTDB cluster:

```
# service ctdb stop
```

1. To verify if the CTDB and SMB services are stopped, execute the following command:

```
ps axf | grep -E '(ctdb|smb|winbind|nmb)[d]'
```

3. Stop the gluster services on the storage server using the following commands:

```
# service glusterd stop  
# pkill glusterfs  
# pkill glusterfsd
```

4. In `/etc/fstab`, comment out the line containing the volume used for CTDB service as shown in the following example:

```
# HostName:/volname /gluster/lock glusterfs defaults,transport=tcp  
0 0
```

5. Update the server using the following command:

```
# yum update
```

6. If SELinux support is required, then enable SELinux by following the steps mentioned in, [Chapter 10, Enabling SELinux](#)

7. After SELinux is enabled, set the following boolean:

For Samba

```
setsebool -P samba_load_libgfapi 1
```

For CTDB

```
setsebool -P use_fusefs_home_dirs 1
```

8. To ensure the `glusterd` service does not start automatically after reboot, execute the following command:

```
# chkconfig glusterd off
```

9. Reboot the server.

10. Update the META=all with the gluster volume information in the following scripts:

```
/var/lib/glusterd/hooks/1/start/post/S29CTDBsetup.sh  
/var/lib/glusterd/hooks/1/stop/pre/S29CTDB-teardown.sh
```

11. In `/etc/fstab`, uncomment the line containing the volume used for CTDB service as shown in the following example:

```
HostName:/volname /gluster/lock glusterfs defaults,transport=tcp 0 0
```

12. To automatically start the `glusterd` daemon every time the system boots, run the following command:

```
# chkconfig glusterd on
```

13. To automatically start the ctdb daemon every time the system boots, run the following command:

```
# chkconfig ctdb on
```

14. Start the **glusterd** service using the following command:

```
# service glusterd start
```

15. If using NFS to access volumes, enable gluster-NFS using below command:

```
# gluster volume set <volname> nfs.disable off
```

For example:

```
# gluster volume set testvol nfs.disable off
volume set: success
```

16. Mount the CTDB volume by running the following command:

```
# mount -a
```

17. Start the CTDB service using the following command:

```
# service ctdb start
```

18. To verify if CTDB is running successfully, execute the following commands:

```
# ctdb status
# ctdb ip
# ctdb ping -n all
```

CTDB Upgrade

After upgrading the Red Hat Gluster Storage server, upgrade the CTDB package by executing the following steps:



NOTE

- Upgrading CTDB on all the nodes must be done simultaneously to avoid any data corruption.
- The following steps have to be performed only when upgrading CTDB from CTDB 1.x to CTDB 4.x.

1. Stop the CTDB service on all the nodes of the CTDB cluster by executing the following command. Ensure it is performed on all the nodes simultaneously as two different versions of CTDB cannot run at the same time in the CTDB cluster:

```
# service ctdb stop
```

2. Perform the following operations on all the nodes used as samba servers:

- o Remove the following soft links:

```
/etc/sysconfig/ctdb  
/etc/ctdb/nodes  
/etc/ctdb/public_addresses
```

- o Copy the following files from the CTDB volume to the corresponding location by executing the following command on each node of the CTDB cluster:

```
cp /gluster/lock/nodes /etc/ctdb/nodes  
cp /gluster/lock/public_addresses /etc/ctdb/public_addresses
```

3. Stop and delete the CTDB volume by executing the following commands on one of the nodes of the CTDB cluster:

```
# gluster volume stop volname
```

```
# gluster volume delete volname
```

4. To update CTDB, execute the following command:

```
# yum update
```

For more information about configuring CTDB on a Red Hat Gluster Storage server, see section *Setting Up CTDB* in the *Red Hat Gluster Storage Administration Guide*

8.2.4.2. Verifying In-Service Software Upgrade

To verify if you have upgraded to the latest version of the Red Hat Gluster Storage server execute the following command:

```
# gluster --version
```

8.2.4.3. Upgrading the Native Client

All clients must use the same version of glusterfs-fuse. Red Hat strongly recommends that you upgrade servers before upgrading clients. For more information regarding installing and upgrading native client, refer the *Red Hat Gluster Storage Administration Guide*

CHAPTER 9. UPGRADING FROM RED HAT GLUSTER STORAGE 3.1 TO RED HAT GLUSTER STORAGE 3.2 IN A RED HAT ENTERPRISE VIRTUALIZATION-RED HAT GLUSTER STORAGE ENVIRONMENT

This section describes the upgrade methods for a Red Hat Gluster Storage and Red Hat Enterprise Virtualization integrated environment. You can upgrade Red Hat Gluster Storage 3.1 to Red Hat Gluster Storage 3.2 using an ISO or yum.



WARNING

Before you upgrade, be aware of changed requirements that exist after Red Hat Gluster Storage 3.1.3. If you want to access a volume being provided by a Red Hat Gluster Storage 3.1.3 or higher server, your client must also be using Red Hat Gluster Storage 3.1.3 or higher. Accessing volumes from other client versions can result in data becoming unavailable and problems with directory operations. This requirement exists because Red Hat Gluster Storage 3.1.3 contained a number of changes that affect how the Distributed Hash Table works in order to improve directory consistency and remove the effects seen in [BZ#1115367](#) and [BZ#1118762](#).



IMPORTANT

In Red Hat Enterprise Linux 7 based Red Hat Gluster Storage, updating to 3.1 or higher reloads firewall rules. All runtime-only changes made before the reload are lost.

9.1. UPGRADING USING AN ISO

1. Using Red Hat Enterprise Virtualization Manager, stop all the virtual machine instances.

The Red Hat Gluster Storage volume on the instances will be stopped during the upgrade.

2. Using Red Hat Enterprise Virtualization Manager, move the data domain of the data center to *Maintenance* mode.
3. Using Red Hat Enterprise Virtualization Manager, stop the volume (the volume used for data domain) containing Red Hat Gluster Storage nodes in the data center.
4. Using Red Hat Enterprise Virtualization Manager, move all Red Hat Gluster Storage nodes to *Maintenance* mode.
5. Perform the ISO Upgrade as mentioned in section [Section 8.2.1.3, “Configuring repo for Upgrading using ISO”](#).
6. Re-install the Red Hat Gluster Storage nodes from Red Hat Enterprise Virtualization Manager.



NOTE

- Re-installation for the Red Hat Gluster Storage nodes should be done from Red Hat Enterprise Virtualization Manager. The newly upgraded Red Hat Gluster Storage 3.2 nodes lose their network configuration and other configuration details, such as firewall configuration, which was done while adding the nodes to Red Hat Enterprise Virtualization Manager. Re-install the Red Hat Gluster Storage nodes to have the bootstrapping done.
- You can re-configure the Red Hat Gluster Storage nodes using the option provided under **Action Items**, as shown in [Figure 9.1, “Red Hat Gluster Storage Node before Upgrade”](#), and perform bootstrapping.

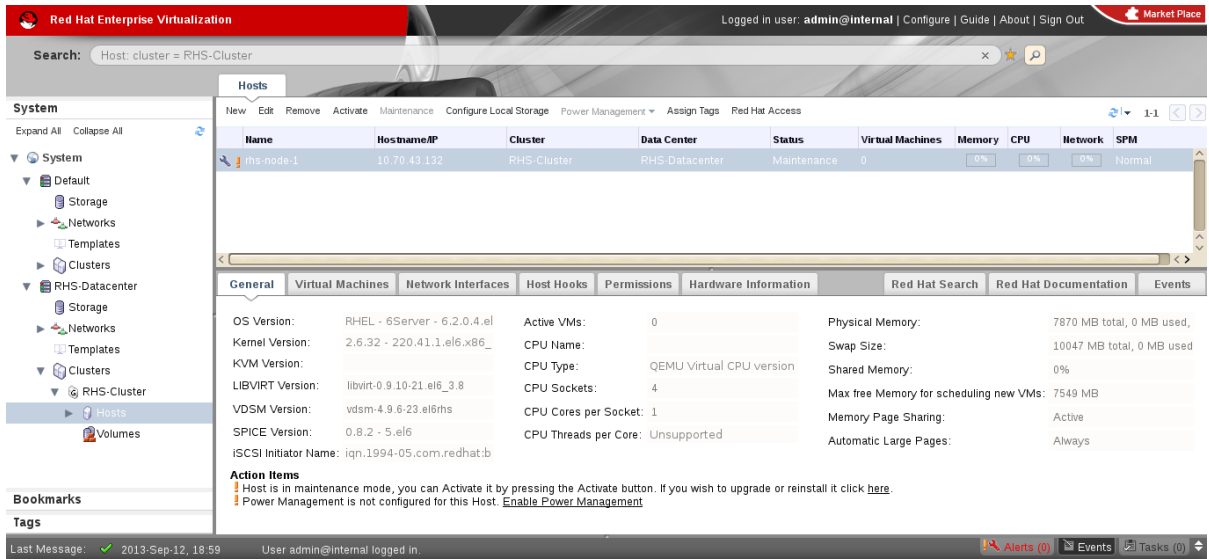


Figure 9.1. Red Hat Gluster Storage Node before Upgrade

7. Perform the steps above in all Red Hat Gluster Storage nodes.
8. Start the volume once all the nodes are shown in **Up** status in Red Hat Enterprise Virtualization Manager.
9. Upgrade the native client bits for Red Hat Enterprise Linux 6 or Red Hat Enterprise Linux 7, depending on the hypervisor you are using.



NOTE

If Red Hat Enterprise Virtualization Hypervisor is used as hypervisor, then install a suitable build of Red Hat Enterprise Virtualization Hypervisor containing the latest native client.

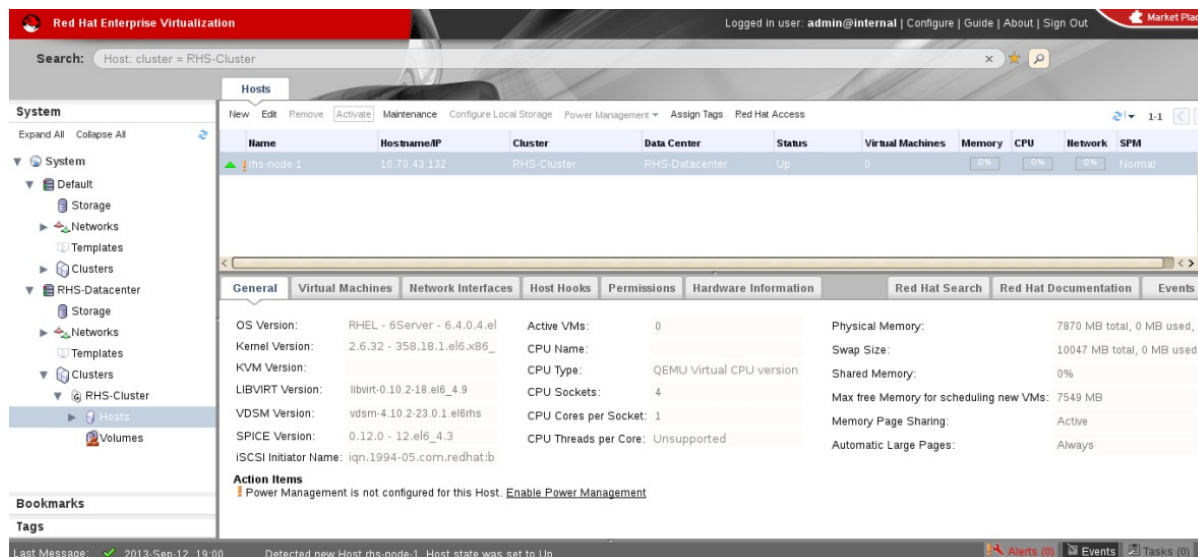


Figure 9.2. Red Hat Gluster Storage Node after Upgrade

- Using Red Hat Enterprise Virtualization Manager, activate the data domain and start all the virtual machine instances in the data center.

9.2. UPGRADING USING YUM

- Using Red Hat Enterprise Virtualization Manager, stop all virtual machine instances in the data center.
- Using Red Hat Enterprise Virtualization Manager, move the data domain backed by gluster volume to *Maintenance* mode.
- Using Red Hat Enterprise Virtualization Manager, move all Red Hat Gluster Storage nodes to *Maintenance* mode.
- Perform `yum` update as mentioned in [Section 8.1.1, “Upgrading from Red Hat Gluster Storage 3.1 to Red Hat Gluster Storage 3.2 for Systems Subscribed to Red Hat Network”](#).
- Once the Red Hat Gluster Storage nodes are rebooted and up, **Activate** them using Red Hat Enterprise Virtualization Manager.



NOTE

Re-installation of Red Hat Gluster Storage nodes is required, as the network configurations and bootstrapping configurations done prior to upgrade are preserved, unlike ISO upgrade.

- Using Red Hat Enterprise Virtualization Manager, start the volume.
- Upgrade the native client bits for Red Hat Enterprise Linux 6 or Red Hat Enterprise Linux 7, based on the Red Hat Enterprise Linux server hypervisor used.



NOTE

If Red Hat Enterprise Virtualization Hypervisor is used as hypervisor, reinstall Red Hat Enterprise Virtualization Hypervisor containing the latest version of Red Hat Gluster Storage native client.

8. Activate the data domain and start all the virtual machine instances.

CHAPTER 10. ENABLING SELINUX

After upgrading to Red Hat Gluster Storage 3.1, follow the steps mentioned here for SELinux support. These steps are not mandatory and are required only if SELinux support is required.

This procedure assumes that the `selinux-policy-targeted`, `selinux-policy`, `libselinux`, `libselinux-python`, `libselinux-utils`, `policycoreutils`, `policycoreutils-python`, `setroubleshoot`, `setroubleshoot-server`, `setroubleshoot-plugins` packages are installed. To verify that the packages are installed, use the following command:

```
# rpm -q package_name
```



IMPORTANT

If the system was initially installed without SELinux, particularly the `selinux-policy` package, one additional step is necessary to enable SELinux. To make sure SELinux is initialized during system startup, the `dracut` utility has to be run to put SELinux awareness into the `initramfs` file system. Failing to do so causes SELinux to not start during system startup.

1. Before SELinux is enabled, each file on the file system must be labeled with an SELinux context. Before this happens, confined domains may be denied access, preventing your system from booting correctly. To prevent this, configure `SELINUX=permissive` in `/etc/selinux/config`:

```
# This file controls the state of SELinux on the system.
# SELINUX= can take one of these three values:
#     enforcing - SELinux security policy is enforced.
#     permissive - SELinux prints warnings instead of enforcing.
#     disabled - No SELinux policy is loaded.
SELINUX=permissive
# SELINUXTYPE= can take one of these two values:
#     targeted - Targeted processes are protected,
#     mls - Multi Level Security protection.
SELINUXTYPE=targeted
```

2. As the Linux root user, reboot the system. During the next boot, file systems are labeled. The label process labels each file with an SELinux context:

```
*** Warning -- SELinux targeted policy relabel is required.
*** Relabeling could take a very long time, depending on file
*** system size and speed of hard drives.
****
```

Each * (asterisk) character on the bottom line represents 1000 files that have been labeled. In the above example, four * characters represent 4000 files have been labeled. The time it takes to label all files depends on the number of files on the system and the speed of hard drives. On modern systems, this process can take as short as 10 minutes.

3. In permissive mode, the SELinux policy is not enforced, but denial messages are still logged for actions that would have been denied in enforcing mode. Before changing to enforcing mode, as the Linux root user, run the following command to confirm that SELinux did not deny actions during the last boot:

```
# grep "SELinux is preventing" /var/log/messages
```

If SELinux did not deny any actions during the last boot, this command returns no output.

4. If there were no denial messages in `/var/log/messages`, configure `SELINUX=enforcing` in `/etc/selinux/config`:

```
# This file controls the state of SELinux on the system.
# SELINUX= can take one of these three values:
#     enforcing - SELinux security policy is enforced.
#     permissive - SELinux prints warnings instead of enforcing.
#     disabled - No SELinux policy is loaded.
SELINUX=enforcing
# SELINUXTYPE= can take one of these two values:
#     targeted - Targeted processes are protected,
#     mls - Multi Level Security protection.
SELINUXTYPE=targeted
```

5. Reboot your system. After reboot, confirm that `getenforce` returns `Enforcing`

```
~]$ getenforce
Enforcing
```

For more information see, https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/6/html/Security-Enhanced_Linux/sect-Security-Enhanced_Linux-Working_with_SELinux-Changing_SELinux_Modes.html#sect-Security-Enhanced_Linux-Enabling_and_Disabling_SELinux-Enabling_SELinux

CHAPTER 11. USING THE GLUSTER COMMAND LINE INTERFACE

The Gluster command line interface (CLI) simplifies configuration and management of the storage environment. The Gluster CLI is similar to the LVM (Logical Volume Manager) CLI or the ZFS CLI, but operates across multiple storage servers. The Gluster CLI can be used when volumes are mounted (active) and not mounted (inactive). Red Hat Gluster Storage automatically synchronizes volume configuration information across all servers.

Use the Gluster CLI to create new volumes, start and stop existing volumes, add bricks to volumes, remove bricks from volumes, and change translator settings. Additionally, the Gluster CLI commands can create automation scripts and use the commands as an API to allow integration with third-party applications.



NOTE

Appending `--mode=script` to any CLI command ensures that the command executes without confirmation prompts.

Running the Gluster CLI

Run the Gluster CLI on any Red Hat Gluster Storage Server by either invoking the commands or running the Gluster CLI in interactive mode. The `gluster` command can be remotely used via SSH.

Run commands directly as follows, after replacing *COMMAND* with the required command:

```
# gluster peer COMMAND
```

The following is an example using the `status` command:

```
# gluster peer status
```

Gluster CLI Interactive Mode

Alternatively, run the Gluster CLI in interactive mode using the following command:

```
# gluster
```

If successful, the prompt changes to the following:

```
gluster>
```

When the prompt appears, execute gluster commands from the CLI prompt as follows:

```
gluster> COMMAND
```

As an example, replace the *COMMAND* with a command such as `status` to view the status of the peer server:

1. Start Gluster CLI's interactive mode:

```
# gluster
```

2. Request the peer server status:

```
█ gluster> status
```

3. The peer server status displays.

The following is another example, replacing the *COMMAND* with a command such as `help` to view the gluster help options.

1. Start Gluster CLI's interactive mode:

```
█ # gluster
```

2. Request the help options:

```
█ gluster> help
```

3. A list of gluster commands and options displays.

APPENDIX A. REVISION HISTORY

Revision 3.2-2**Fri Mar 17 2017****Divya Muntimadugu, Laura Bailey,
Bhavana Mohanraj, Anjana Sriram**

Clarified NFS-Ganesha installation instructions. (BZ#1402253)
Updated details of NFS-Ganesha upgrade process. (BZ#1368444)
Added a note to check for self-heal status before an in-service update. (BZ#1376287)
Updated thin pool examples throughout the guide to give appropriate size recommendations. (BZ#1390509)
Updated upgrade process to account for new shared storage volume requirements. (BZ#1414363)
Updated upgrade process to ensure write-behind behavior is disabled. (BZ#1409569)
Various updates to ensure document refers to version 3.2. (BZ#1406250)
Rewrote update and upgrade procedures for version 3.2. (BZ#1422768)
Applied feedback to improve clarity of rewritten upgrade process. (BZ#1432782)
Updated guide to reflect changed glusterd log file name. (BZ#1365003)
Updated the upgrade chapter remove the offline ISO upgrade. (BZ#1434377)

Revision 3.2-1**Fri Mar 17 2017****Bhavana Mohan**

Version for 3.2 GA.