Installation Guide

Installing Red Hat Gluster Storage 3.4

Gluster Storage Documentation Team, Red Hat `<gluster-docs@redhat.com>`

Installation Guide: Installing Red Hat Gluster Storage 3.4

by Gluster Storage Documentation Team

Abstract

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

Table of Contents

List of Figures

List of Tables

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.

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 3.4 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.

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) or Chrony daemon. For more information about these features see Choosing Between NTP Daemons [https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html-single/system_administrators_guide/#sect-Choosing_between_NTP_daemon]

Configuring time syncronization using Chrony

For information on configuring chrony for time synchronization for all Gluster nodes, see Using chrony [https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html-single/system_administrators_guide/#sect-Using_chrony].

After configuring chrony, verify that time is synchronized for all Gluster nodes by using the following command from the primary node to all other nodes:

```
# clockdiff node-hostname
```

Configuring time syncronization using Network Time Protocol

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*.

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.

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.
2049	For glusterFS's NFS exports (nfsd process).
24007	For glusterd (for management).
24008	For glusterd RDMA port management.
24009	For glustereventsd.
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).
49152 - 49664	Each brick on a host requires its own port for communication with GlusterFS clients. One port is used for each brick, in ascending order, beginning with port 49152. By default, up to 512 ports will be used. This maximum can be increased by editing the port range configuration, as described in Port Range Configuration [https://access.redhat.com/documentation/en-us/red_hat_gluster_storage/3.4/html-single/administration_guide/#port-range-

Port Number	Usage
	configuration]. Consider the maximum number of bricks that you expect to exist simultaneously on your host, and open only the ports that you need for your current configuration to help ensure system security.
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.

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	Ansible
3.4	3.12.2-18	31302	SMB 1, 2.0, 2.1, 3.0, 3.1.1	NFSv3, NFSv4	gdeploy-2.0.2-27	

RHGS version	glusterfs and glusterfs-fuse	RHGS op-version	SMB	NFS	gDeploy	Ansible
3.4 Batch 1 Update	3.12.2-25	31303	SMB 1, 2.0, 2.1, 3.0, 3.1.1	NFSv3, NFSv4	gdeploy-2.0.2-30	
3.4 Batch 2 Update	3.12.2-32	31304	SMB 1, 2.0, 2.1, 3.0, 3.1.1	NFSv3, NFSv4	gdeploy-2.0.2-3	gluster-ansible-infra-1.0.2-2
3.4 Batch 3 Update	3.12.2-40	31305	SMB 1, 2.0, 2.1, 3.0, 3.1.1	NFSv3, NFSv4	gdeploy-2.0.2-3	gluster-ansible-infra-1.0.2-2

Note

For more information regarding the version details of previous Red Hat Gluster Storage releases, see https://access.redhat.com/solutions/543123.

The repository that you install packages from may affect the package version installed which in turn may affect the compatibility of your systems. Red Hat recommends installing the client from the Red Hat Gluster Storage repository:

- For Red Hat Enterprise Linux 7 clients: `rhel-x86_64-server-7-rh-gluster-3-client`

- For Red Hat Enterprise Linux 6 clients: `rhel-x86_64-server-rhsclient-6`

Feature Compatibility Support

Red Hat Gluster Storage supports a number of features. Most features are supported with other features, but there are some exceptions. This section clearly identifies which features are supported and compatible with other features to help you in planning your Red Hat Gluster Storage deployment.

Important

Virtual Data Optimizer (VDO) volumes, which are supported in Red Hat Enterprise Linux 7.5, are not currently supported in Red Hat Gluster Storage. VDO is supported only when used as part of Red Hat Hyperconverged Infrastructure for Virtualization 2.0. See Understanding VDO [https://access.redhat.com/documentation/en-us/red_hat_hyperconverged_infrastructure_for_virtualization/2.0/html-single/deploying_red_hat_hyperconverged_infrastructure_for_virtualization/#understanding-vdo] for more information.

Features in the following table are supported from the specified version and later.

Table 1.6. Features supported by Red Hat Gluster Storage version

Feature	Version
Arbiter bricks	3.2
Bitrot detection	3.1
Erasure coding	3.1
Google Compute Engine	3.1.3

Feature	Version
Metadata caching	3.2
Microsoft Azure	3.1.3
NFS version 4	3.1
SELinux	3.1
Sharding	3.2.0
Snapshots	3.0
Snapshots, cloning	3.1.3
Snapshots, user-serviceable	3.0.3
Tiering	3.1.2
Volume Shadow Copy (VSS)	3.1.3

Table 1.7. Features supported by volume type

Volume Type	Sharding	Tiering	Quota	Snapshots	Geo-Rep	Bitrot
Arbitrated-Replicated	Yes	No	Yes	Yes	Yes	Yes
Distributed	No	Yes	Yes	Yes	Yes	Yes
Distributed-Dispersed	No	Yes	Yes	Yes	Yes	Yes
Distributed-Replicated	Yes	Yes	Yes	Yes	Yes	Yes
Replicated	Yes	Yes	Yes	Yes	Yes	Yes
Sharded	N/A	No	No	No	Yes	No
Tiered	No	N/A	Limited[a]	Limited[a] [https://access.redhat.com/documentation/en-us/red_hat_gluster_storage/3.4/html-single/installation_guide/#ftn.tiering-footnote]	Limited[a] [https://access.redhat.com/documentation/en-us/red_hat_gluster_storage/3.4/html-single/installation_guide/#ftn.tiering-footnote]	Limited[a] [https://access.redhat.com/documentation/en-us/red_hat_gluster_storage/3.4/html-single/installation_guide/#ftn.tiering-footnote]

[a] See Section 17.3. Tiering Limitations [https://access.redhat.com/documentation/en-us/red_hat_gluster_storage/3.4/html-single/administration_guide/#chap-Managing_Data_Tiering-Limitations] in the *Red Hat Gluster Storage 3.4 Administration Guide* for details.

Table 1.8. Features supported by client protocol

Feature	FUSE	Gluster-NFS	NFS-Ganesha	SMB	Swift/S3
Arbiter	Yes	Yes	Yes	Yes	No

Feature	FUSE	Gluster-NFS	NFS-Ganesha	SMB	Swift/S3
Bitrot detection	Yes	Yes	No	Yes	No
dm-cache	Yes	Yes	Yes	Yes	Yes
Encryption (TLS-SSL)	Yes	Yes	Yes	Yes	No
Erasure coding	Yes	Yes	Yes	Yes	No
Export subdirectory	Yes	Yes	Yes	N/A	N/A
Geo-replication	Yes	Yes	Yes	Yes	Yes
Quota	Yes	Yes	Yes	Yes	No
RDMA	Yes	No	No	No	N/A
Snapshots	Yes	Yes	Yes	Yes	Yes
Snapshot cloning	Yes	Yes	Yes	Yes	Yes
Tiering	Yes	Yes	N/A	N/A	N/A

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 3.4 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.

Obtaining Red Hat Gluster Storage

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

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.

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.

Installing from an ISO Image

Installing Red Hat Gluster Storage 3.4 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 the section called "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.

    > ## Important
    >
    > Virtual Data Optimizer (VDO) volumes, which are supported in Red Hat Enterprise Linux 7.5, are not currently supported in Red Hat Gluster Storage. VDO is supported only when used as part of Red Hat Hyperconverged Infrastructure for Virtualization 2.0. See Understanding VDO [https://access.redhat.com/documentation/en-us/red_hat_hyperconverged_infrastructure_for_virtualization/2.0/html-single/deploying_red_hat_hyperconverged_infrastructure_for_virtualization/#understanding-vdo] 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.

   - 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.

   - 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*.

Installing Red Hat Gluster Storage 3.4 on Red Hat Enterprise Linux 7.5 and higher

Important

Virtual Data Optimizer (VDO) volumes, which are supported in Red Hat Enterprise Linux 7.5, are not currently supported in Red Hat Gluster Storage. VDO is supported only when used as part of Red Hat Hyperconverged Infrastructure for Virtualization 2.0. See Understanding VDO [https://access.redhat.com/documentation/en-us/ red_hat_hyperconverged_infrastructure_for_virtualization/2.0/html-single/ deploying_red_hat_hyperconverged_infrastructure_for_virtualization/#understanding-vdo] for more information.

To install Red Hat Gluster Storage 3.4 on Red Hat Enterprise Linux 7.5 and higher:

1. Download the ISO image file for Red Hat Gluster Storage Server as described in the section called "Obtaining Red Hat Gluster Storage"

2. In the Welcome to Red Hat Gluster Storage 3.4 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 3.4 installation program on Red Hat Enterprise Linux 7.5 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.**

  ### Important

  Setting up NTP is mandatory for Gluster installation.
  To configure NTP, perform the following steps:

  a. Click Date & Time and specify a time zone to maintain the accuracy of the system clock.

  b. Toggle the Network Time switch to ON.

     ### Note

     By default, the Network Time switch is enabled if you are connected to the network.

  c. Click the configuration icon to add new NTP servers or select existing servers.

  d. Once you have made your addition or selection, click Done to return to the Installation Summary screen.

     ### Note

     NTP servers might be unavailable at the time of installation. In such a case, enabling them will not set the time automatically. When the servers are available, the date and time will be updated.

- **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 3.4 on Red Hat Enterprise Linux 7.5 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 3.4 on Red Hat Enterprise Linux 7.5 from a DVD to a local hard drive.

  When you install Red Hat Gluster Storage 3.4 on Red Hat Enterprise Linux 7.5 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.

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

Virtual Data Optimizer (VDO) volumes, which are supported in Red Hat Enterprise Linux 7.5, are not currently supported in Red Hat Gluster Storage. VDO is supported only when used as part of Red Hat Hyperconverged Infrastructure for Virtualization 2.0. See Understanding VDO [https://access.redhat.com/documentation/en-us/ red_hat_hyperconverged_infrastructure_for_virtualization/2.0/html-single/ deploying_red_hat_hyperconverged_infrastructure_for_virtualization/#understanding-vdo] for more information.

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.

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 3.4 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-rp
```

For more information about installing Samba, see Chapter 3, *Deploying Samba on Red Hat Gluster Storage*

3. NFS-Ganesha is not supported on Red Hat Enterprise Linux 6 based installations.

For Red Hat Gluster Storage 3.4 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-rp
```

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 repositories:

```
# 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.

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.

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 the section called "Port Information".

Using Red Hat Satellite Server 6.x

Note

For detailed steps for Satellite, see Red Hat Satellite [https://access.redhat.com/documentation/en-us/red_hat_satellite/6.4/] documentation.

1. Create a new manifest file and upload the manifest in the Satellite 6 server.

2. Search for the required Red Hat Gluster Storage repositories and enable them.

3. Synchronize all repositories enabled for Red Hat Gluster Storage.

4. Create a new content view and add all the required products.

5. Publish the content view and create an activation key.

6. Register the required clients

    ```
    # rpm -Uvh satellite-server-host-address/pub/katello-ca-consumer-latest.noarch.
    # subscription-manager register --org="Organization_Name" --activationkey="Acti
    ```

7. Identify available entitlement pools

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

8. Attach entitlement pools to the system

    ```
    # subscription-manager attach --pool=Pool_ID
    ```

9. Subscribe to required channels

    1. Enable the RHEL and Gluster channel

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

    2. If you require Samba, enable its repository

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

    3. If you require NFS-Ganesha, enable its repository

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

    4. If you require HA, enable its repository

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

    5. If you require gdeploy, enable the Ansible repository

        ```
        # subscription-manager repos --enable=rhel-7-server-ansible-2-rpms
        ```

    6. If you require Web Administration, enable its repository

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

10. Install Red Hat Gluster Storage

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

Using Red Hat Satellite Server 5.x

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.

   ### 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:

```
--enableshadow --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*.

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 the section called "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 Se
   # 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 3.4 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-rp
```

3. NFS-Ganesha is not supported on Red Hat Enterprise Linux 6 based installations.

For Red Hat Gluster Storage 3.4 on Red Hat Enterprise Linux 7.5.

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-rp
```

3. For Red Hat Gluster Storage 3.4, 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 3.4, if you require CTDB, then enable the following repository:

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

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 RHEL 6 - 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 3.4 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_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. NFS-Ganesha is not supported on Red Hat Enterprise Linux 6 based installations.

   For Red Hat Gluster Storage 3.4 on Red Hat Enterprise Linux 7.5.

   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 3.4, for NFS-Ganesha enable the following channel:

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

   4. For Red Hat Gluster Storage 3.4, 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
   ```

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 by using the `glusterd` service commands. For more information on creating trusted storage pools, see the *Red Hat Gluster Storage 3.4 Administration Guide*: https://access.redhat.com/documentation/en-us/red_hat_gluster_storage/3.4/html-single/administration_guide/#idm140612789187008.

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.

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
  ```

Important

If `glusterd` crashes, there is no functionality impact to this crash as it occurs during the shutdown. For more information, see Resolving `glusterd` crash [https://access.redhat.com/documentation/en-us/red_hat_gluster_storage/3.4/html-single/administration_guide/#glusterd_crashing]

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 2.5 or higher to use gDeploy.

1. Execute the following command to enable the repository required to install Ansible:

   ```
   # subscription-manager repos --enable=rhel-7-server-ansible-2-rpms
   ```

2. Install `ansible` by executing the following command:

   ```
   # yum install ansible
   ```

Installing Native Client

Native Client is a FUSE-based client in the user space. Native Client is the recommended method for accessing Red Hat Gluster Storage volumes when high concurrency and high write performance

is required. For more information regarding installing native client, refer to section Installing Native Client [https://access.redhat.com/documentation/en-us/red_hat_gluster_storage/3.4/html-single/administration_guide/#Installing_Native_Client] in the *Red Hat Gluster Storage Administration Guide*.

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.

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 3.4 on Red Hat Enterprise Linux 6.x.

   a. 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-rpm
   ```

   b. 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 3.4 on Red Hat Enterprise Linux 7.x.

   a. 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-rpm
```

b. 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
```

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 the section called "Installing from an ISO Image"

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 the section called "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

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.

For Red Hat Gluster Storage 3.4 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
   ```

   To add the HA repository, execute the following command:

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

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
   ```

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

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

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.

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

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*.

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.

Supported Deployment Configuration

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

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 Enterprsie Linux Atomic Host
3.2	1.13	7.3
3.3	1.12	7.4

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/.

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 7.5 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 3.4.x to 3.4.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.

To keep your Red Hat Gluster Storage system up-to-date, associate the system with the Red Hat Network 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 3.4 installation:

- This is an update between minor releases of the same major product version. If you want to upgrade between major product versions, see Chapter 8, *Upgrading to Red Hat Gluster Storage 3.4*.

- 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 a Red Hat Gluster Storage server node that hosts a very large number of bricks or snapshots is updated, cluster management commands may become unresponsive as glusterd attempts to start all brick processes concurrently for all bricks and snapshots. If you have more than 250 bricks or snapshots being hosted by a single node, Red Hat recommends deactivating snapshots until the update is complete.

- After performing inservice upgrade of NFS-Ganesha, the new configuration file is saved as "ganesha.conf.rpmnew" in /etc/ganesha folder. The old configuration file is not overwritten during the inservice upgrade process. However, post upgradation, you have to manually copy any new configuration changes from "ganesha.conf.rpmnew" to the existing ganesha.conf file in /etc/ganesha folder.

Before you update

Make a complete backup using a reliable backup solution before you update your system. This Knowledge Base solution covers one possible approach: https://access.redhat.com/solutions/1484053.

Regardless of the approach that you use:

- Back up to a location that is not on the operating system partition.

- Ensure that you have sufficient space available for a complete backup.

- Copy the .glusterfs directory before copying any data files.

- Ensure that the following configuration directories and files are backed up.

  - /var/lib/glusterd

  - /etc/swift

  - /etc/samba

  - /etc/ctdb

  - /etc/glusterfs

- `/var/lib/samba`

- `/var/lib/ctdb`

- `/var/run/gluster/shared_storage/nfs-ganesha`

If you use NFS-Ganesha, also ensure that the following files from all nodes are backed up:

- `/etc/ganesha/exports/export.*.conf`

- `/etc/ganesha/ganesha.conf`

- `/etc/ganesha/ganesha-ha.conf`

- Ensure that all extended attributes, ACLs, owners, groups, and symbolic and hard links are backed up.

- Ensure that no new files are created on Red Hat Gluster Storage file systems during the backup or update process.

- Check that the backup restores correctly before you continue with the update:

  - Create a new thin provisioned logical volume. For more information, see https://access.redhat.com/site/documentation/en-US/Red_Hat_Enterprise_Linux/6/html/Logical_Volume_Manager_Administration/thinprovisioned_volumes.html

  - Restore the backed up content to the newly created thin provisioned logical volume.

Updating Red Hat Gluster Storage in the Offline Mode

Important

- Offline updates result in server downtime, as volumes are offline during the update process.

- Complete updates to all Red Hat Gluster Storage servers before updating any clients.

Important

The Web Administration update sequence outlined below is part of the overall Red Hat Gluster Storage offline update. Updating Web Administration as a standalone component is not supported.

Red Hat Gluster Storage Web Administration Update

For Red Hat Gluster Storage Web Administration users intending to update to the latest version, see the Red Hat Gluster Storage Web Administration 3.4.x to 3.4.y [https://access.redhat.com/documentation/en-us/red_hat_gluster_storage/3.4/html-single/quick_start_guide/#red_hat_gluster_storage_web_administration_3_4_x_to_3_4_y] section of the Red Hat Gluster Storage Web Administration Quick Start Guide for detailed steps. Follow this sequence to update your Red Hat Gluster Storage system along with the Web Administration environment:

1. Visit the section Red Hat Gluster Storage Web Administration 3.4.x to 3.4.y [https://access.redhat.com/documentation/en-us/red_hat_gluster_storage/3.4/html-single/quick_start_guide/#red_hat_gluster_storage_web_administration_3_4_x_to_3_4_y] and stop all Web

Administration services on the Red Hat Gluster Storage servers outlined in step 1 under the heading *On Red Hat Gluster Storage Servers (Part 1).*

2. Return to this section 7.2. Updating Red Hat Gluster Storage in the Offline Mode [https://access.redhat.com/documentation/en-us/red_hat_gluster_storage/3.4/html-single/installation_guide/#Updating_Red_Hat_Storage_in_the_Offline_Mode] and execute the steps outlined under the heading *Updating Red Hat Gluster Storage 3.4 in the offline mode*.

3. Navigate back to the Red Hat Gluster Storage Web Administration 3.4.x to 3.4.y [https://access.redhat.com/documentation/en-us/red_hat_gluster_storage/3.4/html-single/quick_start_guide/#red_hat_gluster_storage_web_administration_3_4_x_to_3_4_y] section and perform the steps identified under *On Web Administration Server* and *On Red Hat Gluster Storage Servers (Part II)* to complete the Red Hat Gluster Storage and Web Administration update process.

Updating Red Hat Gluster Storage 3.4 in the offline mode

1. Ensure that you have a working backup, as described in the section called "Before you update".

2. Stop all volumes.

   ```
   # for vol in `gluster volume list`; do gluster --mode=script volume stop $vol; s
   ```

3. Run the following commands on one server at a time.

   a. Stop all gluster services.

   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
   ```

   > ## Important
   >
   > If `glusterd` crashes, there is no functionality impact to this crash as it occurs during the shutdown. For more information, see Resolving `glusterd` crash [https://access.redhat.com/documentation/en-us/red_hat_gluster_storage/3.4/html-single/administration_guide/#glusterd_crashing]

   b. If you want to migrate from Gluster NFS to NFS Ganesha as part of this update, perform the following additional steps.

   i. Stop and disable CTDB. This ensures that multiple versions of Samba do not run in the cluster during the update process, and avoids data corruption.

   ```
   # systemctl stop ctdb
   # systemctl disable ctdb
   ```

   ii. Verify that the CTDB and NFS services are stopped:

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

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

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

c. Update the system.

```
# yum update
```

Review the packages to be updated, and enter `y` to proceed with the update when prompted.

Wait for the update to complete.

d. If updates to the kernel package occurred, or if you are migrating from Gluster NFS to NFS Ganesha as part of this update, reboot the system.

e. Start glusterd.

On Red Hat Enterprise Linux 7:

```
# systemctl start glusterd
```

On Red Hat Enterprise Linux 6:

```
# service glusterd start
```

4. When all servers have been updated, run the following command to update the cluster operating version. This helps to prevent any compatibility issues within the cluster.

```
# gluster volume set all cluster.op-version 31305
```

Note

31305 is the `cluster.op-version` value for Red Hat Gluster Storage 3.4 Batch 3 Update. Refer to the section called "Supported Versions of Red Hat Gluster Storage" for the correct `cluster.op-version` value for other versions.

5. If you want to migrate from Gluster NFS to NFS Ganesha as part of this update, install the NFS-Ganesha packages as described in Chapter 4, *Deploying NFS-Ganesha on Red Hat Gluster Storage*, and use the information in the NFS Ganesha [https://access.redhat.com/documentation/en-us/red_hat_gluster_storage/3.4/html/administration_guide/nfs#nfs_ganesha] section of the *Red Hat Gluster Storage 3.4 Administration Guide* to configure the NFS Ganesha cluster.

6. Start all volumes.

```
# for vol in `gluster volume list`; do gluster --mode=script volume start $vol;
```

7. If you did not reboot as part of the update process, run the following command to remount 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
```

8. If you use Gluster 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
```

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

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

You may need to append the `force` parameter to successfully restart in some circumstances. See BZ#1347625 [https://bugzilla.redhat.com/show_bug.cgi?id=1347625] for details.

Note

If you are updating your Web Administration environment, after executing step 9, navigate to the Red Hat Gluster Storage Web Administration 3.4.x to 3.4.y [https://access.redhat.com/documentation/en-us/red_hat_gluster_storage/3.4/html-single/quick_start_guide/#red_hat_gluster_storage_web_administration_3_4_x_to_3_4_y] section and perform the steps identified under *On Web Administration Server* and *On Red Hat Gluster Storage Servers (Part II)* to complete the Red Hat Gluster Storage and Web Administration update process.

In-service Software Update from Red Hat Gluster Storage

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.

Important

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.

Before you update, be aware:

- Complete updates to all Red Hat Gluster Storage servers before updating any clients.

- If geo-replication is in use, complete updates to all slave nodes before updating master nodes.

- Erasure coded (dispersed) volumes can be updated while in-service only if the `disperse.optimistic-change-log` and `disperse.eager-lock` options are set to `off`. Wait for at least two minutes after disabling these options before attempting to upgrade to ensure that these configuration changes take effect for I/O operations.

- 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 in order to receive updates. For more information, see the section called "Subscribing to the Red Hat Gluster Storage Server Channels"

• Do not perform any volume operations while the cluster is being updated.

Updating Red Hat Gluster Storage 3.4 in in-service mode

1. Ensure that you have a working backup, as described in the section called "Before you update".

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

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

   • Stop any geo-replication sessions.

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

   • If this node is part of an NFS-Ganehsa cluster, place the node in standby mode.

     ```
     # pcs cluster standby
     ```

   • Verify that there are no pending self-heals:

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

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

   • If this node is part of an NFS-Ganesha cluster:

     a. Disable the PCS cluster and verify that it has stopped.

        ```
        # pcs cluster disable
        # pcs status
        ```

     b. Stop the nfs-ganesha service.

        ```
        # systemctl stop nfs-ganesha
        ```

   • If you need to update an erasure coded (dispersed) volume, set the `disperse.optimistic-change-log` and `disperse.eager-lock` options to `off`. Wait for at least two minutes after disabling these options before attempting to upgrade to ensure that these configuration changes take effect for I/O operations.

     ```
     # gluster volume set volname disperse.optimistic-change-log off
     # gluster volume set volname disperse.eager-lock off
     ```

   • Stop the gluster services on the storage server using the following commands:

     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
     ```

38

Important

> If `glusterd` crashes, there is no functionality impact to this crash as it occurs during the shutdown. For more information, see Resolving `glusterd` crash [https://access.redhat.com/documentation/en-us/red_hat_gluster_storage/3.4/html-single/administration_guide/#glusterd_crashing]

- If you use Samba:

  a. 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-rp
     ```

     On Red Hat Enterprise Linux 7:

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

  b. 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.

     On Red Hat Enterprise Linux 7:

     ```
     # systemctl stop ctdb
     # systemctl disable ctdb
     ```

     On Red Hat Enterprise Linux 6:

     ```
     # service ctdb stop
     # chkconfig ctdb off
     ```

     This ensures different versions of Samba do not run in the same Samba cluster until all Samba nodes are updated.

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

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

- Update the server using the following command:

  ```
  # yum update
  ```

  Take note of the packages being updated, and wait for the update to complete.

- If a kernel update was included as part of the update process in the previous step, reboot the server.

- If a reboot of the server was not required, then start the gluster services on the storage server using the following command.

  On Red Hat Enterprise Linux 7:

  ```
  # systemctl start glusterd
  ```

  On Red Hat Enterprise Linux 6:

```
# service glusterd start
```

- Verify that you have updated to the latest version of the Red Hat Gluster Storage server.

```
# gluster --version
```

  Compare output with the desired version in the section called "Supported Versions of Red Hat Gluster Storage".

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

```
# gluster volume status
```

- Start self-heal on the volume.

```
# gluster volume heal volname
```

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

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

- Verify that shared storage is mounted.

```
# mount | grep /run/gluster/shared_storage
```

3. When all nodes in the volume 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 31305
```

   ## Note

   31305 is the *cluster.op-version* value for Red Hat Gluster Storage 3.4 Batch 3 Update. Refer to the section called "Supported Versions of Red Hat Gluster Storage" for the correct *cluster.op-version* value for other versions.

4. If you use Samba:

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

```
# mount -a
```

   b. If all servers that host volumes accessed via SMB have been updated, then start and re-enable the CTDB and Samba services by executing the following commands.

      On Red Hat Enterprise Linux 7:

```
# systemctl start ctdb
# systemctl enable ctdb
```

      On Red Hat Enterprise Linux 6:

```
# service ctdb start
# chkconfig ctdb on
```

   c. To verify that the CTDB and SMB services have started, execute the following command:

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

5. 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
```

6. If this node is part of an NFS-Ganesha cluster:

   a. If SELinux is in use, set the ganesha_use_fusefs Boolean to on.

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

   b. Start the nfs-ganesha service:

   ```
   # systemctl start nfs-ganesha
   ```

   c. Enable and start the cluster.

   ```
   # pcs cluster enable
   # pcs cluster start
   ```

   d. Release the node from standby mode.

   ```
   # pcs cluster unstandby
   ```

   e. Verify that the PCS cluster is running and that the volume is exporting correctly.

   ```
   # pcs status
   # showmount -e
   ```

   NFS-ganesha enters a short grace period after performing these steps. I/O operations halt during this grace period. Wait until you see NFS Server Now NOT IN GRACE in the ganesha.log file before continuing.

7. 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 [https://bugzilla.redhat.com/show_bug.cgi?id=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
   ```

8. If you disabled the disperse.optimistic-change-log and disperse.eager-lock options in order to update an erasure-coded (dispersed) volume, re-enable these settings.

   ```
   # gluster volume set volname disperse.optimistic-change-log on
   # gluster volume set volname disperse.eager-lock on
   ```

Note

After performing inservice upgrade of NFS-Ganesha, the new configuration file is saved as "`ganesha.conf.rpmnew`" in `/etc/ganesha` folder. The old configuration file is not overwritten during the inservice upgrade process. However, post upgradation, you have to manually copy any new configuration changes from "`ganesha.conf.rpmnew`" to the existing `ganesha.conf` file in `/etc/ganesha` folder.

Chapter 8. Upgrading to Red Hat Gluster Storage 3.4

This chapter describes the procedure to upgrade to Red Hat Gluster Storage 3.4 from Red Hat Gluster Storage 3.1, 3.2, or 3.3.

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.

- Virtual Data Optimizer (VDO) volumes, which are supported in Red Hat Enterprise Linux 7.5, are not currently supported in Red Hat Gluster Storage. VDO is supported only when used as part of Red Hat Hyperconverged Infrastructure for Virtualization 2.0. See Understanding VDO [https://access.redhat.com/documentation/en-us/red_hat_hyperconverged_infrastructure_for_virtualization/2.0/html-single/deploying_red_hat_hyperconverged_infrastructure_for_virtualization/#understanding-vdo] for more information.

- Servers must be upgraded prior to upgrading clients.

- If you are upgrading from Red Hat Gluster Storage 3.1 Update 2 or earlier, you must upgrade servers and clients simultaneously.

- If you use NFS-Ganesha, your supported upgrade path to Red Hat Gluster Storage 3.4 depends on the version from which you are upgrading. If you are upgrading from version 3.1.x to 3.4, use the section called "Offline Upgrade to Red Hat Gluster Storage 3.4". If you are upgrading from version 3.3 to 3.4, use the section called "In-Service Software Upgrade from Red Hat Gluster Storage 3.3 to Red Hat Gluster Storage 3.4".

Offline Upgrade to Red Hat Gluster Storage 3.4

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 [https://bugzilla.redhat.com/show_bug.cgi?id=1115367] and BZ#1118762 [https://bugzilla.redhat.com/show_bug.cgi?id=1118762].

Important

In Red Hat Enterprise Linux 7 based Red Hat Gluster Storage 3.1 and higher, updating reloads firewall rules. All runtime-only changes made before the reload are lost, so ensure that any changes you want to keep are made persistently.

Upgrading to Red Hat Gluster Storage 3.4 for Systems Subscribed to Red Hat Network

Procedure 8.1. Before you upgrade

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`

    *   `/var/run/gluster/shared_storage/nfs-ganesha`

    If you use NFS-Ganesha, back up the following files from all nodes:

    *   `/etc/ganesha/exports/export.*.conf`

    *   `/etc/ganesha/ganesha.conf`

    *   `/etc/ganesha/ganesha-ha.conf`

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

    ```
    # umount mount-point
    ```

3. If you use NFS-Ganesha, run the following on a gluster server to disable the nfs-ganesha service:

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

4. On a gluster server, disable the shared volume.

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

5. Stop all volumes.

    ```
    # for vol in `gluster volume list`; do gluster --mode=script volume stop $vol;
    ```

6. Verify that all volumes are stopped.

    ```
    # gluster volume info
    ```

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

    ```
    # umount mount-point
    ```

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

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

> ## Important
>
> If glusterd crashes, there is no functionality impact to this crash as it occurs during the shutdown. For more information, see Resolving glusterd crash [https://access.redhat.com/documentation/en-us/red_hat_gluster_storage/3.4/html-single/administration_guide/#glusterd_crashing]

9. Stop the pcsd service.

    ```
    # systemctl stop pcsd
    ```

Procedure 8.2. Upgrade using yum

1. 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
    ```

2. 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-rp
        ```

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

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

    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]'
        ```

3. If you want to migrate from Gluster NFS to NFS Ganesha as part of this upgrade, perform the following additional steps.

    1. Stop and disable CTDB. This ensures that multiple versions of Samba do not run in the cluster during the update process, and avoids data corruption.

        ```
        # systemctl stop ctdb
        ```

```
# systemctl disable ctdb
```

2. Verify that the CTDB and NFS services are stopped:

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

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

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

4. Upgrade the server to Red Hat Gluster Storage 3.4.

```
# yum update
```

Wait for the update to complete.

5. Reboot the server to ensure that kernel updates are applied.

6. Ensure that glusterd and pcsd services are started.

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

7. 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 31305
```

> **Note**
>
> 31305 is the `cluster.op-version` value for Red Hat Gluster Storage 3.4 Batch 3 Update. Refer to the section called "Supported Versions of Red Hat Gluster Storage" for the correct `cluster.op-version` value for other versions.

8. If you want to migrate from Gluster NFS to NFS Ganesha as part of this upgrade, install the NFS-Ganesha packages as described in Chapter 4, *Deploying NFS-Ganesha on Red Hat Gluster Storage*, and use the information in the NFS Ganesha [https://access.redhat.com/documentation/en-us/red_hat_gluster_storage/3.4/html/administration_guide/nfs#nfs_ganesha] section of the *Red Hat Gluster Storage 3.4 Administration Guide* to configure the NFS Ganesha cluster.

9. Start all volumes.

```
# for vol in `gluster volume list`; do gluster --mode=script volume start $vol;
```

10. If you are using NFS-Ganesha:

   a. Copy the volume's export information from your backup copy of `ganesha.conf` to the new `/etc/ganesha/ganesha.conf` file.

   The export information in the backed up file is similar to the following:

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

   b. Copy the backup volume export files from the backup directory to `/etc/ganesha/exports` by running the following command from the backup directory:

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

11. Enable firewall settings for new services and ports. See the *Red Hat Gluster Storage 3.4 Administration Guide* for details: https://access.redhat.com/documentation/en-us/ red_hat_gluster_storage/3.4/html/administration_guide/chap-getting_started.

12. Enable the shared volume.

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

13. Ensure that the shared storage volume is mounted on the server. If the volume is not mounted, run the following command:

```
# mount -t glusterfs hostname:gluster_shared_storage /var/run/gluster/shared_st
```

14. Ensure that the /var/run/gluster/shared_storage/nfs-ganesha directory is created.

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

15. If you use NFS-Ganesha:

    a. Copy the ganesha.conf and ganesha-ha.conf files, and the /etc/ganesha/ exports directory to the /var/run/gluster/shared_storage/nfs-ganesha directory.

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

    b. Update the path of any export entries in the ganesha.conf file.

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

    c. Run the following to clean up any existing cluster configuration:

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

    d. If you have upgraded to Red Hat Enterprise Linux 7.4 or later, set the following SELinux Booleans:

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

16. Start the ctdb service (and nfs-ganesha service, if used) and verify that all nodes are functional.

```
# systemctl start ctdb
# gluster nfs-ganesha enable
```

17. If this deployment uses NFS-Ganesha, enable NFS-Ganesha on all volumes.

```
# gluster volume set volname ganesha.enable on
```

Upgrading to Red Hat Gluster Storage 3.4 for Systems Subscribed to Red Hat Network Satellite Server

Procedure 8.3. Before you upgrade

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`

   - `/var/run/gluster/shared_storage/nfs-ganesha`

   If you use NFS-Ganesha, back up the following files from all nodes:

   - `/etc/ganesha/exports/export.*.conf`

   - `/etc/ganesha/ganesha.conf`

   - `/etc/ganesha/ganesha-ha.conf`

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

   ```
   # umount mount-point
   ```

3. If you use NFS-Ganesha, run the following on a gluster server to disable the nfs-ganesha service:

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

4. On a gluster server, disable the shared volume.

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

5. Stop all volumes.

   ```
   # for vol in `gluster volume list`; do gluster --mode=script volume stop $vol;
   ```

6. Verify that all volumes are stopped.

   ```
   # gluster volume info
   ```

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

   ```
   # umount mount-point
   ```

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

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

Important

If `glusterd` crashes, there is no functionality impact to this crash as it occurs during the shutdown. For more information, see Resolving `glusterd` crash [https://access.redhat.com/documentation/en-us/red_hat_gluster_storage/3.4/html-single/administration_guide/#glusterd_crashing]

9. Stop the pcsd service.

```
# systemctl stop pcsd
```

Procedure 8.4. Upgrade using Satellite

1. Create an Activation Key at the Red Hat Network Satellite Server, and associate it with the following channels. For more information, see the section called "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)
   ```

   • 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)
   ```

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

   1. Log in to the Red Hat Network 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.

3. 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 3.4 channels on the Red Hat Network Satellite Server.

```
# rhnreg_ks --username username --password password --force --activationkey Act
```

4. 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
```

5. Upgrade to Red Hat Gluster Storage 3.4.

```
# yum update
```

6. Reboot the server and run volume and data integrity checks.

7. 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 31305
```

> **Note**
>
> 31305 is the `cluster.op-version` value for Red Hat Gluster Storage 3.4 Batch 3 Update. See the section called "Supported Versions of Red Hat Gluster Storage" for the correct `cluster.op-version` value for other versions.

8. Start all volumes.

```
# for vol in `gluster volume list`; do gluster --mode=script volume start $vol;
```

9. If you are using NFS-Ganesha:

   a. Copy the volume's export information from your backup copy of `ganesha.conf` to the new `/etc/ganesha/ganesha.conf` file.

   The export information in the backed up file is similar to the following:

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

b. Copy the backup volume export files from the backup directory to `/etc/ganesha/exports` by running the following command from the backup directory:

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

10. Enable firewall settings for new services and ports. See the *Red Hat Gluster Storage 3.4 Administration Guide* for details: https://access.redhat.com/documentation/en-us/red_hat_gluster_storage/3.4/html/administration_guide/chap-getting_started.

11. Enable the shared volume.

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

12. Ensure that the shared storage volume is mounted on the server. If the volume is not mounted, run the following command:

```
# mount -t glusterfs hostname:gluster_shared_storage /var/run/gluster/shared_st
```

13. Ensure that the `/var/run/gluster/shared_storage/nfs-ganesha` directory is created.

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

14. If you use NFS-Ganesha:

a. Copy the `ganesha.conf` and `ganesha-ha.conf` files, and the `/etc/ganesha/exports` directory to the `/var/run/gluster/shared_storage/nfs-ganesha` directory.

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

b. Update the path of any export entries in the `ganesha.conf` file.

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

c. Run the following to clean up any existing cluster configuration:

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

d. If you have upgraded to Red Hat Enterprise Linux 7.4 or later, set the following SELinux Booleans:

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

15. Start the ctdb service (and nfs-ganesha service, if used) and verify that all nodes are functional.

```
# systemctl start ctdb
# gluster nfs-ganesha enable
```

16. If this deployment uses NFS-Ganesha, enable NFS-Ganesha on all volumes.

```
# gluster volume set volname ganesha.enable on
```

In-Service Software Upgrade from Red Hat Gluster Storage 3.3 to Red Hat Gluster Storage 3.4

Important

Upgrade all Red Hat Gluster Storage servers before updating clients.

If you are upgrading from Red Hat Gluster Storage 3.1 Update 2 or earlier, upgrade servers and clients simultaneously.

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 continue even when the cluster is being updated.

I/O that uses CTDB may pause for the duration of an upgrade or update. This affects clients using Gluster NFS or Samba.

Note

After performing inservice upgrade of NFS-Ganesha, the new configuration file is saved as "`ganesha.conf.rpmnew`" in `/etc/ganesha` folder. The old configuration file is not overwritten during the inservice upgrade process. However, post upgradation, you have to manually copy any new configuration changes from "`ganesha.conf.rpmnew`" to the existing `ganesha.conf` file in `/etc/ganesha` folder.

Pre-upgrade Tasks

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

Upgrade Requirements for Red Hat Gluster Storage 3.4

The following are the upgrade requirements to upgrade to Red Hat Gluster Storage 3.4 from the preceding update:

- In-service software upgrade is supported only for nodes with replicate, distributed replicate, or erasure coded (dispersed) volumes.

- If you want to use snapshots for your existing environment, each brick must be an independent thin 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 3.4. For more information, see https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/Logical_Volume_Manager_Administration/lv_overview.html#linear_volumes

- 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 is not 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:

```
# subscription-manager repos --enable=repo-name
```

Restrictions for In-Service Software Upgrade

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

- In-service upgrade for NFS-Ganesha clusters is supported only from Red Hat Gluster Storage 3.3 to Red Hat Gluster Storage 3.4. If you are upgrading from Red Hat Gluster Storage 3.1 and you use NFS-Ganesha, use the offline upgrade method instead.

- Erasure coded (dispersed) volumes can be upgraded while in-service only if the `disperse.optimistic-change-log` and `disperse.eager-lock` options are set to `off`. Wait for at least two minutes after disabling these options before attempting to upgrade to ensure that these configuration changes take effect for I/O operations.

- 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 3.3 and Red Hat Gluster Storage 3.4 for a prolonged time.

- Do not combine different upgrade methods.

- It is not recommended to use in-service software upgrade for migrating to thin provisioned volumes, but to use offline upgrade scenario instead. For more information see, the section called "Offline Upgrade to Red Hat Gluster Storage 3.4"

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 3.4 using ISO can only be done from Red Hat Gluster Storage 3.3.1. For a complete list of supported Red Hat Gluster Storage releases, see the section called "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-3.4-RHEL-7-x86_64-dvd-1.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
   ```

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 t
  ---------  -----------    ---------   --------  ---------  ------  --------  --
  10.70.43.198       0       0Bytes       99        0          0    completed
  10.70.43.148      49     196Bytes      100        0          0    completed
  ```

- If you need to upgrade an erasure coded (dispersed) volume, set the disperse.optimistic-change-log and disperse.eager-lock options to off. Wait for at least two minutes after disabling these options before attempting to upgrade to ensure that these configuration changes take effect for I/O operations.

  ```
  # gluster volume set volname disperse.optimistic-change-log off
  # gluster volume set volname disperse.eager-lock off
  ```

- Ensure that there are no pending self-heals by 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
  ```

Service Impact of In-Service Upgrade

In-service software upgrade impacts 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
```

Gluster NFS.   When you use Gluster NFS to 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.

In-Service Software Upgrade

If you have a CTDB environment, see the section called "In-Service Software Upgrade for a CTDB Setup".

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`

   - `/var/run/gluster/shared_storage/nfs-ganesha`

   If you use NFS-Ganesha, back up the following files from all nodes:

   - `/etc/ganesha/exports/export.*.conf`

   - `/etc/ganesha/ganesha.conf`

   - `/etc/ganesha/ganesha-ha.conf`

2. If the node is part of an NFS-Ganesha cluster, place the node in standby mode.

   ```
   # pcs cluster standby
   ```

3. Ensure that there are no pending self-heal operations.

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

4. If this node is part of an NFS-Ganesha cluster:

   a. Disable the PCS cluster and verify that it has stopped.

   ```
   # pcs cluster disable
   # pcs status
   ```

   b. Stop the nfs-ganesha service.

   ```
   # systemctl stop nfs-ganesha
   ```

5. Stop all gluster services on the node and verify that they have stopped.

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

Important

If glusterd crashes, there is no functionality impact to this crash as it occurs during the shutdown. For more information, see Resolving glusterd crash [https://access.redhat.com/documentation/en-us/red_hat_gluster_storage/3.4/html-single/administration_guide/#glusterd_crashing]

6. 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
```

7. Update the server using the following command:

```
# yum update
```

8. If the volumes are thick provisioned, and you plan to use snapshots, perform the following steps to migrate to thin provisioned volumes:

Note

Migrating from thick provisioned volume to thin 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 the section called "Offline Upgrade to Red Hat Gluster Storage 3.4"

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

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

```
# umount mount_point
```

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

```
# lvremove logical_volume_name
```

For example:

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

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

```
# vgremove -ff volume_group_name
```

For example:

```
# vgremove -ff RHS_vg
```

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

```
# pvremove -ff physical_volume
```

e. 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 3.4 Administration Guide*: https://access.redhat.com/documentation/en-us/red_hat_gluster_Storage/3.4/html-single/administration_guide/#Formatting_and_Mounting_Bricks.

f. 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
```

g. Create a thinpool using the following command:

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

For example:

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

h. 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
```

i. 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 thin provisioned volume.

j. Mount the thin 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/glu
```

9. Disable glusterd.

```
# systemctl disable glusterd
```

This prevents it starting during boot time, so that you can ensure the node is healthy before it rejoins the cluster.

10. Reboot the server.

```
# shutdown -r now "Shutting down for upgrade to Red Hat Gluster Storage 3.4"
```

11. 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.

a. Create a FUSE mount point from any server to edit the extended attributes. Extended attributes cannot be edited using the NFS and CIFS mount points.

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

b. 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
```

c. 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
```

d. 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=0x756e636f6e66696e65645f753a6f626a6563745f723a66696c655f743a7
trusted.afr.r2-client-0=0x000000000000000000000000
trusted.afr.r2-client-1=0x000000000000000300000002
trusted.gfid=0x00000000000000000000000000000001
trusted.glusterfs.dht=0x0000000100000000000000007ffffffe
```

```
    trusted.glusterfs.volume-id=0xde822e25ebd049ea83bfaa3c4be2b440
```

12. Start the `glusterd` service.

```
# systemctl start glusterd
```

13. Verify that you have upgraded to the latest version of Red Hat Gluster Storage.

```
# gluster --version
```

14. Ensure that all bricks are online.

```
# 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
```

15. Start self-heal on the volume.

```
# gluster volume heal volname
```

16. Ensure that self-heal on the volume is complete.

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

The following example shows a completed self heal operation.

```
 # 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
```

17. Verify that shared storage is mounted.

```
# mount | grep /run/gluster/shared_storage
```

18. If this node is part of an NFS-Ganesha cluster:

   a. If the system is managed by SELinux, set the `ganesha_use_fusefs` Boolean to `on`.

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

   b. Start the NFS-Ganesha service.

   ```
   # systemctl start nfs-ganesha
   ```

   c. Enable and start the cluster.

   ```
   # pcs cluster enable
   # pcs cluster start
   ```

   d. Release the node from standby mode.

   ```
   # pcs cluster unstandby
   ```

   e. Verify that the pcs cluster is running, and that the volume is being exported correctly after upgrade.

   ```
   # pcs status
   # showmount -e
   ```

   NFS-ganesha enters a short grace period after performing these steps. I/O operations halt during this grace period. Wait until you see `NFS Server Now NOT IN GRACE` in the `ganesha.log` file before continuing.

19. Optionally, enable the glusterd service to start at boot time.

```
# systemctl enable glusterd
```

20. Repeat the above steps on the other node of the replica pair. In the case of a distributed-replicate setup, repeat the above steps on all the replica pairs.

21. 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 31305
```

Note

31305 is the `cluster.op-version` value for Red Hat Gluster Storage 3.4 Batch 3 Update. Refer to the section called "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 3.4 Administration Guide*: https://access.redhat.com/documentation/en-us/red_hat_gluster_Storage/3.4/html-single/administration_guide/#Troubleshooting1.

22. 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
```

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

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

24. 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_

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

25. If you disabled the `disperse.optimistic-change-log` and `disperse.eager-lock` options in order to upgrade an erasure-coded (dispersed) volume, re-enable these settings.

```
# gluster volume set volname disperse.optimistic-change-log on
# gluster volume set volname disperse.eager-lock on
```

Special Consideration for In-Service Software Upgrade

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

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:

```
# systemctl disable ctdb
```

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

```
# systemctl stop ctdb
```

a. 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:

```
# systemctl stop glusterd
```

```
# pkill glusterfs
# pkill glusterfsd
```

Important

If glusterd crashes, there is no functionality impact to this crash as it occurs during the shutdown. For more information, see Resolving glusterd crash [https://access.redhat.com/documentation/en-us/red_hat_gluster_storage/3.4/html-single/administration_guide/#glusterd_crashing]

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 11, *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:

   ```
   # systemctl disable glusterd
   ```

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/s
    ```

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:

    ```
    # systemctl enable glusterd
    ```

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

    ```
    # systemctl enable ctdb
    ```

14. Start the glusterd service using the following command:

    ```
    # systemctl start glusterd
    ```

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:

```
# systemctl start ctdb
```

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 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:

```
# systemctl stop ctdb
```

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

- Remove the following soft links:

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

- 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*

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
```

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. If you are upgrading from Red Hat Gluster Storage 3.1 Update 2 or earlier, you must upgrade servers and clients simultaneously. For more information regarding upgrading native client, refer to section Upgrading Native Client [https://access.redhat.com/documentation/en-us/red_hat_gluster_storage/3.4/html-single/administration_guide/#Upgrading_Native_Client] in the *Red Hat Gluster Storage Administration Guide*.

Chapter 9. Upgrading Red Hat Gluster Storage to Red Hat Enterprise Linux 7

This chapter describes the procedure to upgrade from Red Hat Enterprise Linux 6 based Red Hat Gluster Storage to Red Hat Enterprise Linux 7 based Red Hat Gluster Storage.

Important

Be aware of the following when upgrading a Gluster node from RHEL 6 to RHEL 7:

* Red Hat recommends that the RHEL 6 to RHEL 7 upgrade be performed with assistance from Red Hat.

* Upgrading from RHEL 6 to RHEL 7 is an offline procedure which results in Gluster node downtime, as volumes are offline during the upgrade process.

* Servers must be upgraded prior to upgrading clients.

Preparing System for Upgrade

1. **Migrate from Red Hat Network Classic to Red Hat Subscription Manager**

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

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

   If the system is not on Red Hat Network Classic, skip this step

   If the system is on Red Hat Network Classic, migrate to Red Hat Subscription Manager using: Migrating from RHN to RHSM in Red Hat Enterprise Linux [https://access.redhat.com/solutions/129723].

2. **Register the system with Subscription Manager**

   ### Important

   If the system is already registered using Subscription Manager, skip the registrations steps and start the backup of the Gluster configuration files.

   To register the system with Red Hat Network, execute the following command and enter your Red Hat Network user name and password that have the RHEL entitlements:

   ```
   # subscription-manager register --username=user_name --password=password
   ```

3. **Identify the available entitlement pools**

   Find the entitlement pools containing the RHEL 6 repositories:

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

4. **Attach entitlement pool to the system**

   Use the pool identifier located in the previous step to attach the RHEL 6 entitlements to the system:

```
# subscription-manager attach --pool=pool_ID
```

5. **Enable repositories**

   Enable the RHEL 6, scalefs, and Red Hat Gluster Storage repositories:

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

6. **Backup the Gluster configuration files**

   > **Note**
   >
   > It is recommended to make a complete backup using a reliable backup solution before you update your system. This Knowledge Base solution covers one possible approach: https://access.redhat.com/solutions/1484053.

   1. Ensure that the following configuration directories and files are backed up:

      - `/var/lib/glusterd`

      - `/etc/glusterfs`

   2. For systems with samba-ctdb enabled cluster, create a new directory to store the backup:

      ```
      # mkdir backup_folder_name
      # cd backup_folder_name
      ```

   3. Execute the following command to take a backup samba-ctdb data:

      ```
      for each in `ctdb getdbmap | grep PERSISTENT | cut -d" " -f2 | cut -d":" -f2`
      ```

7. **Stop all Gluster services, volumes, and processes**

   1. Stop any geo-replication sessions:

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

   2. Stop the Nagios process:

      ```
      # service nrpe stop
      ```

   3. Stop all volumes:

      ```
      # for vol in `gluster volume list`; do gluster --mode=script volume stop $vol
      ```

   4. Stop the Gluster processes:

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

      > **Important**
      >
      > If `glusterd` crashes, there is no functionality impact to this crash as it occurs during the shutdown. For more information, see Resolving `glusterd`

> crash [https://access.redhat.com/documentation/en-us/red_hat_gluster_storage/3.4/html-single/administration_guide/#glusterd_crashing]

5. For system with samba-ctdb enabled cluster, stop the CTDB services:

```
# service ctdb stop
```

8. **Update the system and reboot**

   Update the system to the latest minor version of RHEL 6, and reboot after the update is complete:

```
# yum update
# reboot
```

9. **Verify the version number**

   Check the current version number of the updated RHEL 6 system using the following command:

```
# cat /etc/redhat-release
```

Important

The version number should be 6.10.

Performing System Assessment

1. **Install preupgrade-assistant**

   The preupgrade-assistant is a tool that scans through the existing RHEL 6 system and finds out how fit and ready a system is, to be upgraded to RHEL 7.

   Subscribe to the required repositories and install the preupgrade-assistant tool using the following commands:

```
# subscription-manager repos --enable rhel-6-server-extras-rpms --enable rhel-6
# yum install preupgrade-assistant preupgrade-assistant-el6toel7
```

2. **Run the preupgrade tool**

   Launch the preupgrade tool:

```
# preupg -v
```

3. **Assess the results**

   1. View the following result of the preupgrade tool in a browser:

```
/root/preupgrade/result.html
```

   2. Make a note of all the components that are marked as Failed and Needs Action.

   3. Share the findings with the Red Hat executive assisting with the upgrade process.

4. **SELinux policy in RHEL 7**

   SELinux policies have changed between RHEL 6 and RHEL 7. For a working SELinux policy in RHEL 7, execute the following command:

```
# semodule -r sandbox
```

5. **Uninstall Cluster and HA related packages**

> ## Important
>
> Migration is not possible with Cluster and HA related packages installed on the system.

Uninstall Cluster and HA related packages by executing the following command:

```
# yum remove modcluster ricci openais corosync
```

Upgrading from RHEL 6.X to RHEL 7.X

1. **Install migration tool**

Install the tool to perform the migration from RHEL 6 to RHEL 7:

```
# yum install redhat-upgrade-tool
# yum install yum-utils
```

2. **Disable all repositories**

Disable all the enabled repositories:

```
# yum-config-manager --disable \*
```

3. **Download latest RHEL 7 ISO**

Visit the Software & Download Center in the Red Hat Customer Service Portal (https://access.redhat.com/downloads) to obtain the latest RHEL 7 *ISO image file*. 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 RHEL 7 installation files using a Red Hat Subscription or a Red Hat Evaluation Subscription, perform the following steps:

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 Enterprise Linux area, click Download Software to download the latest version of the software.

4. **Upgrade to RHEL 7 using ISO**

Upgrade to RHEL 7 using the Red Hat upgrade tool and reboot after the upgrade process is completed:

```
# redhat-upgrade-tool --iso RHEL7_ISO_filepath --cleanup-post
# reboot
```

> ## Important
>
> The upgrade process is time-cosuming depending on your system's configuration and amount of data.

Upgrading to Red Hat Gluster Storage 3.4

1. **Disable all repositories**

   ```
   # subscription-manager repos --disable='*'
   ```

2. **Subscribe to RHEL 7 channel**

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

3. **Check for stale RHEL 6 packages**

   Make a note of any stale RHEL 6 packages post upgrade:

   ```
   # rpm -qa | grep el6
   ```

4. **Update and reboot**

   Update the RHEL 7 packages and reboot once the update is complete.

   ```
   # yum update
   # reboot
   ```

5. **Verify the version number**

   Check the current version number of the updated RHEL 7 system:

   ```
   # cat /etc/redhat-release
   ```

   ## Important

   The version number should be 7.5.

6. **Subscribe to required channels**

   1. Subscibe to the Gluster channel:

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

   2. If you require Samba, enable its repository:

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

   3. If you require NFS-Ganesha, enable its repository:

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

   4. If you require gdeploy, enable the Ansible repository:

      ```
      # subscription-manager repos --enable=rhel-7-server-ansible-2-rpms
      ```

   5. If you require Nagios, enable its repository:

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

7. **Install and update Gluster**

   1. Install Red Hat Gluster Storage 3.4 using the following command:

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

2. Update Red Hat Gluster Storage to the latest packages using the following command:

```
# yum update
```

8. **Verify the installation and update**

   1. Check the current version number of the updated Red Hat Gluster Storage system:

   ```
   # cat /etc/redhat-storage-release
   ```

   ## Important

   The version number should be 3.4.

   2. Check if any RHEL 6 packages are present:

   ```
   # rpm -qa | grep el6
   ```

   ## Important

   The output of the command should not list any packages of RHEL 6 variant. If the output lists packages of RHEL 6 variant, contact Red Hat Support for further course of action on these packages.

9. **Firewalld installation and configuration**

   1. Install and start the firewall deamon using the following commands:

   ```
   # yum install firewalld
   # systemctl start firewalld
   ```

   2. Add the Gluster process to firewall:

   ```
   # firewall-cmd --zone=public --add-service=glusterfs --permanent
   ```

   3. Add the required services and ports to firewalld, see Considerations for Red Hat Gluster Storage [https://access.redhat.com/documentation/en-us/red_hat_gluster_storage/3.4/html-single/administration_guide/chap-red_hat_storage_volumes#chap-Getting_Started]

   4. Reload the firewall using the following commands:

   ```
   # firewall-cmd --reload
   ```

10. **Start Gluster processes**

    1. Start the **glusterd** process:

    ```
    # systemctl start glusterd
    ```

    2. For a system with Nagios, start and enable the following process:

    ```
    # systemctl start glusterpmd
    ```

    3. Start and enable the Nagios process: 71

```
# systemctl start nrpe
```

4. If the Nagios process fails to start, execute the following commands:

```
# restorecon -Rv /etc/nagios/nrpe.cfg
# systemctl start nrpe
```

11. **Update Gluster op-version**

Update the Gluster op-version to the required maximum version using the following commands:

```
# gluster volume get all cluster.max-op-version
# gluster volume set all cluster.op-version op_version
```

> # Note
>
> 31305 is the `cluster.op-version` value for Red Hat Gluster Storage 3.4 Batch 3
> Update. Refer to the section called "Supported Versions of Red Hat Gluster Storage" for the
> correct `cluster.op-version` value for other versions.

12. **Setup Samba and CTDB**

If the Gluster setup on RHEL 6 had Samba and CTDB configured, you should have the following
available on the updated RHEL 7 system:

- CTDB volume

- `/etc/ctdb/nodes` file

- `/etc/ctdb/public_addresses` file

Perform the following steps to reconfigure Samba and CTDB:

1. Configure the firewall for Samba:

```
# firewall-cmd --zone=public  --add-service=samba --permanent
# firewall-cmd --zone=public  --add-port=4379/tcp --permanent
```

2. Subscribe to the Samba channel:

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

3. Update Samba to the latest packages:

```
# yum update
```

4. Setup CTDB for Samba, see *Configuring CTDB on Red Hat Gluster Storage Server* in Setting up
   CTDB for Samba [https://access.redhat.com/documentation/en-us/red_hat_gluster_storage/3.4/
   html-single/administration_guide/chap-red_hat_storage_volumes#sect-SMB_CTDB]. You must
   skip creating the volume as the volumes present before the upgrade would be persistent after the
   upgrade.

5. In the following files, replace `all` in the statement `META="all"` with the volume name:

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

For example, the volume name is ctdb_volname, the META="all" in the files should be changed to META="ctdb_volname".

6. Restart the CTDB volume using the following commands:

```
# gluster volume stop volume_name
# gluster volume start volume_name
```

7. Start the CTDB process:

```
# systemctl start ctdb
```

8. Share the volume over Samba as required, see Sharing Volumes over SMB [https://access.redhat.com/documentation/en-us/red_hat_gluster_storage/3.4/html-single/administration_guide/chap-red_hat_storage_volumes#sect-Sharing_Volumes_over_SMB].

13. **Start volumes and geo-replication**

1. Start the required volumes using the following command:

```
# gluster volume start volume_name
```

2. 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/   gluster
```

3. Restore the geo-replication session:

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

For more information on geo-replication, see Preparing to Deploy Geo-replication [https://access.redhat.com/documentation/en-us/red_hat_gluster_storage/3.4/html-single/administration_guide/chap-red_hat_storage_volumes#sect-Preparing_to_Deploy_Geo-replication].

Chapter 10. Upgrading from Red Hat Gluster Storage 3.3 to Red Hat Gluster Storage 3.4 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 3.3 to Red Hat Gluster Storage 3.4 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 [https://bugzilla.redhat.com/show_bug.cgi?id=1115367] and BZ#1118762 [https://bugzilla.redhat.com/show_bug.cgi?id=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.

Prerequisites

- Verify that no self-heal operations are in progress.

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

Upgrading an in-service integrated environment

Follow the instructions below to upgrade an integrated Red Hat Virtualization and Red Hat Gluster Storage environment that is online and in-service.

1. **Upgrade the Red Hat Gluster Storage nodes**

   Perform the following steps for each Red Hat Gluster Storage node, one node at a time. If you have multiple replica sets, upgrade each node in a replica set before moving on to another replica set.

a. Stop all gluster services on the storage node that you want to upgrade by running the following commands as the root user:

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

Important

If `glusterd` crashes, there is no functionality impact to this crash as it occurs during the shutdown. For more information, see Resolving `glusterd` crash [https://access.redhat.com/documentation/en-us/red_hat_gluster_storage/3.4/html-single/administration_guide/#glusterd_crashing]

b. In the Administration Portal, click Storage $\rightarrow$ Hosts and select the storage node to upgrade.

c. Click Management $\rightarrow$ Maintenance and click OK.

d. Upgrade the storage node using the procedure at the section called "In-Service Software Upgrade from Red Hat Gluster Storage 3.3 to Red Hat Gluster Storage 3.4".

e. Reboot the storage node.

f. In the Administration Portal, click Storage $\rightarrow$ Hosts and select the storage node.

g. Click Management $\rightarrow$ Activate.

h. Click the name of the storage node $\rightarrow$ Bricks, and verify that the Self-Heal Info column of all bricks is listed as OK before upgrading the next storage node.

2. **Upgrade Gluster client software on virtualization hosts**

After upgrading all storage nodes, update the client software on all virtualization hosts.

a. **For Red Hat Virtualization hosts**

i. In the Administration Portal, click Compute $\rightarrow$ Hosts and select the node to update.

ii. On Red Hat Virtualization versions earlier than 4.2, click Management $\rightarrow$ Maintenance and click OK.

iii. Click Installation $\rightarrow$ Upgrade and click OK.

The node is automatically updated with the latest packages and reactivated when the update is complete.

b. **For Red Hat Enterprise Linux hosts**

i. In the Administration Portal, click Compute $\rightarrow$ Hosts and select the node to update.

ii. Click Management $\rightarrow$ Maintenance and click OK.

iii. Subscribe the host to the correct repositories to receive updates.

iv. Update the host.

Upgrading from Red Hat
Gluster Storage 3.3 to Red Hat
Gluster Storage 3.4 in a Red Hat
yum update Enterprise Virtualization-Red
Hat Gluster Storage Environment

v.   In the Administration Portal, click Management $\rightarrow$ Activate.

Upgrading an offline integrated environment

Follow the instructions below to upgrade an integrated Red Hat Virtualization and Red Hat Gluster Storage environment that is offline and out-of-service.

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 the section called "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 3.4 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 10.1, "Red Hat Gluster Storage Node before Upgrade ", and perform bootstrapping.

     Figure 10.1. 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	Upgrading from Red Hat Gluster Storage 3.3 to Red Hat Gluster Storage 3.4 in a Red Hat Enterprise Virtualization-Red Hat Gluster Storage Environment

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 10.2. Red Hat Gluster Storage Node after Upgrade

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

Upgrading using yum

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

2. Using Red Hat Enterprise Virtualization Manager, move the data domain backed by gluster volume to *Maintenance* mode.

3. Using Red Hat Enterprise Virtualization Manager, move all Red Hat Gluster Storage nodes to *Maintenance* mode.

4. Perform **yum** update as mentioned in the section called "Upgrading to Red Hat Gluster Storage 3.4 for Systems Subscribed to Red Hat Network".

5. 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.

6. Using Red Hat Enterprise Virtualization Manager, start the volume.

7. 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 11. 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 12. 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 COMMAND
```

The following is an example using the **peer 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 **peer status** to view the status of the peer server:

1. Start Gluster CLI's interactive mode:

   ```
   # gluster
   ```

2. Request the peer server status:

   ```
   gluster> peer 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 History

Revision 3.4-3	Mon Feb 04 2019	Red Hat Gluster StorageDocumentation Team<gluster-docs@redhat.com>

BZ#1659944 - Documented new default value for op-version.
BZ#1644480 - Added warnings about glusterd crash.

Revision 3.4-2	Wed Oct 31 2018	Red Hat Gluster StorageDocumentation Team<gluster-docs@redhat.com>

Content updates for Red Hat Gluster Storage 3.4 Batch 1 Update.

Revision 3.4-1	Tue Sep 04 2018	Red Hat Gluster StorageDocumentation Team<gluster-docs@redhat.com>

Updates for Red Hat Gluster Storage 3.4 release.