



Red Hat Gluster Storage 3

Installation Guide

Installing Red Hat Storage 3

Red Hat Gluster Storage 3 Installation Guide

Installing Red Hat Storage 3

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

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

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

Legal Notice

Copyright © 2014-2015 Red Hat, Inc.

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

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

Abstract

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

Table of Contents

CHAPTER 1. INTRODUCTION	4
CHAPTER 2. OBTAINING RED HAT STORAGE	5
2.1. OBTAINING RED HAT STORAGE SERVER FOR ON-PREMISE	5
2.2. OBTAINING RED HAT STORAGE SERVER FOR PUBLIC CLOUD	5
CHAPTER 3. PLANNING RED HAT STORAGE INSTALLATION	6
3.1. PREREQUISITES	6
3.2. HARDWARE COMPATIBILITY	7
3.3. PORT INFORMATION	7
CHAPTER 4. INSTALLING RED HAT STORAGE	10
4.1. INSTALLING FROM AN ISO IMAGE	10
4.2. INSTALLING RED HAT STORAGE SERVER ON RED HAT ENTERPRISE LINUX (LAYERED INSTALL)	21
4.3. INSTALLING FROM A PXE SERVER	23
4.4. INSTALLING FROM RED HAT SATELLITE SERVER	24
CHAPTER 5. SUBSCRIBING TO THE RED HAT STORAGE SERVER CHANNELS	27
CHAPTER 6. UPGRADING RED HAT STORAGE	29
6.1. UPGRADING FROM RED HAT STORAGE 2.1 TO RED HAT STORAGE 3.0 USING AN ISO	29
6.2. UPGRADING FROM RED HAT STORAGE 2.1 TO RED HAT STORAGE 3.0 FOR SYSTEMS SUBSCRIBED TO RED HAT NETWORK	37
6.3. UPGRADING FROM RED HAT STORAGE 2.1 TO RED HAT STORAGE 3.0 FOR SYSTEMS SUBSCRIBED TO RED HAT SATELLITE SERVER	39
6.4. UPGRADING FROM RED HAT STORAGE 2.1 TO RED HAT STORAGE 3.0 IN A RED HAT ENTERPRISE VIRTUALIZATION-RED HAT STORAGE ENVIRONMENT	41
CHAPTER 7. DEPLOYING SAMBA ON RED HAT STORAGE	45
7.1. PREREQUISITES	45
7.2. INSTALLING SAMBA USING ISO	45
7.3. INSTALLING SAMBA USING YUM	46
CHAPTER 8. DEPLOYING THE HORTONWORKS DATA PLATFORM 2.1 ON RED HAT STORAGE	47
8.1. PREREQUISITES	47
8.2. INSTALLING THE HADOOP FILESYSTEM PLUGIN FOR RED HAT STORAGE	52
8.3. ADDING AND REMOVING USERS	59
8.4. DISABLING A VOLUME FOR USE WITH HADOOP	59
8.5. VERIFYING THE CONFIGURATION	60
8.6. TROUBLESHOOTING	61
CHAPTER 9. SETTING UP SOFTWARE UPDATES	63
9.1. UPDATING RED HAT STORAGE IN THE OFFLINE MODE	63
9.2. IN-SERVICE SOFTWARE UPGRADE TO UPGRADE FROM RED HAT STORAGE 2.1 TO RED HAT STORAGE 3.0	64
9.3. UPDATING YOUR CURRENT SYSTEM	78
CHAPTER 10. MANAGING THE GLUSTERD SERVICE	81
10.1. MANUALLY STARTING AND STOPPING GLUSTERD	81
CHAPTER 11. USING THE GLUSTER COMMAND LINE INTERFACE	82
PART I. APPENDIX	84

CHAPTER 12. DEPLOYING THE HORTONWORKS DATA PLATFORM 2.0.6 ON RED HAT STORAGE	...	85
12.1. PREREQUISITES		85
12.2. INSTALLING THE HADOOP FILESYSTEM PLUGIN FOR RED HAT STORAGE		89
12.3. ADDING AND REMOVING USERS		96
12.4. DISABLING A VOLUME FOR USE WITH HADOOP		96
12.5. VERIFYING THE CONFIGURATION		97
12.6. TROUBLESHOOTING		97
APPENDIX A. REVISION HISTORY	98

CHAPTER 1. INTRODUCTION

Red Hat Storage is software-only, scale-out storage that provides flexible and affordable unstructured data storage for the enterprise. Red Hat Storage 3.0 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 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 Storage can be deployed in the private cloud or datacenter using Red Hat Storage Server for On-Premise. Red Hat Storage can be installed on commodity servers and storage hardware resulting in a powerful, massively scalable, and highly available NAS environment. Additionally, Red Hat Storage can be deployed in the public cloud using Red Hat 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 Storage Server for On-Premise

Red Hat 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 Storage Server for Public Cloud

Red Hat 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.

CHAPTER 2. OBTAINING RED HAT STORAGE

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

2.1. OBTAINING RED HAT 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 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 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 Storage Server area, click **Download Software** to download the latest version of the software.

2.2. OBTAINING RED HAT STORAGE SERVER FOR PUBLIC CLOUD

Red Hat 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>.

CHAPTER 3. PLANNING RED HAT STORAGE INSTALLATION

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

3.1. 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 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 Storage servers using the Network Time Protocol (NTP) daemon.

3.1.1. Network Time Protocol Setup

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

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

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

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

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

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

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

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

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

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

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

```
# service ntpd restart
```

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

```
# chkconfig ntpd on
```

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

3.2. 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 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.

3.3. PORT INFORMATION

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

Table 3.1. TCP Port Numbers

Port Number	Usage
22	For sshd used by geo-replication.
111	For rpc port mapper.
139	For netbios service.
445	For CIFS protocol.
965	For NFS's Lock Manager (NLM).

Port Number	Usage
2049	For glusterFS's NFS exports (nfsd process).
24007	For glusterd (for management).
24009 - 24108	For client communication with Red Hat Storage 2.0.
38465	For NFS mount protocol.
38466	For NFS mount protocol.
38468	For NFS's Lock Manager (NLM).
38469	For NFS's ACL support.
39543	For oVirt (Red Hat Storage-Console).
49152 - 49251	For client communication with Red Hat Storage 2.1 and for brick processes depending on the availability of the ports. The total number of ports required to be open depends on the total number of bricks exported on the machine.
55863	For oVirt (Red Hat Storage-Console).

Table 3.2. 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 3.3. TCP Port Numbers for Nagios Monitoring

Port Number	Usage
80	For HTTP protocol (required only if Nagios server is running on a Red Hat Storage node).

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

Table 3.4. UDP Port Numbers


Port Number	Usage
111	For RPC Bind.
963	For NFS's Lock Manager (NLN).

CHAPTER 4. INSTALLING RED HAT STORAGE

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

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

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



IMPORTANT

- Technology preview packages will also be installed with this installation of Red Hat Storage Server. For more information about the list of technology preview features, see *Chapter 4. Technology Previews* in the *Red Hat Storage 3.0 Release Notes*.
- While cloning a Red Hat Storage Server installed on a virtual machine, the `/var/lib/glusterd/glusterd.info` file will be cloned to the other virtual machines, hence causing all the cloned virtual machines to have the same UUID. Ensure to remove the `/var/lib/glusterd/glusterd.info` file before the virtual machine is cloned. The file will be automatically created with a UUID on initial start-up of the glusterd daemon on the cloned virtual machines.

4.1. INSTALLING FROM AN ISO IMAGE

To install Red Hat Storage Server from the ISO image:


1. Download an ISO image file for Red Hat Storage Server as described in [Chapter 2, Obtaining Red Hat Storage](#).

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



Figure 4.1. Installation

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




NOTE

For some hypervisors, while installing Red Hat 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 Storage does not support clients over IPv6.




Figure 4.2. Configure TCP/IP

3. The **Welcome** screen displays.




Figure 4.3. Welcome

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**.

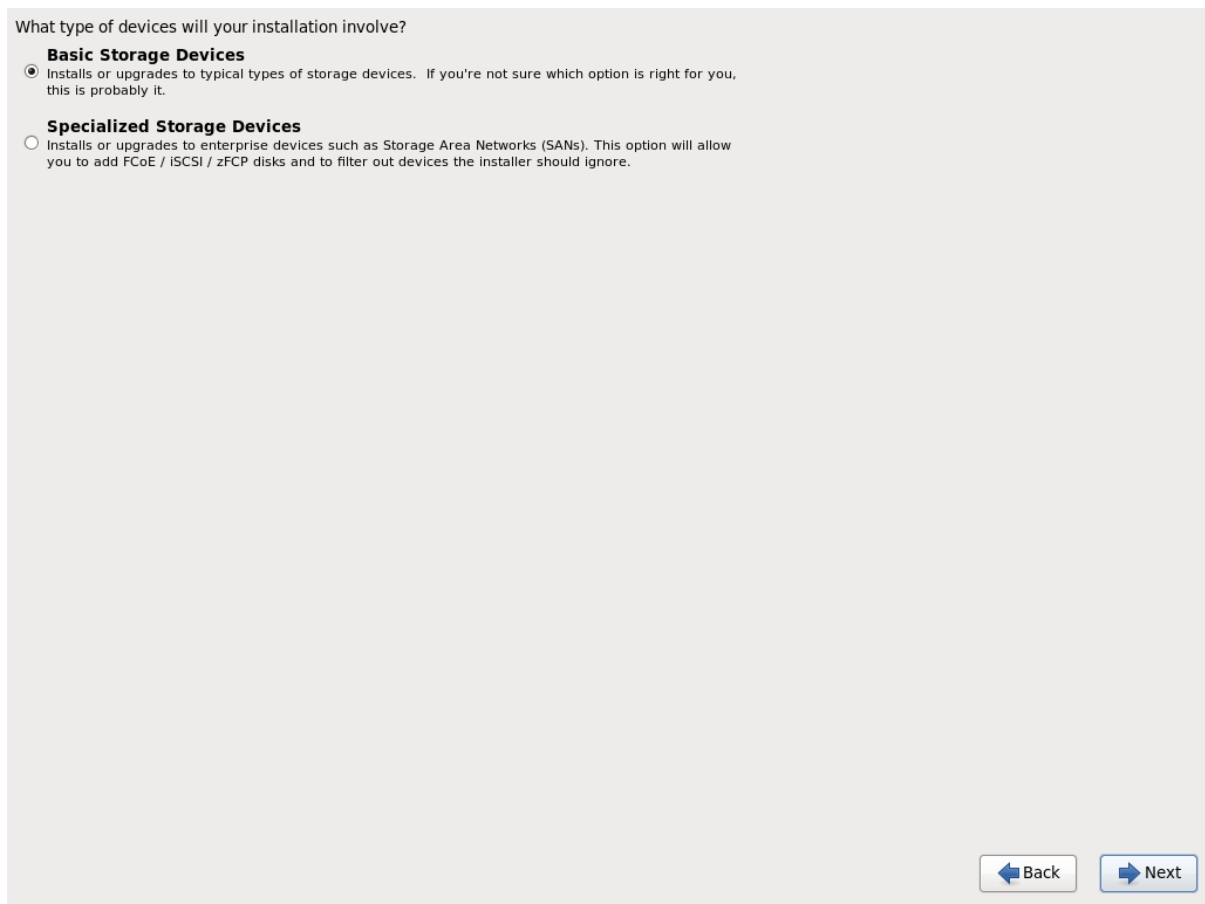


Figure 4.4. Storage Devices

Click **Next**.

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



Figure 4.5. Hostname

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

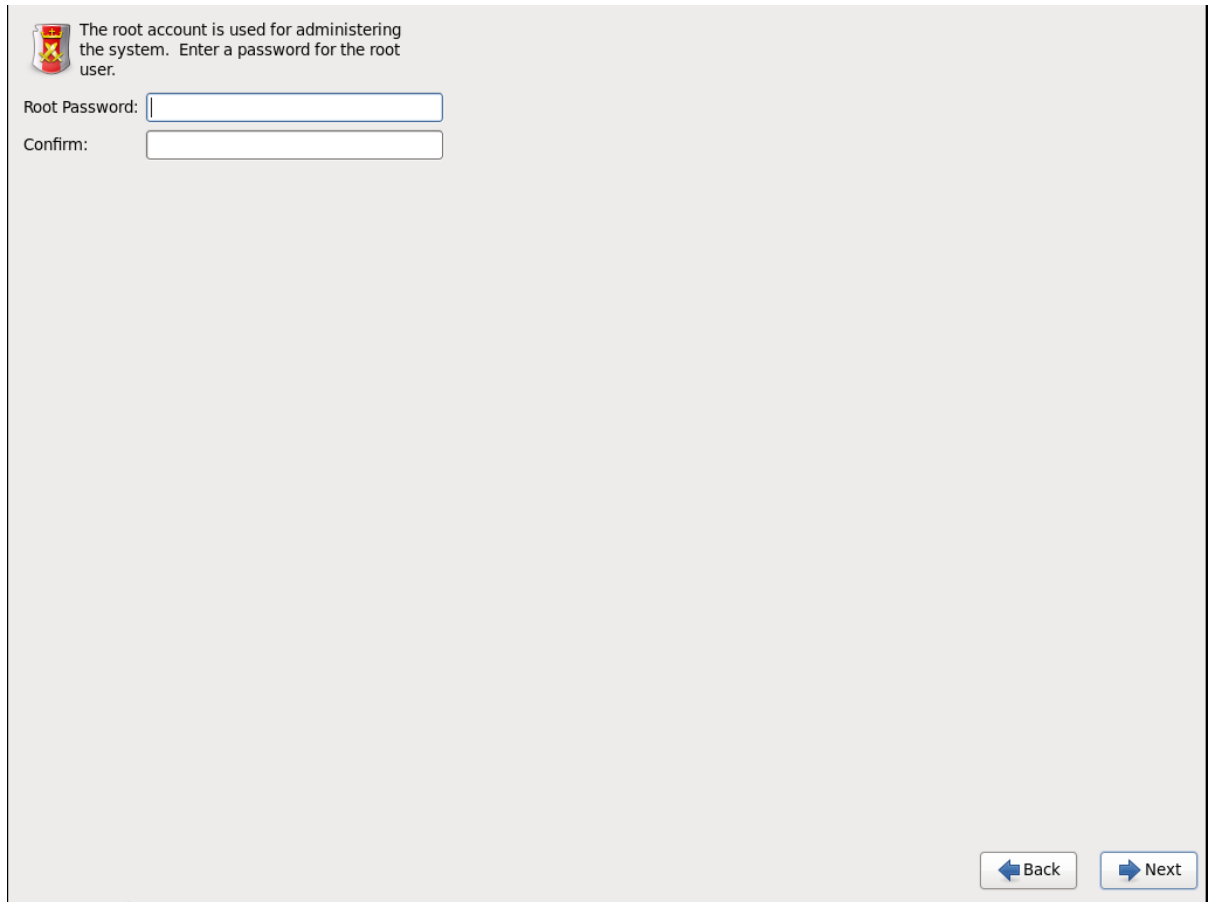
8. The **Time Zone Configuration** screen displays. Set your time zone by selecting the city closest to your computer's physical location.
9. The **Set Root Password** screen displays.

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 root account is used for administering the system. Enter a password for the root user.

Root Password:

Confirm:

Back Next

Figure 4.6. Set Root Password

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

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

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

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

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

If you do not feel 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:

- o Use All Space
- o Replace Existing Linux System(s)
- o Shrink Current System
- o 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.

Which type of installation would you like?

- Use All Space**
Removes all partitions on the selected device(s). This includes partitions created by other operating systems.
Tip: This option will remove data from the selected device(s). Make sure you have backups.
- Replace Existing Linux System(s)**
Removes only Linux partitions (created from a previous Linux installation). This does not remove other partitions you may have on your storage device(s) (such as VFAT or FAT32).
Tip: This option will remove data from the selected device(s). Make sure you have backups.
- Shrink Current System**
Shrinks existing partitions to create free space for the default layout.
- Use Free Space**
Retains your current data and partitions and uses only the unpartitioned space on the selected device(s), assuming you have enough free space available.
- Create Custom Layout**
Manually create your own custom layout on the selected device(s) using our partitioning tool.

Encrypt system
 Review and modify partitioning layout

Figure 4.7. Partitioning Type

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*.



NOTE

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.

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

Click **Next**.

- The **Minimal Selection** screen displays.

The default installation of Red Hat Storage Server includes a set of software applicable for general internet usage. You can optionally select a different set of software now.

Minimal

Please select any additional repositories that you want to use for software installation.

Installation Repo

You can further customize the software selection now, or after install via the software management application.

Customize later Customize now

Figure 4.8. Minimal Selection

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.

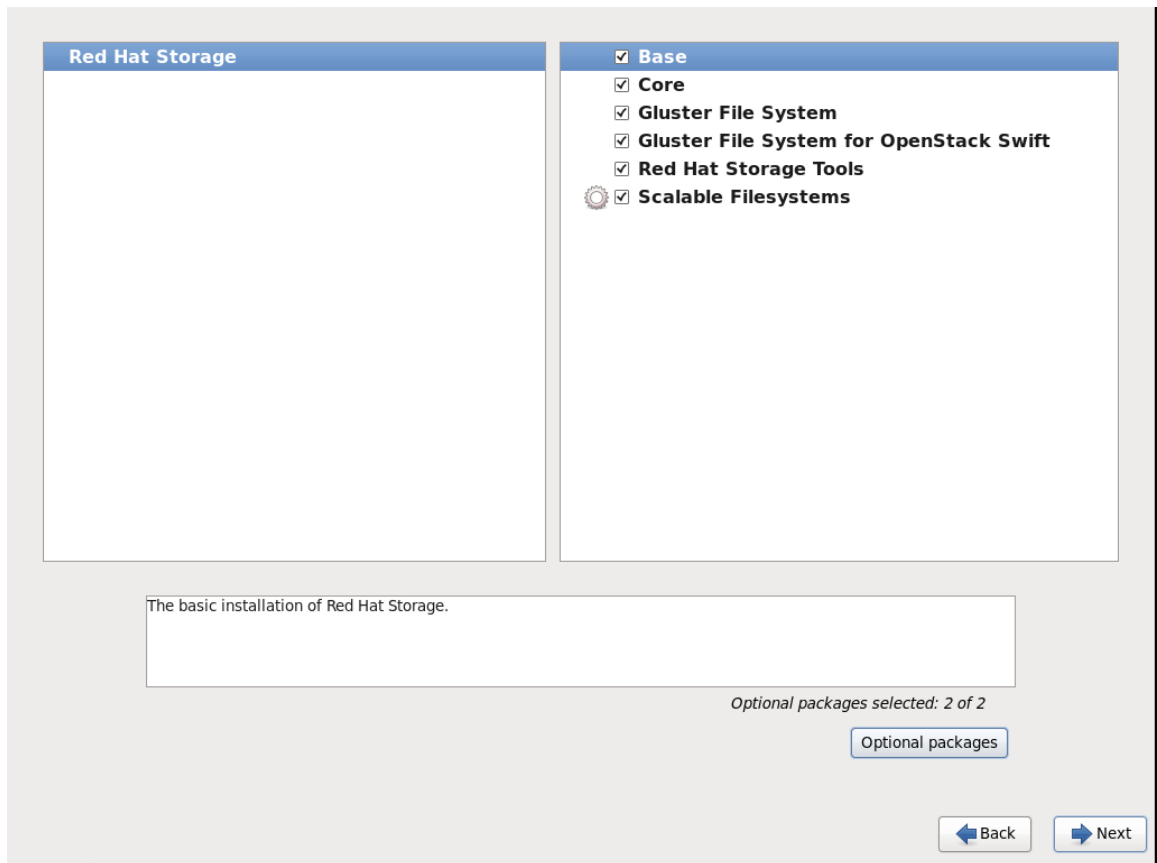


Figure 4.9. Customize Packages

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

- For Red Hat Storage 3.0.4 or later, if you require the Samba packages, ensure you select the **Samba (SMB) server for gluster** component, in the **Customizing the Software Selection** screen. If you require samba active directory integration with gluster, ensure you select **Active Directory Integration** component.

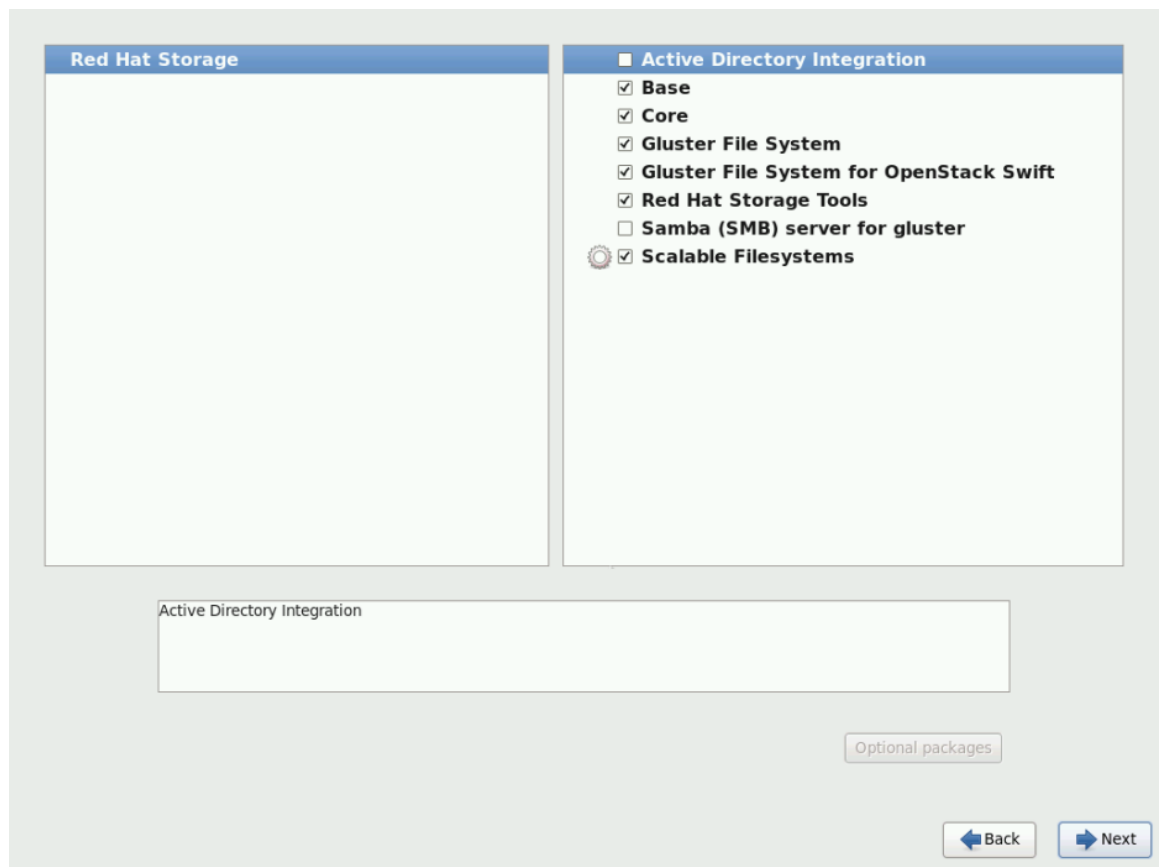


Figure 4.10. Customize Packages

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

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

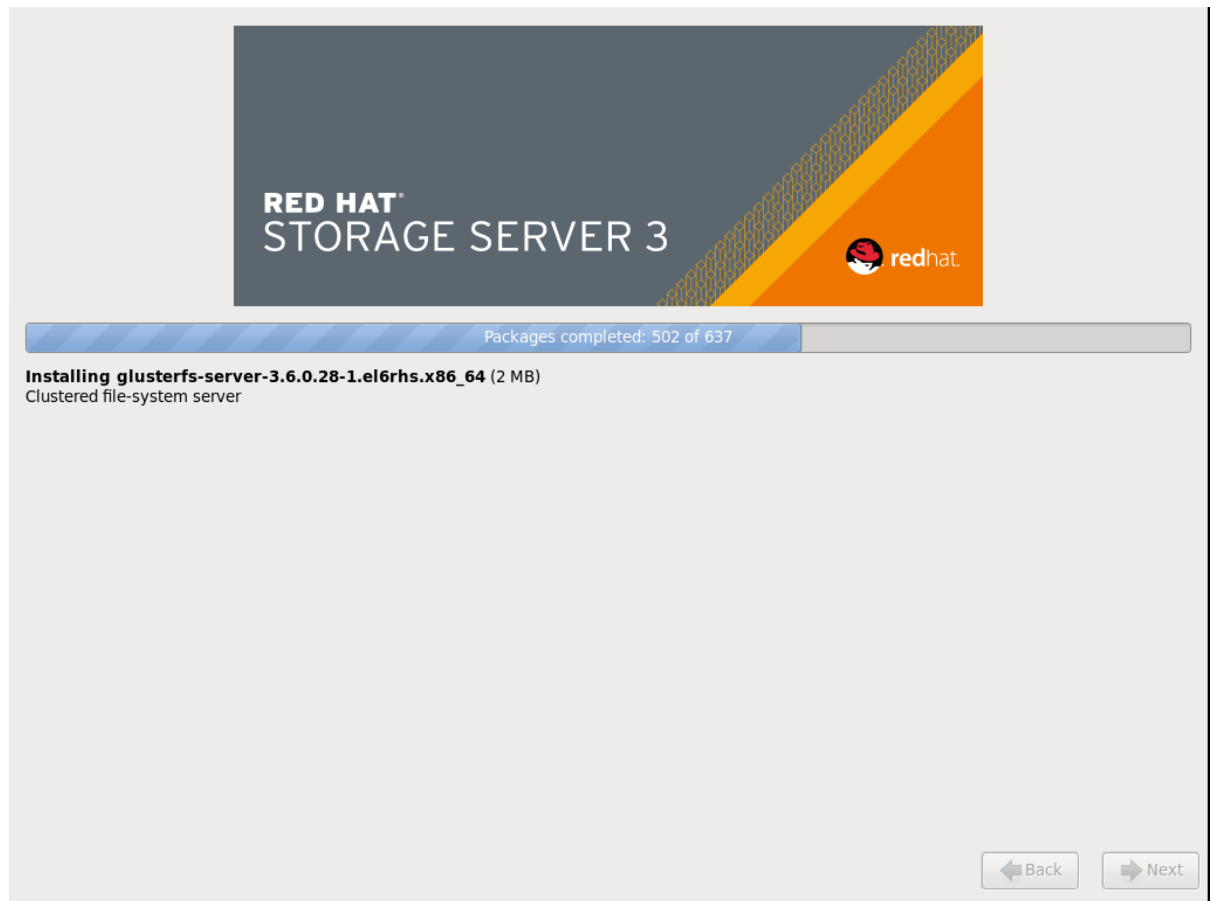


Figure 4.11. Package Installation

14. On successful completion, the **Installation Complete** screen displays.

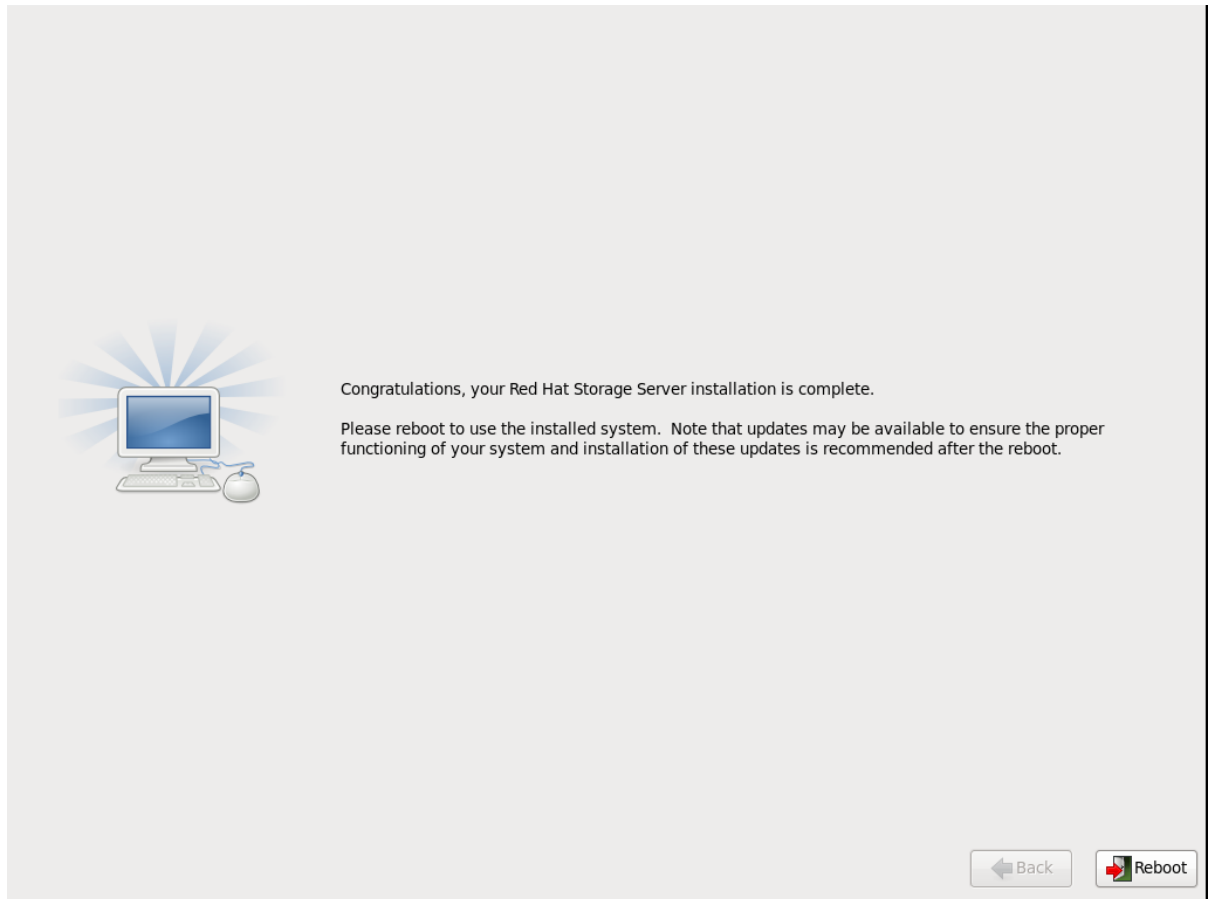


Figure 4.12. Installation Complete

15. Click **Reboot** to reboot the system and complete the installation of Red Hat Storage Server.

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

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

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

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



IMPORTANT

Layered Installation is not available to all customers. Contact your Red Hat representative for more details on whether you can use it.

1. **Perform a base install of Red Hat Enterprise Linux Server version 6.5 or 6.6.**
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 channels required to install Red Hat Storage:

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

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 Storage** entitlements to the system. Run the following command to attach the entitlements:

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

5. Enable the Required Channels

Run the following commands to enable the channels required to install Red Hat 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
```

1. For Red Hat Storage 3.0.4 and later, if you require Samba, then enable the following channel:

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

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 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. Install the Required Kernel Version

From the previous step, if the kernel version is found to be lower than kernel-2.6.32-431.17.1.el6, install the kernel-2.6.32-431.17.1.el6 or later:

1. To install kernel-2.6.32-431.17.1.el6 version, run the command:

```
# yum install kernel-2.6.32-431.17.1.el6
```

- To install the latest kernel version, run the command:

```
# yum update kernel
```

9. Install Red Hat Storage

Run the following command to install Red Hat Storage:

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

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

```
# yum groupinstall "Samba (SMB) server for gluster"
```

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

```
# yum groupinstall "Active Directory Integration"
```

10. Reboot the system



WARNING

Red Hat Storage server currently does not support SELinux. You must reboot the system after the layered install is complete, in order to disable SELinux on the system.

4.3. 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 Storage Server installation system without any other media.

To boot a computer from a PXE server:

- 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.
- Switch on the computer.
- 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.



IMPORTANT

Check the Security-Enhanced Linux (SELinux) status on the Red Hat Storage Server after installation. You must ensure that SELinux is disabled if it is found to be enforced or permissive. For more information on enabling and disabling SELinux, see https://access.redhat.com/knowledge/docs/en-US/Red_Hat_Enterprise_Linux/6/html/Security-Enhanced_Linux/sect-Security-Enhanced_Linux-Enabling_and_Disabling_SELinux-Disabling_SELinux.html.

4.4. INSTALLING FROM RED HAT SATELLITE SERVER

Ensure that the firewall settings are configured so that the required ports are open. For a list of port numbers, see [Section 3.3, “Port Information”](#).

Creating the Activation Key

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

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

The screenshot shows the Red Hat Satellite web interface. At the top, there's a navigation bar with 'Overview', 'Systems', 'Environments', 'Channels', 'Audit', 'Configuration', 'Schedules', 'Users', 'Admin', and 'Help'. A search bar is on the right. Below the navigation bar, a message states 'Activation key RHS3_ActivationKey has been created.' The main content area is titled 'RHS3_ActivationKey' and has a 'delete key' button. The 'Details' tab is selected, showing 'Activation Key Details'. A note says 'Systems registered with this activation key will inherit the settings listed below.' The form fields are:

- Description:** RHS3_ActivationKey (with a tip: 'Use this to describe what kind of settings this key will reflect on systems that use it. If left blank, this field will be filled in "None"').
- Key:** 1- [a5adcaaf4d1d15b7c75f8d77] (with a tip: 'Leave blank for automatic key generation. Note that the prefix is an indication of the Red Hat Satellite organization the key is associated with.').
- Usage:** (with a tip: 'Leave blank for unlimited use.').
- Base Channels:** Red Hat Enterprise Linux Server (v. 6 for 64-bit x86_64) (with a tip: 'Choose "Red Hat Satellite Default" to allow systems to register to the default Red Hat provided channel that corresponds to their installed version of Red Hat Enterprise Linux. You may also choose particular Red Hat provided channels or custom base channels here, but please note if a system using this key is not compatible with the selected channel, it will fall back to the Red Hat default channel.').
- Add-On Entitlements:** Monitoring, Provisioning, Virtualization, Virtualization Platform (all unchecked).
- Configuration File Deployment:** (unchecked, with a tip: 'Configuration File Deployment is currently disabled. Select "Provisioning" under "Add-On Entitlements" to enable configuration management facilities for this activation key.').
- Universal Default:** (unchecked, with a tip: 'Only one universal default activation key may be set for this organization. By setting this key as universal default, you will remove universal default status from the current universal default key if it exists. If this key is set as universal default, then newly-registered systems to your organization will inherit the properties of this key.').

 An 'Update Activation Key' button is at the bottom right.

Figure 4.13. 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 Storage Server 3 (RHEL 6 for x86_64)

For Red Hat 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)

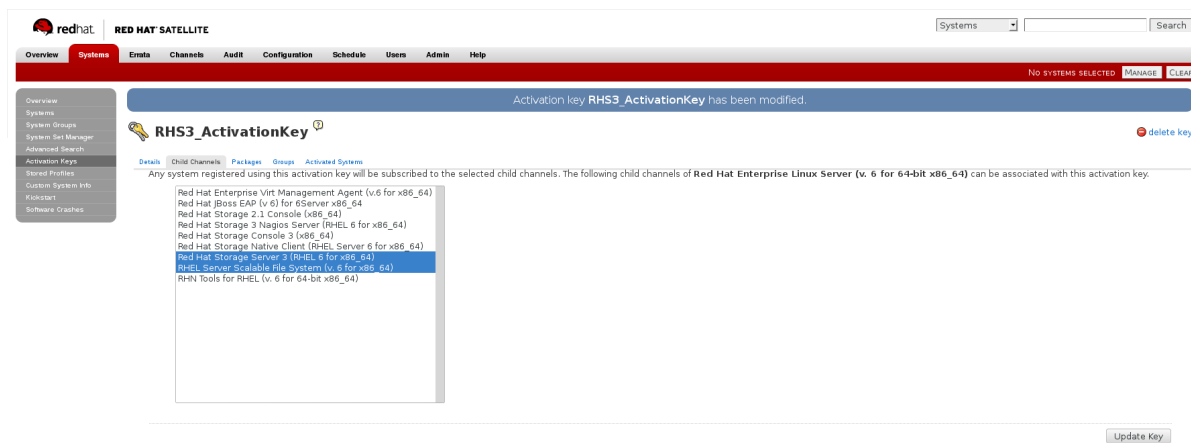


Figure 4.14. Child Channels

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

```
redhat - storage-server
```

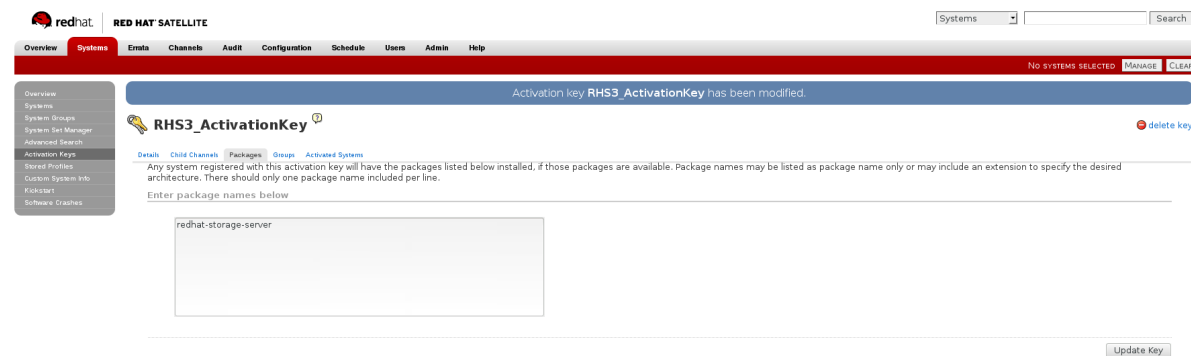


Figure 4.15. Package

- For Red Hat 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 Storage Server using the Kickstart Profile

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

CHAPTER 5. SUBSCRIBING TO THE RED HAT STORAGE SERVER CHANNELS

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



NOTE

If you used Red Hat Satellite or Layered Installation method to install Red Hat Storage 3.0, then you can skip the steps mentioned in this chapter. For more information regarding layered installation, see [Section 4.2, “Installing Red Hat 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 --auto-attach
```

2. Enable the Required Channels

Run the following commands to enable the repos required to install Red Hat 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
```

1. For Red Hat Storage 3.0.4 or later, if you require Samba, then enable the following channel:

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

3. Verify if the Channels are Enabled

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

```
# yum repolist
```

Using Red Hat Satellite Server

1. Configure the Client System to Access Red Hat Satellite

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

2. Register to the Red Hat Satellite Server

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

```
# rhn_register
```

3. Register to the Standard Base Channel

In the select operating system release page, select **All available updates** and follow the prompts to register the system to the standard base channel for RHEL6 - rhel-x86_64-server-6

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

Run the following command to subscribe the system to the required Red Hat Storage server channels:

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

1. For Red Hat 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
```

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

For Red Hat Storage 3.0.4 or later, if you have enabled the Samba channel, then you will also see the following channel while verifying:

```
rhel-x86_64-server-6-rh-gluster-3-samba
```


CHAPTER 6. UPGRADING RED HAT STORAGE

This chapter describes the procedure to upgrade to Red Hat Storage 3.0 from Red Hat Storage 2.1.

6.1. UPGRADING FROM RED HAT STORAGE 2.1 TO RED HAT STORAGE 3.0 USING AN ISO

This method re-images the software in the storage server by keeping the data intact after a backup-restore of the configuration files. This method is quite invasive and should only be used if a local yum repository or an Internet connection to access Red Hat Network is not available.

The preferable method to upgrade is using the **yum** command. For more information, refer to [Section 6.2, “Upgrading from Red Hat Storage 2.1 to Red Hat Storage 3.0 for Systems Subscribed to Red Hat Network”](#).



NOTE

- Ensure that you perform the steps listed in this section on all the servers.
- In the case of a geo-replication set-up, perform the steps listed in this section on all the master and slave servers.
- You cannot access data during the upgrade process, and a downtime should be scheduled with applications, clients, and other end-users.

1. Get the volume information and peer status using the following commands:

```
# gluster volume info
```

The command displays the volume information similar to the following:

```
Volume Name: volname
Type: Distributed-Replicate
Volume ID: d6274441-65bc-49f4-a705-fc180c96a072
Status: Started
Number of Bricks: 2 x 2 = 4
Transport-type: tcp
Bricks:
Brick1: server1:/rhs/brick1/brick1
Brick2: server2:/rhs/brick1/brick2
Brick3: server3:/rhs/brick1/brick3
Brick4: server4:/rhs/brick1/brick4
Options Reconfigured:
geo-replication.indexing: on
```

```
# gluster peer status
```

The command displays the peer status information similar to the following:

```
# gluster peer status
Number of Peers: 3

Hostname: server2
Port: 24007
```

```

Uuid: 2dde2c42-1616-4109-b782-dd37185702d8
State: Peer in Cluster (Connected)

```

```

Hostname: server3
Port: 24007
Uuid: 4224e2ac-8f72-4ef2-a01d-09ff46fb9414
State: Peer in Cluster (Connected)

```

```

Hostname: server4
Port: 24007
Uuid: 10ae22d5-761c-4b2e-ad0c-7e6bd3f919dc
State: Peer in Cluster (Connected)

```

**NOTE**

Make a note of this information to compare with the output after upgrading.

2. In case of a geo-replication set-up, stop the geo-replication session using the following command:

```

# gluster volume geo-replication master_volname
slave_node::slave_volname stop

```

3. In case of a CTDB/Samba set-up, stop the CTDB service using the following command:

```

# service ctdb stop ;Stopping the CTDB service also stops the SMB
service

```

1. Verify if the CTDB and the SMB services are stopped using the following command:

```

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

```

4. In case of an object store set-up, turn off object store using the following commands:

```

# service gluster-swift-proxy stop
# service gluster-swift-account stop
# service gluster-swift-container stop
# service gluster-swift-object stop

```

5. Stop all the gluster volumes using the following command:

```

# gluster volume stop volname

```

6. Stop the **glusterd** services on all the nodes using the following command:

```

# service glusterd stop

```

7. If there are any gluster processes still running, terminate the process using **kill**.

8. Ensure all gluster processes are stopped using the following command:

-

```
# pgrep gluster
```

- Back up the following configuration directory and files on the backup directory:

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

Ensure that the backup directory is not the operating system partition.

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

Also, back up any other files or configuration files that you might require to restore later. You can create a backup of everything in `/etc/`.

- Locate and unmount the data disk partition that contains the bricks using the following command:

```
# mount | grep backend-disk  
# umount /dev/device
```

For example, use the `gluster volume info` command to display the `backend-disk` information:

```
Volume Name: volname  
Type: Distributed-Replicate  
Volume ID: d6274441-65bc-49f4-a705-fc180c96a072  
Status: Started  
Number of Bricks: 2 x 2 = 4  
Transport-type: tcp  
Bricks:  
Brick1: server1:/rhs/brick1/brick1  
Brick2: server2:/rhs/brick1/brick2  
Brick3: server3:/rhs/brick1/brick3  
Brick4: server4:/rhs/brick1/brick4  
Options Reconfigured:  
geo-replication.indexing: on
```

In the above example, the `backend-disk` is mounted at `/rhs/brick1`

```
# findmnt /rhs/brick1  
TARGET          SOURCE          FSTYPE OPTIONS  
/rhs/brick1 /dev/mapper/glustervg-brick1 xfs  
rw,relatime,attr2,delaylog,no  
# umount /rhs/brick1
```

- Insert the DVD with Red Hat Storage 3.0 ISO and reboot the machine. The installation starts automatically. You must install Red Hat Storage on the system with the same network credentials, IP address, and host name.



WARNING

During installation, while creating a custom layout, ensure that you choose **Create Custom Layout** to proceed with installation. If you choose **Replace Existing Linux System(s)**, it formats all disks on the system and erases existing data.

Select **Create Custom Layout**. Click **Next**.

Which type of installation would you like?

Use All Space
Removes all partitions on the selected device(s). This includes partitions created by other operating systems.
Tip: This option will remove data from the selected device(s). Make sure you have backups.

Replace Existing Linux System(s)
Removes only Linux partitions (created from a previous Linux installation). This does not remove other partitions you may have on your storage device(s) (such as VFAT or FAT32).
Tip: This option will remove data from the selected device(s). Make sure you have backups.

Shrink Current System
Shrinks existing partitions to create free space for the default layout.

Use Free Space
Retains your current data and partitions and uses only the unpartitioned space on the selected device(s), assuming you have enough free space available.

Create Custom Layout
Manually create your own custom layout on the selected device(s) using our partitioning tool.

Encrypt system
 Review and modify partitioning layout

[← Back](#) [Next →](#)

Figure 6.1. Custom Layout Window

12. Select the disk on which to install Red Hat Storage. Click **Next**.

For Red Hat Storage to install successfully, you must select the same disk that contained the operating system data previously.



WARNING

While selecting your disk, do not select the disks containing bricks.

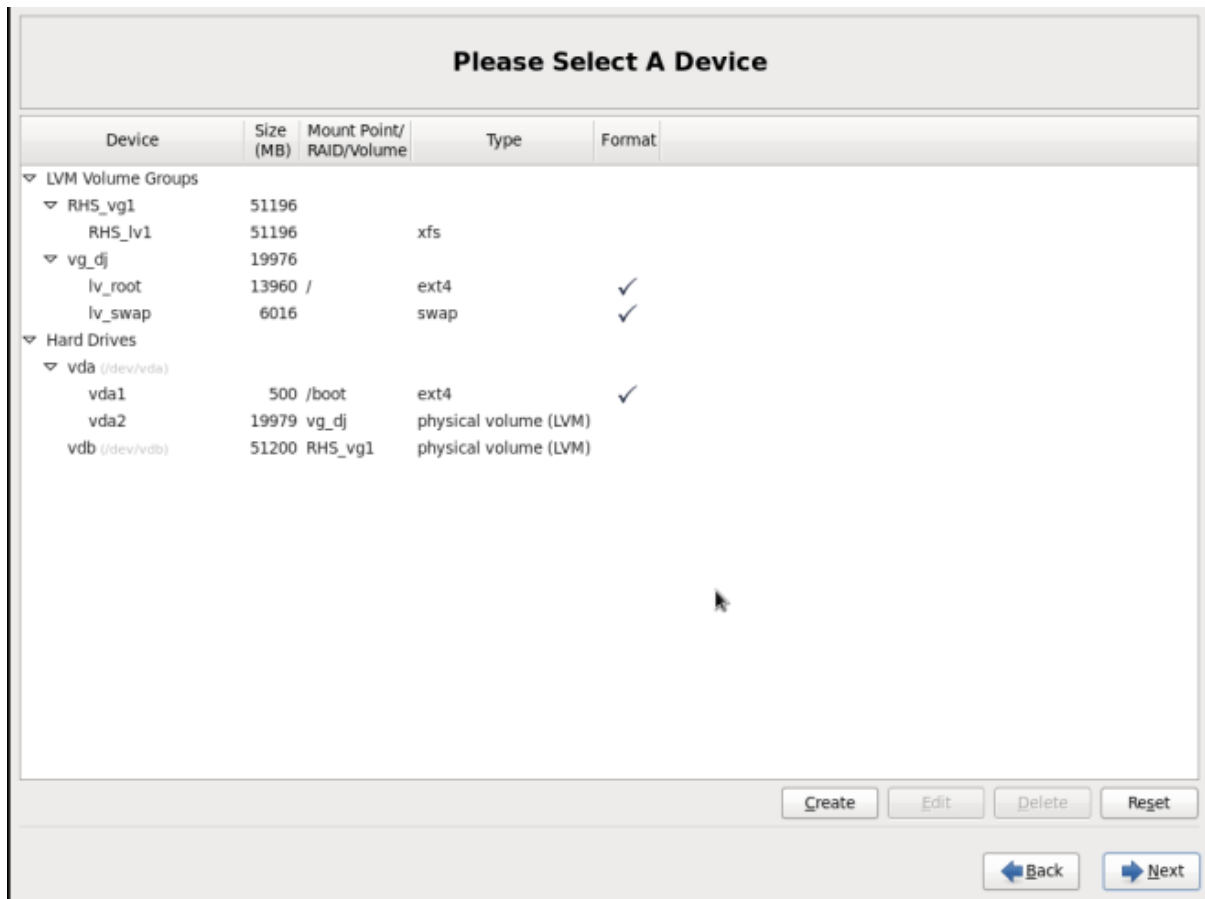


Figure 6.2. Select Disk Partition Window

- After installation, ensure that the host name and IP address of the machine is the same as before.



WARNING

If the IP address and host name are not the same as before, you will not be able to access the data present in your earlier environment.

- After installation, the system automatically starts **glusterd**. Stop the gluster service using the following command:

```
# service glusterd stop
Stopping glusterd: [OK]
```

- Add entries to **/etc/fstab** to mount data disks at the same path as before.



NOTE

Ensure that the mount points exist in your trusted storage pool environment.

- Mount all data disks using the following command:


```
# mount -a
```

17. Back up the latest **glusterd** using the following command:

```
# cp -a /var/lib/glusterd /var/lib/glusterd-backup
```

18. Copy **/var/lib/glusterd** and **/etc/glusterfs** from your backup disk to the OS disk.

```
# cp -a /backup-disk/glusterd/* /var/lib/glusterd
# cp -a /backup-disk/glusterfs/* /etc/glusterfs
```



NOTE

Do not restore the swift, samba and ctdb configuration files from the backup disk. However, any changes in swift, samba, and ctdb must be applied separately in the new configuration files from the backup taken earlier.

19. Copy back the latest hooks scripts to **/var/lib/glusterd/hooks**.

```
# cp -a /var/lib/glusterd-backup/hooks /var/lib/glusterd
```

20. Ensure you restore any other files from the backup that was created earlier.

21. You must restart the **glusterd** management daemon using the following commands:

```
# glusterd --xlator-option *.upgrade=yes -N
# service glusterd start
Starting glusterd: [OK]
```

22. Start the volume using the following command:

```
# gluster volume start volname force
volume start: volname : success
```



NOTE

Repeat the above steps on all the servers in your trusted storage pool environment.

23. In case you have a pure replica volume (1*n) where n is the replica count, perform the following additional steps:

1. Run the **fix-layout** command on the volume using the following command:

```
# gluster volume rebalance volname fix-layout start
```

2. Wait for the **fix-layout** command to complete. You can check the status for completion using the following command:

```
# gluster volume rebalance volname status
```

- 3. Stop the volume using the following command:

```
# gluster volume stop volname
```


- 4. Force start the volume using the following command:

```
# gluster volume start volname force
```

- 24. In case of an Object Store set-up, any configuration files that were edited should be renamed to end with a **.rpmsave** file extension, and other unedited files should be removed.
- 25. Re-configure the Object Store. For information on configuring Object Store, refer to *Section 18.5* in *Chapter 18. Managing Object Store* of the *Red Hat Storage Administration Guide*.
- 26. Get the volume information and peer status of the created volume using the following commands:

```
# gluster volume info
# gluster peer status
```

Ensure that the output of these commands has the same values that they had before you started the upgrade.



NOTE

In Red Hat Storage 3.0, the **gluster peer status** output does not display the port number.

- 27. Verify the upgrade.
 - 1. If all servers in the trusted storage pool are not upgraded, the **gluster peer status** command displays the peers as disconnected or rejected.

The command displays the peer status information similar to the following:

```
# gluster peer status
Number of Peers: 3

Hostname: server2
Uuid: 2dde2c42-1616-4109-b782-dd37185702d8
State: Peer Rejected (Connected)

Hostname: server3
Uuid: 4224e2ac-8f72-4ef2-a01d-09ff46fb9414
State: Peer in Cluster (Connected)

Hostname: server4

Uuid: 10ae22d5-761c-4b2e-ad0c-7e6bd3f919dc
State: Peer Rejected (Disconnected)
```

2. If all systems in the trusted storage pool are upgraded, the **gluster peer status** command displays peers as connected.

The command displays the peer status information similar to the following:

```
# gluster peer status
Number of Peers: 3

Hostname: server2
Uuid: 2dde2c42-1616-4109-b782-dd37185702d8
State: Peer in Cluster (Connected)

Hostname: server3
Uuid: 4224e2ac-8f72-4ef2-a01d-09ff46fb9414
State: Peer in Cluster (Connected)

Hostname: server4
Uuid: 10ae22d5-761c-4b2e-ad0c-7e6bd3f919dc
State: Peer in Cluster (Connected)
```

3. If all the volumes in the trusted storage pool are started, the **gluster volume info** command displays the volume status as started.

```
Volume Name: volname
Type: Distributed-Replicate
Volume ID: d6274441-65bc-49f4-a705-fc180c96a072
Status: Started
Number of Bricks: 2 x 2 = 4
Transport-type: tcp
Bricks:
Brick1: server1:/rhs/brick1/brick1
Brick2: server2:/rhs/brick1/brick2
Brick3: server3:/rhs/brick1/brick3
Brick4: server4:/rhs/brick1/brick4
Options Reconfigured:
geo-replication.indexing: on
```

28. If you have a geo-replication setup, re-establish the geo-replication session between the master and slave using the following steps:

1. Run the following commands on any one of the master nodes:

```
# cd /usr/share/glusterfs/scripts/
# sh generate-gfid-file.sh localhost:${master-vol} $PWD/get-gfid.sh /tmp/tmp.atyEmKyCjo/upgrade-gfid-values.txt
# scp /tmp/tmp.atyEmKyCjo/upgrade-gfid-values.txt root@${slavehost}:/tmp/
```

2. Run the following commands on a slave node:

```
# cd /usr/share/glusterfs/scripts/
# sh slave-upgrade.sh localhost:${slave-vol} /tmp/tmp.atyEmKyCjo/upgrade-gfid-values.txt $PWD/gsync-sync-gfid
```


**NOTE**

If the SSH connection for your setup requires a password, you will be prompted for a password for all machines where the bricks are residing.

3. Re-create and start the geo-replication sessions.

For information on creating and starting geo-replication sessions, refer to *Managing Geo-replication* in the *Red Hat Storage Administration Guide*.

**NOTE**

It is recommended to add the child channel of Red Hat Enterprise Linux 6 containing the native client, so that you can refresh the clients and get access to all the new features in Red Hat Storage 3.0. For more information, refer to the *Upgrading Native Client* section in the *Red Hat Storage Administration Guide*.

29. Remount the volume to the client and verify for data consistency. If the gluster volume information and gluster peer status information matches with the information collected before migration, you have successfully upgraded your environment to Red Hat Storage 3.0.

6.2. UPGRADING FROM RED HAT STORAGE 2.1 TO RED HAT STORAGE 3.0 FOR SYSTEMS SUBSCRIBED TO RED HAT NETWORK

Pre-Upgrade Steps:

1. Unmount the clients using the following command:

```
umount mount-point
```

2. Stop the volumes using the following command:

```
gluster volume stop volname
```

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

```
umount mount-point
```

4. To verify if the volume status is stopped, use the following command:

```
# gluster volume info
```

If there is more than one volume, stop all of the volumes.

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

```
# service glusterd stop
```

yum Upgrade Steps:



IMPORTANT

- You can upgrade to Red Hat Storage 3.0 only from Red Hat Storage 2.1 Update 4 release. If your current version is lower than Update 4, then upgrade it to Update 4 before upgrading to Red Hat Storage 3.0.
- Upgrade the servers before upgrading the clients.

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

```
# pkill gluster
```

2. To check the system's current subscription status run the following command:

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

3. Install the required packages using the following command:

```
# yum install subscription-manager-migration  
# yum install subscription-manager-migration-data
```

4. Execute the following command to migrate from Red Hat Network Classic to Red Hat Subscription Manager

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

5. To enable the Red Hat Storage 3.0 repos, execute the following command:

```
# migrate-rhs-classic-to-rhsm --upgrade --version 3
```

1. For Red Hat Storage 3.0.4 or later, if you require Samba, then enable the following channel:

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

**WARNING**

- The Samba version 3 is being deprecated from Red Hat 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.
- 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. Stop the CTDB and SMB services across all nodes in the Samba cluster using the following command. This is because different versions of Samba cannot run in the same Samba cluster.

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

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

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

6. To verify if the migration from Red Hat Network Classic to Red Hat Subscription Manager is successful, execute the following command:

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

7. To upgrade the server from Red Hat Storage 2.1 to 3.0, use the following command:

```
# yum update
```

**NOTE**

It is recommended to add the child channel of Red Hat Enterprise Linux 6 that contains the native client to refresh the clients and access the new features in Red Hat Storage 3.0. For more information, refer to *Installing Native Client* in the *Red Hat Storage Administration Guide*.

8. Reboot the servers. This is required as the kernel is updated to the latest version.

6.3. UPGRADING FROM RED HAT STORAGE 2.1 TO RED HAT STORAGE 3.0 FOR SYSTEMS SUBSCRIBED TO RED HAT SATELLITE SERVER

1. Unmount all the clients using the following command:

```
umount mount-name
```

2. Stop the volumes using the following command:

```
# gluster volume stop volname
```

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

```
umount mount-point
```

4. Ensure that the Red Hat Storage 2.1 server is updated to Red Hat Storage 2.1 Update 4, by running the following command:

```
# yum update
```

5. Create an Activation Key at the Red Hat Satellite Server, and associate it with the following channels. For more information, refer to [Section 4.4, “Installing from Red Hat Satellite Server”](#)

```
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 Storage Server 3 (RHEL 6 for x86_64)
```

1. For Red Hat Storage 3.0.4 or later, if you require the Samba package add the following child channel:

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

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

1. Log in to the Red Hat Satellite server.
2. Click on the **Systems** tab in the top navigation bar and then the name of the old or duplicated system in the **System List**.
3. Click the **delete system** link in the top-right corner of the page.
4. To confirm the system profile deletion by clicking the **Delete System** button.

7. On the updated Red Hat Storage 2.1 Update 4 server, run the following command:

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

This uses the prepared Activation Key and re-registers the system to the Red Hat Storage 3.0 channels on the Red Hat Satellite Server.

8. Verify if the channel subscriptions have changed to the following:

■

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

For Red Hat Storage 3.0.4 or later, if you have enabled the Samba channel, then verify if you have the following channel:

```
rhel-x86_64-server-6-rh-gluster-3-samba
```

- Run the following command to upgrade to Red Hat Storage 3.0.

```
# yum update
```

- Reboot, and run volume and data integrity checks.

6.4. UPGRADING FROM RED HAT STORAGE 2.1 TO RED HAT STORAGE 3.0 IN A RED HAT ENTERPRISE VIRTUALIZATION-RED HAT STORAGE ENVIRONMENT

This section describes the upgrade methods for a Red Hat Storage and Red Hat Enterprise Virtualization integrated environment. You can upgrade Red Hat Storage 2.1 to Red Hat Storage 3.0 using an ISO or `yum`.

6.4.1. Upgrading using an ISO

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

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



NOTE

Ensure you stop the volume, as rolling upgrade is not supported in Red Hat Storage.

- Using Red Hat Enterprise Virtualization Manager, move the data domain of the data center to *Maintenance* mode.
- Using Red Hat Enterprise Virtualization Manager, stop the volume (the volume used for data domain) containing Red Hat Storage nodes in the data center.
- Using Red Hat Enterprise Virtualization Manager, move all Red Hat Storage nodes to *Maintenance* mode.
- Perform the ISO Upgrade as mentioned in [Section 6.1, “Upgrading from Red Hat Storage 2.1 to Red Hat Storage 3.0 using an ISO”](#).

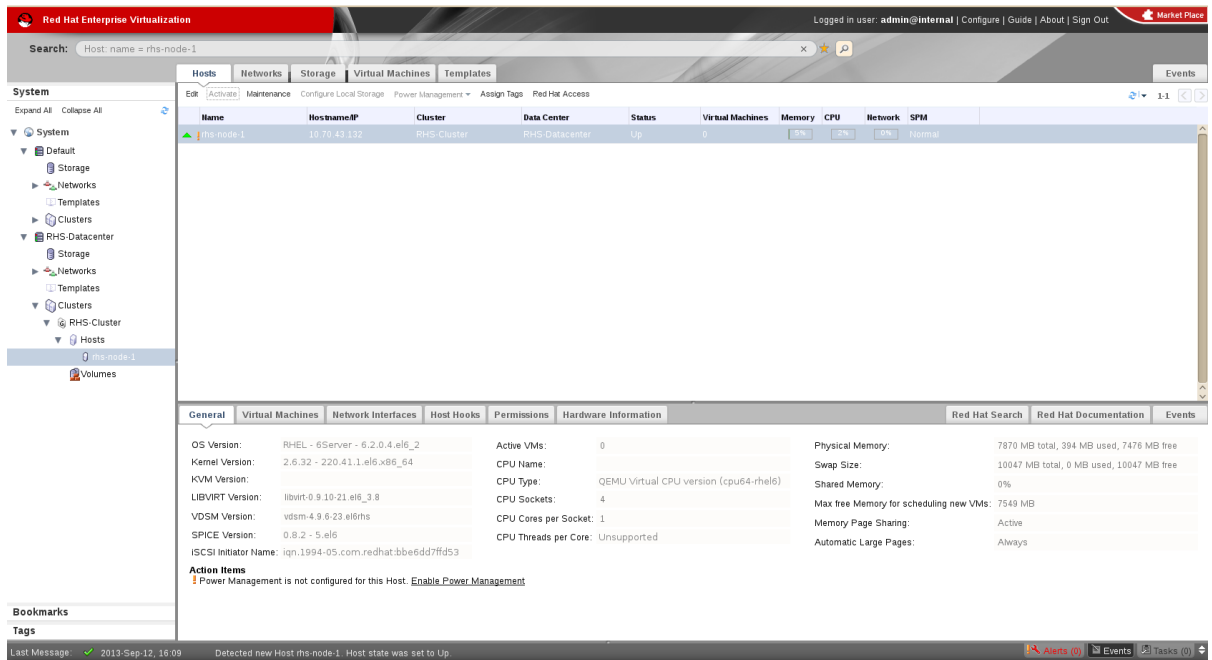


Figure 6.3. Red Hat Storage Node

6. Re-install the Red Hat Storage nodes from Red Hat Enterprise Virtualization Manager.



NOTE

- o Re-installation for the Red Hat Storage nodes should be done from Red Hat Enterprise Virtualization Manager. The newly upgraded Red Hat Storage 3.0 nodes lose their network configuration and other configurations, such as iptables configuration, done earlier while adding the nodes to Red Hat Enterprise Virtualization Manager. Re-install the Red Hat Storage nodes to have the bootstrapping done.
- o You can re-configure the Red Hat Storage nodes using the option provided under **Action Items**, as shown in [Figure 6.4, “Red Hat Storage Node before Upgrade”](#), and perform bootstrapping.

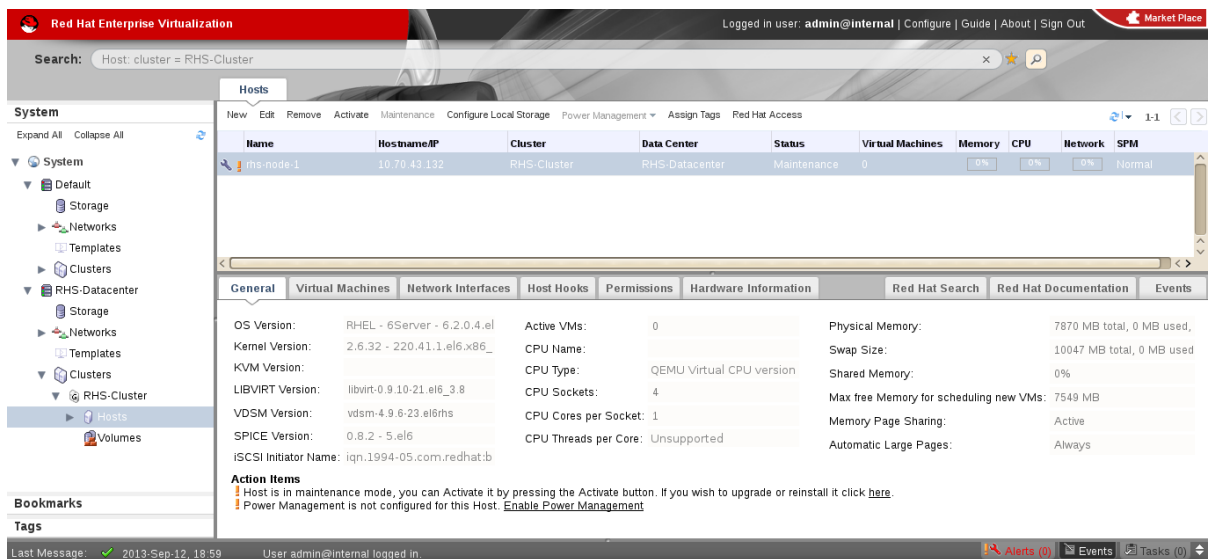


Figure 6.4. Red Hat Storage Node before Upgrade

7. Perform the steps above in all Red Hat 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 6.4, if Red Hat Enterprise Linux 6.4 is used as hypervisor.



NOTE

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

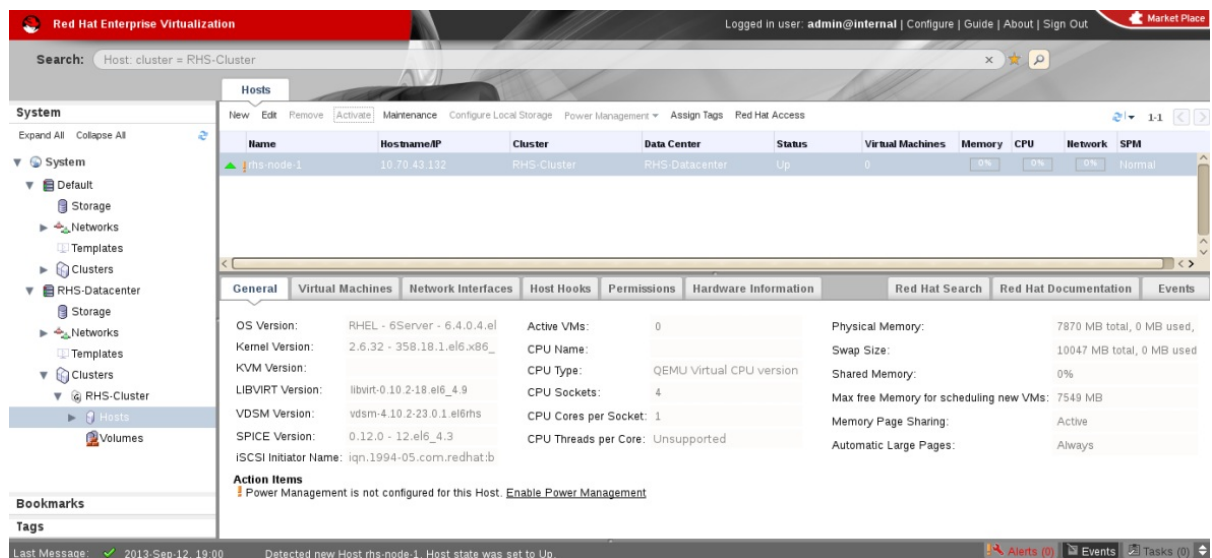


Figure 6.5. Red Hat Storage Node after Upgrade

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

6.4.2. 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 Storage nodes to *Maintenance* mode.
4. Perform **yum** update as mentioned in [Section 6.2, “Upgrading from Red Hat Storage 2.1 to Red Hat Storage 3.0 for Systems Subscribed to Red Hat Network”](#).
5. Once the Red Hat Storage nodes are rebooted and up, **Activate** them using Red Hat Enterprise Virtualization Manager.



NOTE

Re-installation of Red Hat 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 on Red Hat Enterprise Linux 6.4, in case Red Hat Enterprise Linux 6.4 is used as hypervisor.



NOTE

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

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

CHAPTER 7. DEPLOYING SAMBA ON RED HAT STORAGE

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

7.1. PREREQUISITES

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

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



WARNING

- o The Samba version 3 is being deprecated from Red Hat 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.
- o Downgrade of Samba from Samba 4.x to Samba 3.x is not supported.
- o 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:

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

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

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

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

7.2. INSTALLING SAMBA USING ISO

During the installation of Red Hat Storage, ensure you select the **Samba (SMB) server for gluster** component, in the **Customizing the Software Selection** screen:

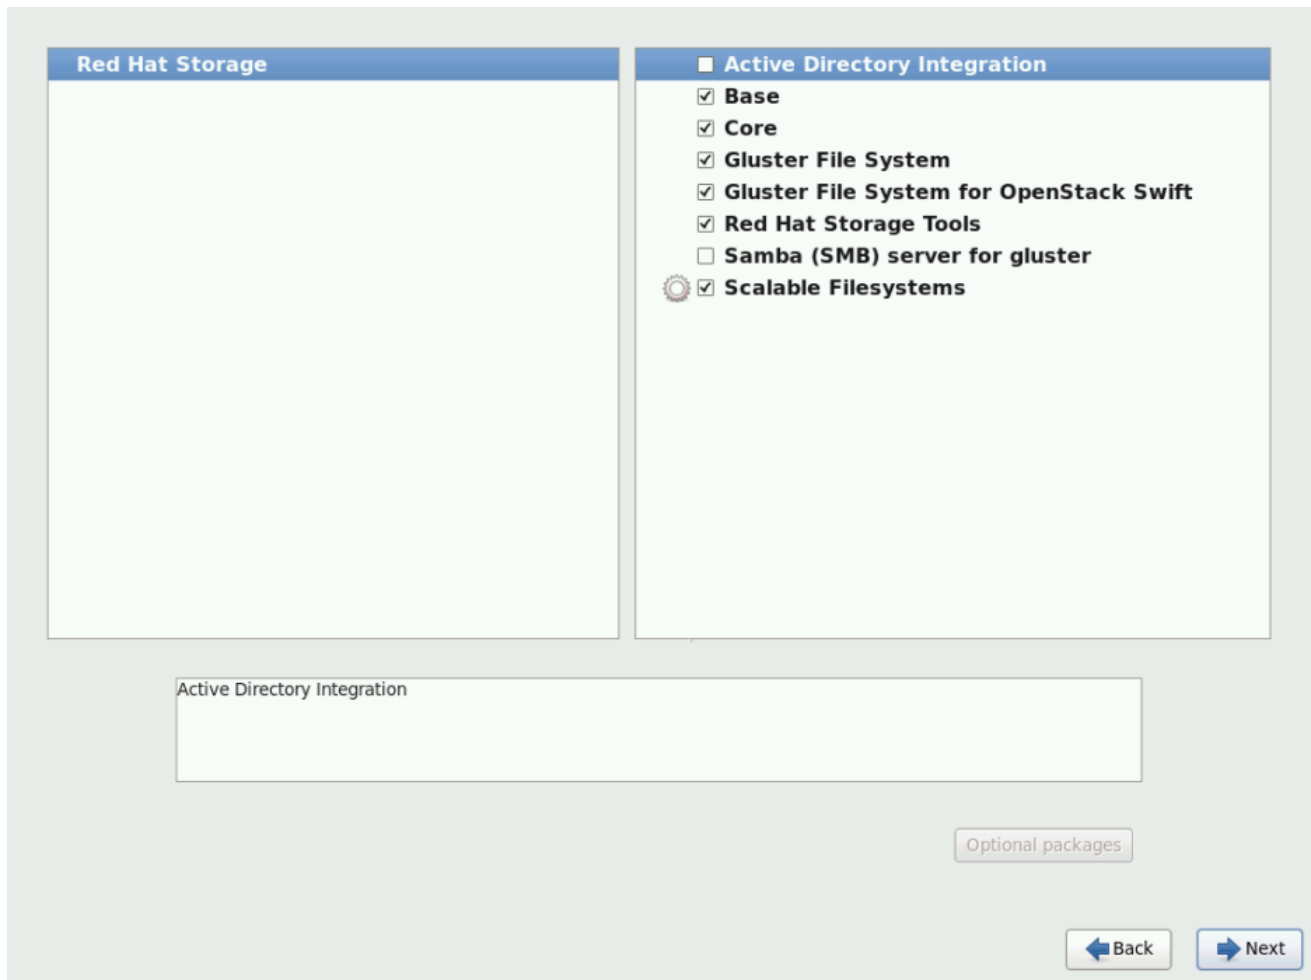


Figure 7.1. Customize Packages

For more information about installing Red Hat Storage using an ISO, see [Section 4.1, “Installing from an ISO Image”](#)

7.3. INSTALLING SAMBA USING YUM

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

```
# yum groupinstall "Samba (SMB) server for gluster"
```

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

```
# yum groupinstall "Active Directory Integration"
```

For more information about installing Red Hat Storage using yum, see [Section 4.2, “Installing Red Hat Storage Server on Red Hat Enterprise Linux \(Layered Install\)”](#).

CHAPTER 8. DEPLOYING THE HORTONWORKS DATA PLATFORM 2.1 ON RED HAT STORAGE

Red Hat Storage provides compatibility for Apache Hadoop and it uses the standard file system APIs available in Hadoop to provide a new storage option for Hadoop deployments. Red Hat has created a Hadoop File System plug-in that enables Hadoop Distributions to run on Red Hat Storage.

8.1. PREREQUISITES

Before you begin installation, you must establish the basic infrastructure required to enable Hadoop to run on Red Hat Storage.

8.1.1. Supported Versions

The following table lists the supported versions of HDP and Ambari with Red Hat Storage Server.

Table 8.1. Red Hat Storage Server Support Matrix

Red Hat Storage Server version	HDP version	Ambari version
3.0.4	2.1	1.6.1
3.0.2	2.1	1.6.1
3.0.0	2.0.6	1.4.4

8.1.2. Software and Hardware Requirements

You must ensure that all the servers used in this environment meet the following requirements:

- Must have at least the following hardware specification:
 - 2 x 2 GHz 4 core processors
 - 32 GB RAM
 - 500 GB of storage capacity
 - 1 x 1 GbE NIC
- Must have iptables disabled.
- Must use fully qualified domain names (FQDN). For example rhs-1.server.com is acceptable, but rhs-1 is not allowed.
- Either, all servers must be configured to use a DNS server and must be able to use DNS for FQDN resolution or all the storage nodes must have the FQDN of all of the servers in the cluster listed in their `/etc/hosts` file.
- Must have the following users and groups available on all the servers.

User	Group
yarn	hadoop
mapred	hadoop
hive	hadoop
hcat	hadoop
ambari-qa	hadoop
hbase	hadoop
tez	hadoop
zookeeper	hadoop
oozie	hadoop
falcon	hadoop

The specific UIDs and GIDs for the respective users and groups are up to the Administrator of the trusted storage pool, but they must be consistent across the trusted storage pool. For example, if the "hadoop" user has a UID as 591 on one server, the hadoop user must have UID as 591 on all other servers. This can be quite a lot of work to manage using Local Authentication and it is common and acceptable to install a central authentication solution such as LDAP or Active Directory for your cluster, so that users and groups can be easily managed in one place. However, to use local authentication, you can run the script below on each server to create the users and groups and ensure they are consistent across the cluster:

```
groupadd hadoop -g 590; useradd -u 591 mapred -g hadoop; useradd -u 592 yarn -g hadoop; useradd -u 594 hcat -g hadoop; useradd -u 595 hive -g hadoop; useradd -u 590 ambari-qa -g hadoop; useradd -u 593 tez -g hadoop; useradd -u 596 oozie -g hadoop; useradd -u 597 zookeeper -g hadoop; useradd -u 598 falcon -g hadoop; useradd -u 599 hbase -g hadoop
```

8.1.3. Existing Red Hat Storage Trusted Storage Pool

If you have an existing Red Hat Storage trusted storage pool, you need to add two additional servers to run the Hortonworks Ambari Management Services and the YARN Master Services, respectively. For more information on recommended deployment topologies, see *Administering the Hortonworks Data Platform on Red Hat Storage* chapter in *Red Hat Storage Administration Guide*.

In addition, all nodes within the Red Hat Storage Trusted Storage Pool that contain volumes that are to be used with Hadoop must contain a local glusterfs-fuse mount of that volume. The path of the mount for each volume must be consistent across the cluster.

For information on expanding your trusted storage pool by adding servers, see section *Expanding Volumes* in the *Red Hat Storage 3.0 Administration Guide*.



NOTE

The supported volume configuration for Hadoop is Distributed Replicated volume with replica count 2 or 3.



IMPORTANT

New Red Hat Storage and Hadoop Clusters use the naming conventions of `/mnt/brick1` as the mount point for Red Hat Storage bricks and `/mnt/glusterfs/volname` as the mount point for Red Hat Storage volume. It is possible that you have an existing Red Hat Storage volume that has been created with different mount points for the Red Hat Storage bricks and volumes. If the mount points differ from the convention, replace the prefix listed in this installation guide with the prefix that you have.

Information on how to mount and configure bricks and volumes with required parameters and description of required local mount of gluster volume are available in [Section 8.2.5, “Enabling Existing Volumes for use with Hadoop”](#)

8.1.4. New Red Hat Storage Trusted Storage Pool

You must create a Red Hat Storage trusted storage pool with at least four bricks for two-way replication and with six bricks for three-way replication. The servers on which these bricks reside must have the Red Hat Storage installed on them. The number of bricks must be a multiple of the replica count for a distributed replicated volume.

For more information on installing Red Hat Storage see [Chapter 4, *Installing Red Hat Storage*](#) or for upgrading to Red Hat Storage 3.0, see [Chapter 6, *Upgrading Red Hat Storage*](#).

Red Hat recommends that you have an additional two servers set aside to run the Hortonworks Ambari Management Services and the YARN Master Services, respectively. Alternate deployment topologies are also possible, for more information on various supported deployment topologies, see *Administering the Hortonworks Data Platform on Red Hat Storage* chapter in *Red Hat Storage Administration Guide*.

For information on expanding your trusted storage pool by adding servers, see section *Expanding Volumes* in the *Red Hat Storage 3.0 Administration Guide*.



NOTE

The supported volume configuration for Hadoop is Distributed Replicated volume with replica count 2 or 3.

8.1.5. Red Hat Storage Server Requirements

You must install Red Hat Storage Server this server. While installing the server, you must ensure to specify a fully qualified domain name (FQDN). A hostname alone will not meet the requirements for the Hortonworks Data Platform Ambari deployment tool.

You must also enable the `rhs-big-data-3-for-rhel-6-server-rpms` channel on this server.

- If you have registered your machine using Red Hat Subscription Manager, enable the channel by running the following command:

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

- 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-rhs-bigdata-3
```

8.1.6. Hortonworks Ambari Server Requirements

You must install Red Hat Enterprise Linux 6.6 on the servers. You can also choose to install Red Hat Storage Console on this server as well, but this is optional. This allows all aspects of the Red Hat Storage trusted pool to be managed from a single server. While installing the server, you must ensure to specify a fully qualified domain name (FQDN). A hostname alone will not meet the requirements for the Horton Data Platform Ambari deployment tool. **It is mandatory to setup a passwordless-SSH connection from the Ambari Server to all other servers within the trusted storage pool.** Instructions for installing and configuring Hortonworks Ambari is provided in the further sections of this chapter.

If the Hortonworks Ambari server is installed on a different node than Red Hat Storage Server, you must also enable the **rhel-6-server-rh-common-rpms** channel on this server.

- If you have registered your machine using Red Hat Subscription Manager, enable the channel by running the following command:

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

- 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-rh-common-6
```



WARNING

Red Hat Storage Console enables Nagios Alerting for Red Hat Storage. The Nagios Client libraries are shipped with Red Hat Storage and are on each Red Hat Storage Server. This causes a conflict with the Nagios System that is bundled with the Hortonworks Data Platform (HDP). Hence, using Ambari to deploy and manage HDP Nagios is not supported.



NOTE

If you are using one of the condensed deployment topologies listed in the *Administration Guide* and you have elected to place the Ambari Management server on the same node as a Red Hat Storage Server, you must only enable the **rhs-big-data-3-for-rhel-6-server-rpms** channel on that server.

- If you have registered your machine using Red Hat Subscription Manager, enable the channel by running the following command:

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

- 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-rhs-bigdata-3
```

8.1.7. YARN Master Server Requirements

You must install the Red Hat Enterprise Linux 6.6 on this server. While installing the server, you must ensure to specify a fully qualified domain name (FQDN). A hostname alone will not meet the requirements for the Horton Data Platform Ambari deployment tool.

If the YARN Master server is installed on a different node than Red Hat Storage Server, you must also enable the **rhel-6-server-rh-common-rpms** and **rhel-6-server-rhs-client-1-rpms** channels on the YARN server.

- If you have registered your machine using Red Hat Subscription Manager, enable the channel by running the following command:

```
# subscription-manager repos --enable=rhel-6-server-rh-common-rpms --enable=rhel-6-server-rhs-client-1-rpms
```

- 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-rh-common-6
# rhn-channel --add --channel rhel-x86_64-server-rhsclient-6
```



NOTE

If you are using one of the condensed deployment topologies listed in the *Administration Guide* and you have elected to place the YARN Master server on the same node as a Red Hat Storage Server, you must only enable the **rhs-big-data-3-for-rhel-6-server-rpms** channel on that server.

- If you have registered your machine using Red Hat Subscription Manager, enable the channel by running the following command:

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

- 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-rhs-bigdata-3
```

8.2. INSTALLING THE HADOOP FILESYSTEM PLUGIN FOR RED HAT STORAGE

8.2.1. Adding the Hadoop Installer for Red Hat Storage

You must have the big-data channel added and the hadoop components installed on all the servers to use the Hadoop feature on Red Hat Storage. Run the following command on the Ambari Management Server, the YARN Master Server and all the servers within the Red Hat Storage trusted storage pool:

```
# yum install rhs-hadoop rhs-hadoop-install
```

8.2.2. Configuring the Trusted Storage Pool for use with Hadoop

Red Hat Storage provides a series of utility scripts that allows you to quickly prepare Red Hat Storage for use with Hadoop, and install the Ambari Management Server. You must first run the Hadoop cluster configuration initial script to install the Ambari Management Server, prepare the YARN Master Server to host the Resource Manager and Job History Server services for Red Hat Storage and build a trusted storage pool if it does not exist.



NOTE

You must run the script given below irrespective of whether you have an existing Red Hat Storage trusted storage pool or not.

To run the Hadoop configuration initial script:

1. Open the terminal window of the server designated to be the Ambari Management Server and navigate to the **/usr/share/rhs-hadoop-install/** directory.
2. Run the hadoop cluster configuration script as given below:


```
setup_cluster.sh [-y] [--hadoop-mgmt-node <node>] [--yarn-master
<node>] [--profile <profile>] [--ambari-repo <url>] <node-list-
spec>
```

where <node-list-spec> is


```
<node1>:<brickmnt1>:<blkdev1> <node2>[:<brickmnt2>][:<blkdev2>]
[<node3>[:<brickmnt3>][:<blkdev3>]] ... [<nodeN>[:<brickmntN>][:
<blkdevN>]]
```

where

- <brickmnt> is the name of the XFS mount for the above <blkdev>, for example, **/mnt/brick1** or **/external/HadoopBrick**. When a Red Hat Storage volume is created its bricks has the volume name appended, so <brickmnt> is a prefix for the volume's bricks. Example: If a new volume is named **HadoopVol** then its brick list would be:
<node>:/mnt/brick1/HadoopVol or
<node>:/external/HadoopBrick/HadoopVol.
- <blkdev> is the name of a Logical Volume device path, for example, **/dev/VG1/LV1** or **/dev/mapper/VG1-LV1**. Since LVM is a prerequisite for Red Hat Storage, the <blkdev> is not expected to be a raw block path, such as **/dev/sdb**.

Given below is an example of running the *setup_cluster.sh* script on the Ambari Management server and four Red Hat Storage Nodes which have the same logical volume and mount point intended to be used as a Red Hat Storage brick.

```
./setup_cluster.sh --yarn-master yarn.hdp rhs-
1.hdp:/mnt/brick1:/dev/rhs_vg1/rhs_lv1 rhs-2.hdp rhs-3.hdp rhs-4.hdp
```



NOTE

If a brick mount is omitted, the brick mount of the first node is used and if one block device is omitted, the block device of the first node is used.

8.2.3. Creating Volumes for use with Hadoop



NOTE

If an existing Red Hat Storage volume is used with Hadoop, skip this section and continue with the instruction in the next section.

Whether you have a new or existing Red Hat Storage trusted storage pool, to create a volume for use with Hadoop, the volume need to be created in such a way as to support Hadoop workloads. The supported volume configuration for Hadoop is Distributed Replicated volume with replica count 2 or 3. You must not name the Hadoop enabled Red Hat Storage volume as **hadoop** or **mapredlocal**.


Run the script given below to create new volumes that you intend to use with Hadoop. The script provides the necessary configuration parameters to the volume as well as updates the Hadoop Configuration to make the volume accessible to Hadoop.

1. Open the terminal window of the server designated to be the Ambari Management Server and navigate to the `/usr/share/rhs-hadoop-install/` directory.
2. Run the hadoop cluster configuration script as given below:

```
create_vol.sh [-y] <volName> [--replica count] <volMountPrefix>
<node-list>
```

where

- o `<--replica count>` is the replica count. You can specify the replica count as 2 or 3. By default, the replica count is 2. The number of bricks must be a multiple of the replica count. The order in which bricks are specified determines how bricks are mirrored with each other. For example, first `n` bricks, where `n` is the replica count.
- o `<node-list>` is: `<node1>:<brickmnt> <node2>[:<brickmnt2>] <node3>[:<brickmnt3>] ... [<nodeN>[:<brickmntN>]`
- o `<brickmnt>` is the name of the XFS mount for the block devices used by the above nodes, for example, `/mnt/brick1` or `/external/HadoopBrick`. When a Red Hat Storage volume is created its bricks will have the volume name appended, so `<brickmnt>` is a prefix for the volume's bricks. For example, if a new volume is named **HadoopVol** then its brick list would be: `<node>:/mnt/brick1/HadoopVol` or `<node>:/external/HadoopBrick/HadoopVol`.



NOTE

The *node-list* for `create_vol.sh` is similar to the `node-list-spec` used by `setup_cluster.sh` except that a block device is not specified in `create_vol`.

Given below is an example on how to create a volume named *HadoopVol*, using 4 Red Hat Storage Servers, each with the same brick mount and mount the volume on `/mnt/glusterfs`

```
./create_vol.sh HadoopVol /mnt/glusterfs rhs-1.hdp:/mnt/brick1 rhs-
2.hdp rhs-3.hdp rhs-4.hdp
```

8.2.4. Deploying and Configuring the HDP 2.1 Stack on Red Hat Storage using Ambari Manager

Prerequisite

Before deploying and configuring the HDP stack, perform the following steps:

1. Open the terminal window of the server designated to be the Ambari Management Server and replace the **HDP 2.1.GlusterFS repoinfo.xml** file by the **HDP 2.1 repoinfo.xml** file.

```
cp /var/lib/ambari-
server/resources/stacks/HDP/2.1/repos/repoinfo.xml /var/lib/ambari-
server/resources/stacks/HDP/2.1.GlusterFS/repos/
```

You will be prompted to overwrite `/2.1.GlusterFS/repos/repoinfo.xml` file, type **yes** to overwrite the file.

2. Restart the Ambari Server.

```
# ambari-server restart
```


Perform the following steps to deploy and configure the HDP stack on Red Hat Storage:



IMPORTANT

This section describes how to deploy HDP on Red Hat Storage. Selecting **HDFS** as the storage selection in the HDP 2.1.GlusterFS stack is not supported. If you want to deploy HDFS, then you must select the HDP 2.1 stack (not HDP 2.1.GlusterFS) and follow the instructions of the Hortonworks documentation.

1. Launch a web browser and enter **http://hostname:8080** in the URL by replacing *hostname* with the hostname of your Ambari Management Server.



NOTE

If the Ambari Console fails to load in the browser, it is usually because iptables is still running. Stop iptables by opening a terminal window and run **service iptables stop** command.

2. Enter **admin** and **admin** for the username and password.
3. Assign a name to your cluster, such as **MyCluster**.
4. Select the **HDP 2.1 GlusterFS Stack** (if not already selected by default) and click **Next**.
5. On the **Install Options** screen:
 1. For **Target Hosts**, add the YARN server and all the nodes in the trusted storage pool.
 2. Select **Provide your SSH Private Key to automatically register hosts** and provide your Ambari Server private key that was used to set up passwordless-SSH across the cluster.
 3. Click **Register and Confirm** button. It may take a while for this process to complete.
6. For **Confirm Hosts**, it may take awhile for all the hosts to be confirmed.
 1. After this process is complete, you can ignore any warnings from the Host Check related to **File and Folder Issues**, **Package Issues** and **User Issues** as these are related to customizations that are required for Red Hat Storage.
 2. Click **Next** and ignore the Confirmation Warning.
7. For **Choose Services**, unselect HDFS and as a minimum select GlusterFS, Ganglia, YARN+MapReduce2, ZooKeeper and Tez.

**NOTE**

- The use of Storm and Falcon have not been extensively tested and as yet are not supported.
- Do not select the Nagios service, as it is not supported. For more information, see subsection *21.1. Deployment Scenarios* of chapter *21. Administering the Hortonworks Data Platform on Red Hat Storage* in the *Red Hat Storage 3.0 Administration Guide*.
- This section describes how to deploy HDP on Red Hat Storage. Selecting **HDFS** as the storage selection in the HDP 2.1 GlusterFS stack is not supported. If users wish to deploy HDFS, then they must select the HDP 2.1 (not HDP 2.1.GlusterFS) and follow the instructions in the Hortonworks documentation.

8. For **Assign Masters**, set all the services to your designated YARN Master Server.
 1. For ZooKeeper, select your YARN Master Server and at least 2 additional servers within your cluster.
 2. Click **Next** to proceed.
9. For **Assign Slaves and Clients**, select all the nodes as **NodeManagers** except the YARN Master Server.
 1. Click **Client** checkbox for each selected node.
 2. Click **Next** to proceed.
10. On the **Customize Services** screen:
 1. Click **YARN** tab, scroll down to the **yarn.nodemanager.log-dirs** and **yarn.nodemanager.local-dirs** properties and remove any entries that begin with **/mnt/glusterfs/**.

**IMPORTANT**

New Red Hat Storage and Hadoop Clusters use the naming convention of **/mnt/glusterfs/volname** as the mount point for Red Hat Storage volumes. If you have existing Red Hat Storage volumes that has been created with different mount points, then remove the entries of those mount points.


2. Update the following property on the YARN tab - Application Timeline Server section:

Key	Value
yarn.timeline-service.leveldb-timeline-store.path	/tmp/hadoop/yarn/timeline

3. Review other tabs that are highlighted in red. These require you to enter additional information, such as passwords for the respective services.

11. On the **Review** screen, review your configuration and then click **Deploy** button.
12. On the **Summary** screen, click the **Complete** button and ignore any warnings and the **Starting Services failed** statement. This is normal as there is still some addition configuration that is required before we can start the services.
13. Click **Next** to proceed to the Ambari Dashboard. Select the YARN service on the top left and click **Stop-All**. Do not click **Start-All** until you perform the steps in section [Section 8.5, “Verifying the Configuration”](#).

8.2.5. Enabling Existing Volumes for use with Hadoop



IMPORTANT

This section is mandatory for every volume you intend to use with Hadoop. It is not sufficient to run the `create_vol.sh` script, you must follow the steps listed in this section as well.

If you have a volume that you would like to analyze with Hadoop, and the volume was created by the above `create_vol.sh` script, then it must be *enabled* to support Hadoop workloads. Execute the `enable_vol.sh` script below to validate the volume's setup and to update Hadoop's `core-site.xml` configuration file which makes the volume accessible to Hadoop.

If you have a volume that was not created by the above `create_vol.sh` script, it is important to ensure that both the bricks and the volumes that you intend to use are properly mounted and configured. If they are not, the `enable_vol.sh` script will throw an exception and not run. Perform the following steps to mount and configure bricks and volumes with required parameters on all storage servers:

1. Bricks need to be an XFS formatted logical volume and mounted with the *noatime* and *inode64* parameters. For example, if we assume the logical volume path is `/dev/rhs_vg1/rhs_lv1` and that path is being mounted on `/mnt/brick1` then the `/etc/fstab` entry for the mount point should look as follows:

```
/dev/rhs_vg1/rhs_lv1    /mnt/brick1  xfs    noatime,inode64    0 0
```

2. Volumes must be mounted with the *entry-timeout=0*, *attribute-timeout=0*, *use-readdirp=no*, *_netdev* settings. Assuming your volume name is `HadoopVol`, the server's FQDN is `rhs-1.hdp` and your intended mount point for the volume is `/mnt/glusterfs/HadoopVol` then the `/etc/fstab` entry for the mount point of the volume must be as follows:

```
rhs-1.hdp:/HadoopVol /mnt/glusterfs/HadoopVol glusterfs entry-
timeout=0,attribute-timeout=0,use-readdirp=no,_netdev 0 0
```

Volumes that are to be used with Hadoop also need to have specific volume level parameters set on them. In order to set these, shell into a node within the appropriate volume's trusted storage pool and run the following commands (the examples assume the volume name is `HadoopVol`):

```
# gluster volume set HadoopVol performance.stat-prefetch off
# gluster volume set HadoopVol cluster.eager-lock on
# gluster volume set HadoopVol performance.quick-read off
```

3. Perform the following to create several Hadoop directories on that volume:

1. Open the terminal window of one of the Red Hat Storage nodes in the trusted storage pool and navigate to the `/usr/share/rhs-hadoop-install` directory.
2. Run the `bin/add_dirs.sh volume-mount-dir list-of-directories`, where `volume-mount-dir` is the path name for the glusterfs-fuse mount of the volume you intend to enable for Hadoop (including the name of the volume) and `list-of-directories` is the list generated by running `bin/gen_dirs.sh -d` script. For example:

```
# bin/add_dirs.sh /mnt/glusterfs/HadoopVol $(bin/gen_dirs.sh -d)
```

After completing these 3 steps, you are now ready to run the `enable_vol.sh` script.

Red Hat Storage-Hadoop has the concept of a **default** volume, which is the volume used when input and/or output URIs are unqualified. Unqualified URIs are common in Hadoop jobs, so defining the default volume, which can be set by `enable_vol.sh` script, is important. The default volume is the first volume appearing in the `fs.glusterfs.volume` property in the `/etc/hadoop/conf/core-site.xml` configuration file. The `enable_vol.sh` supports the `--make-default` option which, if specified, causes the supplied volume to be pre-pended to the above property, and thus, become the default volume. The default behavior for `enable_vol.sh` is to not make the target volume the default volume, meaning the volume name is appended, rather than prepended, to the above property value.

The `--user` and `--pass` options are required for the `enable_vol.sh` script to login into Ambari instance of the cluster to reconfigure Red Hat Storage volume related configuration.



NOTE

The supported volume configuration for Hadoop is Distributed Replicated volume with replica count 2 or 3. Also, when you run the `enable_vol` script for the first time, you must ensure to specify the `--make-default` option.

1. Open the terminal window of the server designated to be the Ambari Management Server and navigate to the `/usr/share/rhs-hadoop-install/` directory.
2. Run the Hadoop Trusted Storage pool configuration script as given below:

```
# enable_vol.sh [-y] [--make-default] [--hadoop-mgmt-node node] [--user admin-user] [--pass admin-password] [--port mgmt-port-num] [--yarn-master yarn-node] [--rhs-node storage-node] volName
```

For Example;

```
# enable_vol.sh --yarn-master yarn.hdp --rhs-node rhs-1.hdp
HadoopVol --make-default
```

**NOTE**

If `--yarn-master` and/or `--rhs-node` options are omitted then the default of `localhost` (the node from which the script is being executed) is assumed. Example:

```
./enable_vol.sh --yarn-master yarn.hdp --rhs-node rhs-1.hdp HadoopVol --make-default
```

If this is the first time you are running `enable_vol` script, you will see a warning **WARN: Cannot find configured default volume on node: rhs-1.hdp: "fs.glusterfs.volumes" property value is missing from /etc/hadoop/conf/core-site.xml**. This is normal and the system will proceed to set the volume you are enabling as the default volume. You will not see this message when subsequently enabling additional volume.

8.3. ADDING AND REMOVING USERS

Only users that are part of the *Hadoop* group (that was created in the prerequisites section) will be able to submit Hadoop Jobs. This can be relatively easily if you are using LDAP for authentication, but if you are not, you need to run the command given below on each server in the trusted storage pool and the YARN Master Server, for each user you add.

```
# useradd -u 1005 -g hadoop tom
```

**NOTE**

The UID of 1005 is arbitrary, you can specify the UID of your choice, but it must be both unique and consistent across the trusted storage pools. Also, the Hadoop Container Executor default properties require that all UIDs be greater or equal to 1000. Thus, 999 is not acceptable as user ID, but 1000 is acceptable. If you want to lower the default minimum acceptable UID, modify the `min.user.id` value in the `/etc/hadoop/conf/container-executor.cfg` file on every Red Hat Storage server that is running a NodeManager.

After adding a user who is part of the *hadoop* group, you need to create a user directory for that user within the default Red Hat Storage Hadoop Volume. The default Red Hat Storage Hadoop Volume is the first volume that was created and enabled for Hadoop and is usually called **HadoopVol** according to the examples given in installation instructions.

Run the following commands from a server within the Red Hat Storage trusted storage pool (replacing `user_name` with the actual user name) for each user that you are adding. You must run this command only once on the default Volume. If you add subsequent volumes, you do not need to repeat this step.

```
# mkdir /mnt/glusterfs/HadoopVol/user/<username>
# chown <username>:hadoop /mnt/glusterfs/HadoopVol/user/<username>
# chmod 0755 /mnt/glusterfs/HadoopVol/user/<username>
```

Removing Users

To disable a user from submitting Hadoop Jobs, remove the user from the Hadoop group.

8.4. DISABLING A VOLUME FOR USE WITH HADOOP

To keep a volume available for Red Hat Storage workloads, but not accessible to Hadoop jobs, you can disable the volume for use with Hadoop. Disabling a volume means that Hadoop jobs no longer have access to any data contained in the target volume; however, non-Hadoop workloads will continue to have access to this volume. The target volume is not deleted nor in anyway made unavailable to Red Hat Storage access. At a later time, disabled volumes can be re-enabled by executing the `enable_vol.sh` script.

For information on enabling volumes using `enable_vol.sh` script, see [Section 8.2.5, “Enabling Existing Volumes for use with Hadoop”](#).

A volume is disabled by modifying the `/etc/hadoop/conf/core-site.xml` file. Specifically, the volume's name is removed from the `fs.glusterfs.volumes` property list, and the `fs.glusterfs.volume.fuse.volname` property is deleted. All Ambari services are automatically restarted.

Perform the following steps to disable the volume:

1. Open the terminal window of the server designated to be the Ambari Management Server and navigate to the `/usr/share/rhs-hadoop-install/` directory.
2. Run the Hadoop cluster configuration script as shown below:

```
disable_vol.sh [-y] [--hadoop-mgmt-node node] [--user admin-user] [-pass admin-password] [--port mgmt-port-num] --yarn-master node [--rhs-node storage-node] volname
```

For example,

```
disable_vol.sh --rhs-node rhs-1.hdp --yarn-master yarn.hdp HadoopVol
```

8.5. VERIFYING THE CONFIGURATION

Open a terminal session on the YARN Master Server and run the following commands:

```
# chown -R yarn:hadoop /mnt/brick1/hadoop/yarn/
# chmod -R 0755 /mnt/brick1/hadoop/yarn/
```

Prior to submitting any jobs, ensure that the trusted storage pool is running. Launch the Ambari Dashboard (<http://ambari-server-hostname:8080>) and select the YARN service and then click the **Start - All** button.



NOTE

Stopping and starting the services takes some time. If one of the services fails to start, it will often start if you select the service and restart it.

The default volume (usually HadoopVol) must always be running when you are running Hadoop Jobs on other volumes. This is because the user directories for all the deployed Hadoop processes are stored on this volume. For example, if you have created and enabled 3 volumes for use with Hadoop (HadoopVol, MyVolume1, MyVolume2) and you are running a Hadoop Job that reads from MyVolume 1 and writes to MyVolume 2, then HadoopVol must still be running.

To test your trusted storage pool, shell into the YARN Master server and navigate to the `/usr/lib/hadoop/` directory. Then `su` to one of the users you have enabled for Hadoop (such as tom) and submit a Hadoop Job:

```
# su tom
# cd /usr/lib/hadoop
# bin/hadoop jar /usr/lib/hadoop-mapreduce/hadoop-mapreduce-examples-
2.4.0.2.1.7.0-784.jar teragen 1000 in
```

TeraGen only generates data. TeraSort reads and sorts the output of TeraGen. In order to fully test the cluster is operational, one needs to run TeraSort as well.

```
# bin/hadoop jar /usr/lib/hadoop-mapreduce/hadoop-mapreduce-examples-
2.4.0.2.1.7.0-784.jar terasort in out
```

For more information on using specific components within the Hadoop Ecosystem, see *Chapter 2. Understanding the Hadoop Ecosystem* in the *Hortonworks Data Platform* documentation.

8.6. TROUBLESHOOTING

This section describes the most common troubleshooting scenarios related to Hadoop and Red Hat Storage integration.

Deployment of HDP 2.1 on an LDAP enabled cluster fails with "Execution of 'groupmod hadoop' returned 10. groupmod: group 'hadoop' does not exist in /etc/group"

This is due to a bug caused by Ambari expecting a local hadoop group on an LDAP enabled cluster. Due to the fact the users and groups are centrally managed with LDAP, Ambari is not able to find the group. In order to resolve this issue:

1. Shell into the Ambari Server and navigate to `/var/lib/ambari-server/resources/scripts`
2. Replace the `$AMBARI-SERVER-FQDN` with the FQDN of your Ambari Server and the `$AMBARI-CLUSTER-NAME` with the cluster name that you specified for your cluster within Ambari and run the following command:

```
./configs.sh set $AMBARI-SERVER-FQDN $AMBARI-CLUSTER-NAME global
ignore_groupsusers_create "true"
```

3. In the Ambari console, click **Retry** in the **Cluster Installation Wizard**.

The WebHCAT service does not start

This is due to a permissions bug in WebHCAT. In order to start the service, it must be restarted multiple times and requires several file permissions to be changed. To resolve this issue, begin by starting the service. After each start attempt, WebHCAT will attempt to copy a different jar with root permissions. Every time it does this you need to `chmod 755` the jar file in `/mnt/glusterfs/HadoopVolumeName/apps/webhcat`. The three files it copies to this directory are **hadoop-streaming-2.4.0.2.1.5.0-648.jar**, **HDP-webhcat/hive.tar.gz** and **HDP-webhcat/pig.tar.gz**. After you have set the permissions on all three files, the service will start and be operational on the fourth attempt.

Exception stating that “job.jar changed on src file system” or “job.xml changed on src file system”.

This error occurs if the clocks are not synchronized across the trusted storage pool. The time in all the servers must be uniform in the trusted storage pool. It is recommended to set up a NTP (Network Time Protocol) service to keep the bricks' time synchronized, and avoid out-of-time synchronization effects.

For more information on configuring NTP, see https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/6/html/Migration_Planning_Guide/sect-Migration_Guide-Networking-NTP.html

While running a Hadoop job, if *FileNotFoundException* exception is displayed with *jobtoken does not exist* message:

This error occurs when the user IDs (UID) and group IDs (GID) are not consistent across the trusted storage pool. For example, user "tom" has a UID of 1002 on server1, but on server2, the user *tom* has a UID of 1003. The simplest and recommended approach is to leverage LDAP authentication to resolve this issue. After creating the necessary users and groups on an LDAP server, the servers within the trusted storage pool can be configured to use the LDAP server for authentication. For more information on configuring authentication, see *Chapter 12. Configuring Authentication of Red Hat Enterprise Linux 6 Deployment Guide*.

CHAPTER 9. SETTING UP SOFTWARE UPDATES

Red Hat strongly recommends you update your Red Hat Storage software regularly with the latest security patches and upgrades. Associate your system with a content server to update existing content or to install new content. This ensures that your system is up-to-date with security updates and upgrades.

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



NOTE

- Asynchronous errata update releases of Red Hat Storage include all fixes that were released asynchronously since the last release as a cumulative update.
- When there are large number of snapshots, ensure to deactivate the snapshots before performing an upgrade. The snapshots can be activated after the upgrade is complete. For more information, see *Chapter 4.1 Starting and Stopping the glusterd service* in the *Red Hat Storage 3 Administration Guide*.

9.1. UPDATING RED HAT STORAGE IN THE OFFLINE MODE

If you have a distributed volume then you must opt for an offline upgrade. Red Hat Storage supports in-service software upgrade from Red Hat Storage 2.1 only for replicate and distributed-replicate volume. For more information about in-service software upgrade, see [Section 9.2.3, “In-Service Software Upgrade”](#)



IMPORTANT

Offline upgrade results in a downtime as the volume is offline during upgrade.

If the volumes are thickly provisioned, then migrate the volume to a thinly provisioned volume. To migrate the volume to a thinly provisioned volume in the offline mode, perform the following steps:

1. Take a backup of the bricks using reliable backup solution.
2. Delete this Logical Volume (LV) and recreate a new thinly provisioned LV. For more information, see https://access.redhat.com/site/documentation/en-US/Red_Hat_Enterprise_Linux/6/html/Logical_Volume_Manager_Administration/thinprovisioned_v
3. Restore the content to the newly created thinly provisioned LV.

To update Red Hat Storage in the offline mode, execute the following steps:


```
# yum update
```

The update process automatically restarts the **glusterd** management daemon. The glusterfs server processes, **glusterfsd** is not restarted by default since restarting this daemon affects the active read and write operations.

After all the nodes in the cluster are updated, the volume must be restarted for the changes to be applied. Red Hat recommends that you restart the system when there are no active read and write operations running on the cluster.

To restart the volume, run the following commands:

```
# gluster volume stop volname  
# gluster volume start volname
```



IMPORTANT

After upgrading all the nodes, ensure to bump up the op version of the cluster by executing the following command:

```
# gluster volume set all cluster.op-version 30000
```

9.2. IN-SERVICE SOFTWARE UPGRADE TO UPGRADE FROM RED HAT STORAGE 2.1 TO RED HAT STORAGE 3.0

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

In-service software upgrade procedure is supported only from Red Hat Storage 2.1 to subsequent updates. To upgrade to Red Hat Storage 3.0 ensure you are on the immediate preceding update before proceeding with the following steps.

9.2.1. Pre-upgrade Tasks

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

9.2.1.1. Upgrade Requirements for Red Hat Storage 3.0

The following are the upgrade requirements to upgrade to Red Hat Storage 3.0 from the latest preceding update:

- In-service software upgrade is supported only for nodes with replicate and distributed replicate volumes.
- Each brick must be independent thinly provisioned logical volume(LV).
- The Logical Volume which contains the brick must not be used for any other purpose.
- Only linear LVM is supported with Red Hat Storage 3.0. For more information, see https://access.redhat.com/site/documentation/en-US/Red_Hat_Enterprise_Linux/4/html-single/Cluster_Logical_Volume_Manager/#lv_overview
- **Recommended Setup**

In addition to the following list, you must ensure to read *Chapter 9 Configuring Red Hat Storage for Enhancing Performance* in the *Red Hat Storage 3.0 Administration Guide* for enhancing performance:

- For each brick, create a dedicated thin pool that contains the brick of the volume and its (thin) brick snapshots. With the current thinly provisioned volume design, avoid placing the bricks of different gluster volumes in the same thin pool.
- The recommended thin pool chunk size is 256KB. There might be exceptions to this in cases where we have a detailed information of the customer's workload.
- The recommended pool metadata size is 0.1% of the thin pool size for a chunk size of 1MB or larger. In special cases, where we recommend a chunk size less than 256KB, use a pool metadata size of 0.5% of thin pool size.
- When server-side quorum is enabled, ensure that bringing one node down does not violate server-side quorum. Add dummy peers to ensure the server-side quorum will not be violated until the completion of rolling upgrade using the following command:

```
# gluster peer probe DummyNodeName
```

For Example 1

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 which 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.

For Example 2

In a three node cluster, if the server-side quorum percentage is set to 77%, then 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.

- If the client-side quorum is enabled then, run the following command to disable the client-side quorum:

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



NOTE

When the client-side quorum is disabled, there are chances that the files might go into split-brain.

- If there are any geo-replication sessions running between the master and slave, then stop this session by executing the following command:

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

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

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

To subscribe to the channels, run the following command:

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

9.2.1.2. Restrictions for In-Service Software Upgrade

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

- Do not perform in-service software upgrade when the I/O or load is high on the Red Hat Storage server.
- Do not perform any volume operations on the Red Hat Storage server
- Do not change the hardware configurations.
- Do not run mixed versions of Red Hat Storage for an extended period of time. For example, do not have a mixed environment of Red Hat Storage 2.1 and Red Hat Storage 2.1 Update 1 for a prolonged time.
- Do not combine different upgrade methods.
- It is not recommended to use in-service software upgrade for migrating to thinly provisioned volumes, but to use offline upgrade scenario instead. For more information see, [Section 9.1, “Updating Red Hat Storage in the Offline Mode”](#)

9.2.1.3. Configuring repo for Upgrading using ISO

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



NOTE

Upgrading Red Hat Storage using ISO can be performed only from the previous release.

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 RHSS-2.1U1-RC3-20131122.0-RHS-x86_64-DVD1.iso /mnt
```

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

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

3. Add the following information to the repo file:

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

9.2.1.4. Preparing and Monitoring the Upgrade Activity

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

- Check the peer status using the following command:

```
# gluster peer status
```

For example:

```
# gluster peer status
Number of Peers: 2

Hostname: 10.70.42.237
Uuid: 04320de4-dc17-48ff-9ec1-557499120a43
State: Peer in Cluster (Connected)

Hostname: 10.70.43.148
Uuid: 58fc51ab-3ad9-41e1-bb6a-3efd4591c297
State: Peer in Cluster (Connected)
```

- Check the volume status using the following command:

```
# gluster volume status
```

For example:

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

Gluster process                                     Port
Online Pid
-----
Brick 10.70.43.198:/brick/r2_0                       49152  Y
32259
Brick 10.70.42.237:/brick/r2_1                       49152  Y
25266
Brick 10.70.43.148:/brick/r2_2                       49154  Y
2857
Brick 10.70.43.198:/brick/r2_3                       49153  Y
32270
NFS Server on localhost                             2049   Y
25280
Self-heal Daemon on localhost                       N/A    Y
25284
NFS Server on 10.70.43.148                           2049   Y
2871
Self-heal Daemon on 10.70.43.148                     N/A    Y
2875
NFS Server on 10.70.43.198                           2049   Y
32284
Self-heal Daemon on 10.70.43.198                     N/A    Y
32288

Task Status of Volume r2
```

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

- Check the rebalance status using the following command:

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

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

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

The following example shows a self-heal completion:

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

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

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

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

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

9.2.2. Upgrade Process with Service Impact

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

SWIFT

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

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


NFS

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

Samba / CTDB

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

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

Distribute Volume

In-service software upgrade is not supported for distributed volume. In case you have a distributed volume in the cluster, stop that volume using the following command:

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

Start the volume after in-service software upgrade is complete using the following command:

```
# gluster volume start <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.

9.2.3. In-Service Software Upgrade

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

1. Back up the following configuration directory and files on the backup directory:

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

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



NOTE

- o If you have a CTDB environment, then to upgrade to Red Hat Storage 3.0, see [Section 9.2.4.1, “In-Service Software Upgrade for a CTDB Setup”](#).

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

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

3. To check the system's current subscription status run the following command:

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

4. Install the required packages using the following command:

```
# yum install subscription-manager-migration
# yum install subscription-manager-migration-data
```

5. Execute the following command to migrate from Red Hat Network Classic to Red Hat Subscription Manager

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

6. To enable the Red Hat Storage 3.0 repos, execute the following command:

```
# migrate-rhs-classic-to-rhsm --upgrade --version 3
```

7. To verify if the migration from Red Hat Network Classic to Red Hat Subscription Manager is successful, execute the following command:

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

8. Update the server using the following command:

```
# yum update
```

9. If the volumes are thickly provisioned, then perform the following steps to migrate to thinly provisioned volumes:



NOTE

Migrating from thickly provisioned volume to thinly provisioned volume during in-service software upgrade takes a significant amount of time based on the data you have in the bricks. You must migrate *only* if you plan on using snapshots for your existing environment and plan to be online during the upgrade. If you do not plan to use snapshots, you can skip the migration steps from LVM to thinp. However, if you plan to use snapshots on your existing environment, the offline method to upgrade is recommended. For more information regarding offline upgrade, see [Section 9.1, “Updating Red Hat Storage in the Offline Mode”](#)

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

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

```
# umount mount point
```

For example:

```
# umount /dev/RHS_vg/brick1
```

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

```
# lvremove logical_volume_name
```

For example:

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

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

```
# vgremove -ff volume_group_name
```

For example:

```
vgremove -ff RHS_vg
```

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

```
# pvremove -ff physical_volume
```

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

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

For more information refer *Section 12.1 Prerequisites* in the *Red Hat Storage 3 Administration Guide*

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

```
# vgcreate volume_group_name disk
```

For example:

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

7. Create a thinpool using the following command:

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

For example:

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

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

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

For example:

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

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

```
mkfs.xfs -i size=512 thin pool device
```

For example:

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

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

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

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

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

```
# chkconfig glusterd off
```

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

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

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

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

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

3. Delete the directory and set the extended attributes.

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

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

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

```
#getfattr -d -m. -e hex /dev/RHS_vg/brick2 # file:
/dev/RHS_vg/brick2

security.selinux=0x756e636f6e66696e65645f753a6f626a6563745f723a66
696c655f743a733000
    trusted.afr.r2-client-0=0x000000000000000000000000
    trusted.afr.r2-client-1=0x00000000000000000300000002
    trusted.gfid=0x00000000000000000000000000000001
    trusted.glusterfs.dht=0x00000001000000000000000007ffffffe
    trusted.glusterfs.volume-
id=0xde822e25ebd049ea83bfaa3c4be2b440
```

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

```
# service glusterd start
```

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

```
# chkconfig glusterd on
```

15. Start self-heal on the volume.

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

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

```
# gluster --version
```

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

```
# gluster volume status
```

For example:

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

Gluster process                                     Port
Online Pid
-----
-----
Brick 10.70.43.198:/brick/r2_0                       49152  Y
32259
```

```

Brick 10.70.42.237:/brick/r2_1          49152  Y
25266
Brick 10.70.43.148:/brick/r2_2        49154  Y
2857
Brick 10.70.43.198:/brick/r2_3        49153  Y
32270
NFS Server on localhost                2049   Y
25280
Self-heal Daemon on localhost          N/A    Y
25284
NFS Server on 10.70.43.148             2049   Y
2871
Self-heal Daemon on 10.70.43.148      N/A    Y
2875
NFS Server on 10.70.43.198            2049   Y
32284
Self-heal Daemon on 10.70.43.198      N/A    Y
32288

```

```

Task Status of Volume r2
-----
-----

```

```

There are no active volume tasks

```

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

```

# gluster volume heal volname info

```

The following example shows self heal completion:

```

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

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

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

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

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

```

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



NOTE

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

- After upgrading all the nodes, ensure to bump up the op-version of the cluster by executing the following command:

```
# gluster volume set all cluster.op-version 30000
```



NOTE

If you want to enable Snapshot, see *Section 12.4. Troubleshooting* in the *Red Hat Storage 3 Administration Guide*.

- 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
```

- If the geo-replication session between master and slave was disabled before upgrade, then restart the session by executing the following command:

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

9.2.4. Special Consideration for In-Service Software Upgrade

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

9.2.4.1. In-Service Software Upgrade for a CTDB Setup

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

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

```
# chkconfig ctdb off
```

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

```
# service ctdb stop
```

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

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

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

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

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

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

5. Update the server using the following command:

```
# yum update
```

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

```
# chkconfig glusterd off
```

7. Reboot the server.

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

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

9. 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
```

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

```
# chkconfig glusterd on
```

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

```
# chkconfig ctdb on
```

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

```
# service glusterd start
```

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

```
# mount -a
```

14. Start the CTDB service using the following command:

```
# service ctdb start
```

15. 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 Storage server, upgrade the CTDB package by executing the following steps:



NOTE

- Upgrading CTDB on all the nodes must be done simultaneously to avoid any data corruption.
- The following steps have to be performed only when upgrading CTDB from CTDB 1.x to CTDB 2.x.

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

```
# service ctdb stop
```

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

- 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 remove the existing CTDB package execute the following command:

```
# yum remove ctdb
```

5. To install CTDB, execute the following command:

```
# yum install ctdb2.5
```

For more information about configuring CTDB on a Red Hat Storage server, see *Section 7.5.1 Setting Up CTDB* in the *Red Hat Storage 3 Administration Guide*

9.2.4.2. Verifying In-Service Software Upgrade

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

```
# gluster --version
```

9.2.4.3. Upgrading the Native Client

All the clients must be of same version. Red Hat strongly recommends you to upgrade the servers before you upgrading the clients. For more information regarding installing and upgrading native client refer *Section 9.2 Native Client* in the *Red Hat Storage Administration Guide*.

9.3. UPDATING YOUR CURRENT SYSTEM

To update your system to the latest update version of Red Hat Storage 3, follow these steps. The following steps have to be performed on each node of the replica pair:



NOTE

Ensure that the system is registered to the Red Hat Network. For more information refer to, [Chapter 5, Subscribing to the Red Hat Storage Server Channels](#)



WARNING

- The Samba version 3 is being deprecated from Red Hat 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.
- 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.

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

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

1. For Red Hat Storage 3.0.4 or later, if you require Samba, then enable the following channel:

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

2. Stop the CTDB and SMB services across all nodes in the Samba cluster using the following command. This is because different versions of Samba cannot run in the same Samba cluster.

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

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

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

2. Update the server using the following command:

```
# yum update
```

3. Reboot the server if a kernel update was included as part of the update process in the previous step.



NOTE

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

```
# service glusterd start
```

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

```
# service ctdb start
```

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

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

4. Start self-heal on the volume.

```
# gluster volume heal volname full
```

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

```
# gluster --version
```

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

```
# gluster volume status
```

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

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

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



NOTE

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

CHAPTER 10. MANAGING THE GLUSTERD SERVICE

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

Use Red Hat 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.

10.1. MANUALLY STARTING AND STOPPING GLUSTERD

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

- Manually start **glusterd** as follows:

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

or

```
# service glusterd start
```

- Manually stop **glusterd** as follows:

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

or

```
# service glusterd stop
```

CHAPTER 11. USING THE GLUSTER COMMAND LINE INTERFACE

The Gluster command line interface (CLI) simplifies configuration and management of the storage environment. The Gluster CLI is similar to the LVM (Logical Volume Manager) CLI or the ZFS CLI, but operates across multiple storage servers. The Gluster CLI can be used when volumes are mounted (active) and not mounted (inactive). Red Hat 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 Storage Server by either invoking the commands or running the Gluster CLI in interactive mode. The `gluster` command can be remotely used via SSH.

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

```
# gluster peer COMMAND
```

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

```
# gluster peer status
```

Gluster CLI Interactive Mode

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

```
# gluster
```

If successful, the prompt changes to the following:

```
gluster>
```

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

```
gluster> COMMAND
```

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

1. Start Gluster CLI's interactive mode:

```
# gluster
```

2. Request the peer server status:

```
█ gluster> status
```

3. The peer server status displays.

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

1. Start Gluster CLI's interactive mode:

```
█ # gluster
```

2. Request the help options:

```
█ gluster> help
```

3. A list of gluster commands and options displays.

PART I. APPENDIX

CHAPTER 12. DEPLOYING THE HORTONWORKS DATA PLATFORM 2.0.6 ON RED HAT STORAGE

Red Hat Storage provides compatibility for Apache Hadoop and it uses the standard file system APIs available in Hadoop to provide a new storage option for Hadoop deployments. Red Hat has created a Hadoop File System plug-in that enables Hadoop Distributions to run on Red Hat Storage.

12.1. PREREQUISITES

Before you begin installation, you must establish the basic infrastructure required to enable Hadoop to run on Red Hat Storage.

12.1.1. Supported Versions

Red Hat Storage 3.0 can be integrated successfully with Hortonworks Data Platform (HDP) version 2.0.6.

12.1.2. Software and Hardware Requirements

You must ensure that all the servers used in this environment meet the following requirements:

- Must have at least the following hardware specification:
 - 2 x 2GHz 4 core processors
 - 32 GB RAM
 - 500 GB of storage capacity
 - 1 x 1GbE NIC
- Must have iptables disabled.
- Must use fully qualified domain names (FQDN). For example rhs-1.server.com is acceptable, but rhs-1 is not allowed.
- Must have access to a DNS server and must be able to use DNS for FQDN resolution.
- Must have the following users and groups available on all the servers.

User	Group
yarn	hadoop
mapred	hadoop
hive	hadoop
hcat	hadoop
ambari-qa	hadoop

The specific UIDs and GIDs for the respective users and groups are up to the Administrator of the trusted storage pool, but they must be consistent across the trusted storage pool. For example, if the "hadoop" user has a UID as 591 on one server, the hadoop user must have UID as 591 on all other servers. This can be quite a lot of work to manage using Local Authentication and it is common and acceptable to install a central authentication solution such as LDAP or Active Directory for your cluster, so that users and groups can be easily managed in one place. However, to use local authentication, you can run the script below on each server to create the users and groups and ensure they are consistent across the cluster:

```
groupadd hadoop -g 590; useradd -u 591 mapred -g hadoop; useradd -u 592 yarn -g hadoop; useradd -u 594 hcat -g hadoop; useradd -u 595 hive -g hadoop; useradd -u 596 ambari-qa -g hadoop
```

12.1.3. Existing Red Hat Storage Trusted Storage Pool

If you have an existing Red Hat Storage trusted storage pool, you need to add two additional servers to run the Hortonworks Ambari Management Services and the YARN Master Services, respectively. For more information on recommended deployment topologies, see *Administering the Hortonworks Data Platform on Red Hat Storage* chapter in *Red Hat Storage Administration Guide*.

For information on expanding your trusted storage pool by adding servers, see section *Expanding Volumes* in the *Red Hat Storage 3.0 Administration Guide*.



NOTE

The supported volume configuration for Hadoop is Distributed Replicated volume with replica count 2.

12.1.4. New Red Hat Storage Trusted Storage Pool

You must create a Red Hat Storage trusted storage pool with at least four bricks. The servers on which these bricks reside must have the Red Hat Storage 3.0 installed on them and must have at least one RAID 6 block device per server. With these bricks, you must create a Distributed Replicated volume.

For more information on installing Red Hat Storage 3.0, see [Chapter 4, Installing Red Hat Storage](#) or for upgrading to Red Hat Storage 3.0, see [Chapter 6, Upgrading Red Hat Storage](#).

Red Hat recommends that you have an additional two servers set aside to run the Hortonworks Ambari Management Services and the YARN Master Services, respectively. Alternate deployment topologies are also possible, for more information on various supported deployment topologies, see *Administering the Hortonworks Data Platform on Red Hat Storage* chapter in *Red Hat Storage Administration Guide*.

For information on expanding your trusted storage pool by adding servers, see section *Expanding Volumes* in the *Red Hat Storage 3.0 Administration Guide*.



NOTE

The supported volume configuration for Hadoop is Distributed Replicated volume with replica count 2.

12.1.5. Red Hat Storage Server Requirements

You must install Red Hat Storage Server 3 on this server. While installing the server, you must ensure to specify a fully qualified domain name (FQDN). A hostname alone will not meet the requirements for the Horton Data Platform Ambari deployment tool.

You must also enable the `rhs-big-data-3-for-rhel-6-server-rpms` channel on this server.

- If you have registered your machine using Red Hat Subscription Manager, enable the channel by running the following command:

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

- 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-rhs-bigdata-3
```

12.1.6. Hortonworks Ambari Server Requirements

You must install Red Hat Enterprise Linux 6.5 on the servers. You can also choose to install Red Hat Storage Console on this server as well, but this is optional. This allows all aspects of the Red Hat Storage trusted pool to be managed from a single server. While installing the server, you must ensure to specify a fully qualified domain name (FQDN). A hostname alone will not meet the requirements for the Horton Data Platform Ambari deployment tool. **It is mandatory to setup a passwordless-SSH connection from the Ambari Server to all other servers within the trusted storage pool.** Instructions for installing and configuring Hortonworks Ambari is provided in the further sections of this chapter.

If the Hortonworks Ambari server is installed on a different node than Red Hat Storage Server, you must also enable the `rhel-6-server-rh-common-rpms` channel on this server.

- If you have registered your machine using Red Hat Subscription Manager, enable the channel by running the following command:

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

- 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-rh-common-6
```



WARNING

Red Hat Storage Console enables Nagios Alerting for Red Hat Storage. The Nagios Client libraries are shipped with Red Hat Storage and are on each Red Hat Storage Server. This causes a conflict with the Nagios System that is bundled with the Hortonworks Data Platform (HDP). As such, using HDP 2.0.6 to deploy and manage Nagios is not supported.



NOTE

If you are using one of the condensed deployment topologies listed in the *Administration Guide* and you have elected to place the Ambari Management server on the same node as a Red Hat Storage Server, you must only enable the **rhs-big-data-3-for-rhel-6-server-rpms** channel on that server.

- If you have registered your machine using Red Hat Subscription Manager, enable the channel by running the following command:

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

- 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-rhs-bigdata-3
```

12.1.7. YARN Master Server Requirements

You must install the Red Hat Enterprise Linux 6.5 on this server. While installing the server, you must ensure to specify a fully qualified domain name (FQDN). A hostname alone will not meet the requirements for the Horton Data Platform Ambari deployment tool.

If the YARN Master server is installed on a different node than Red Hat Storage Server, you must also enable the **rhel-6-server-rh-common-rpms** and **rhel-6-server-rhs-client-1-rpms** channels on the YARN server.

- If you have registered your machine using Red Hat Subscription Manager, enable the channel by running the following command:

```
# subscription-manager repos --enable=rhel-6-server-rh-common-rpms /  
--enable=rhel-6-server-rhs-client-1-rpms
```

- 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-rh-common-6  
# rhn-channel --add --channel rhel-x86_64-server-rhsclient-6
```



NOTE

If you are using one of the condensed deployment topologies listed in the *Administration Guide* and you have elected to place the YARN Master server on the same node as a Red Hat Storage Server, you must only enable the **rhs-big-data-3-for-rhel-6-server-rpms** channel on that server.

- If you have registered your machine using Red Hat Subscription Manager, enable the channel by running the following command:

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

- 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-rhs-bigdata-3
```

12.2. INSTALLING THE HADOOP FILESYSTEM PLUGIN FOR RED HAT STORAGE

12.2.1. Adding the Hadoop Installer for Red Hat Storage

You must have the big-data channel added and the hadoop components installed on all the servers to use the Hadoop feature on Red Hat Storage. Run the following command on the Ambari Management Server, the YARN Master Server and all the servers within the Red Hat Storage trusted storage pool:

```
# yum install rhs-hadoop rhs-hadoop-install
```

On the YARN Master Server

The YARN Master Server is required to FUSE Mount all Red Hat Storage Volumes that is used with Hadoop. It must have the Red Hat Storage Client Channel enabled so that the *setup_cluster* script can install the Red Hat Storage Client Libraries on it.

- If you have registered your machine using Red Hat Subscription Manager, enable the channel by running the following command:

```
# subscription-manager repos --enable=rhel-6-server-rhs-client-1-rpms
```

- 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-rhsclient-6
```

12.2.2. Configuring the Trusted Storage Pool for use with Hadoop

Red Hat Storage provides a series of utility scripts that allows you to quickly prepare Red Hat Storage for use with Hadoop, and install the Ambari Management Server. You must first run the Hadoop cluster configuration initial script to install the Ambari Management Server, prepare the YARN Master Server to

host the Resource Manager and Job History Server services for Red Hat Storage and build a trusted storage pool if it does not exist.



NOTE

You must run the script given below irrespective of whether you have an existing Red Hat Storage trusted storage pool or not.

To run the Hadoop configuration initial script:

1. Open the terminal window of the server designated to be the Ambari Management Server and navigate to the `/usr/share/rhs-hadoop-install/` directory.
2. Run the hadoop cluster configuration script as given below:

```
setup_cluster.sh [-y] [--hadoop-mgmt-node <node>] [--yarn-master
<node>] <node-list-spec>
```

where `<node-list-spec>` is


```
<node1>:<brickmnt1>:<blkdev1> <node2>[:<brickmnt2>][:<blkdev2>]
[<node3>[:<brickmnt3>][:<blkdev3>]] ... [<nodeN>[:<brickmntN>][:
<blkdevN>]]
```

where

- o `<brickmnt>` is the name of the XFS mount for the above `<blkdev>`, for example, `/mnt/brick1` or `/external/HadoopBrick`. When a Red Hat Storage volume is created its bricks has the volume name appended, so `<brickmnt>` is a prefix for the volume's bricks. Example: If a new volume is named `HadoopVol1` then its brick list would be:
`<node>:/mnt/brick1/HadoopVol1` or
`<node>:/external/HadoopBrick/HadoopVol1`.
- o `<blkdev>` is the name of a Logical Volume device path, for example, `/dev/VG1/LV1` or `/dev/mapper/VG1-LV1`. Since LVM is a prerequisite for Red Hat Storage, the `<blkdev>` is not expected to be a raw block path, such as `/dev/sdb`.

Given below is an example of running `setup_cluster.sh` script on a the YARN Master server and four Red Hat Storage Nodes which has the same logical volume and mount point intended to be used as a Red Hat Storage Brick.

```
./setup_cluster.sh --yarn-master yarn.hdp rhs-
1.hdp:/mnt/brick1:/dev/rhs_vg1/rhs_lv1 rhs-2.hdp rhs-3.hdp rhs-4.hdp
```



NOTE

If a brick mount is omitted, the brick mount of the first node is used and if one block device is omitted, the block device of the first node is used.

12.2.3. Creating Volumes for use with Hadoop

**NOTE**

To use an existing Red Hat Storage Volume with Hadoop, skip this section and continue with the section *Adding the User Directories for the Hadoop Processes on the Red Hat Storage Volume*.

Whether you have a new or existing Red Hat Storage trusted storage pool, to create a volume for use with Hadoop, the volume need to be created in such a way as to support Hadoop workloads. The supported volume configuration for Hadoop is Distributed Replicated volume with replica count 2. You must not name the Hadoop enabled Red Hat Storage volume as **hadoop** or **mapredlocal**.

Run the script given below to create new volumes that you intend to use with Hadoop. The script provides the necessary configuration parameters to the volume as well as updates the Hadoop Configuration to make the volume accessible to Hadoop.

1. Open the terminal window of the server designated to be the Ambari Management Server and navigate to the `/usr/share/rhs-hadoop-install/` directory.
2. Run the hadoop cluster configuration script as given below:

```
create_vol.sh [-y] <volName> <volMountPrefix> <node-list>
```

where

- o `<node-list>` is: `<node1>:<brickmnt> <node2>[:<brickmnt2>] <node3>[:<brickmnt3>] ... [<nodeN>[:<brickmntN>]`
- o `<brickmnt>` is the name of the XFS mount for the block devices used by the above nodes, for example, `/mnt/brick1` or `/external/HadoopBrick`. When a RHS volume is created its bricks will have the volume name appended, so `<brickmnt>` is a prefix for the volume's bricks. For example, if a new volume is named **HadoopVol** then its brick list would be: `<node>:/mnt/brick1/HadoopVol` or `<node>:/external/HadoopBrick/HadoopVol`.

**NOTE**

The `node-list` for `create_vol.sh` is similar to the `node-list-spec` used by `setup_cluster.sh` except that a block device is not specified in `create_vol`.

Given below is an example on how to create a volume named *HadoopVol*, using 4 Red Hat Storage Servers, each with the same brick mount and mount the volume on `/mnt/glusterfs`

```
./create_vol.sh HadoopVol /mnt/glusterfs rhs-1.hdp:/mnt/brick1 rhs-2.hdp rhs-3.hdp rhs-4.hdp
```

12.2.4. Adding the User Directories for the Hadoop Processes on the Red Hat Storage Volume

After creating the volume, you need to setup the user directories for all the Hadoop ecosystem component users that you created in the prerequisites section. This is required for completing the Ambari distribution successfully.

**NOTE**

Perform the steps given below only when the volume is created and enabled to be used with Hadoop.


Open the terminal window of the Red Hat Storage server within the trusted storage pool and run the following commands:

```
# mkdir /mnt/glusterfs/HadoopVol/user/mapred
# mkdir /mnt/glusterfs/HadoopVol/user/yarn
# mkdir /mnt/glusterfs/HadoopVol/user/hcat
# mkdir /mnt/glusterfs/HadoopVol/user/hive
# mkdir /mnt/glusterfs/HadoopVol/user/ambari-qa
```

```
# chown ambari-qa:hadoop /mnt/glusterfs/HadoopVol/user/ambari-qa
# chown hive:hadoop /mnt/glusterfs/HadoopVol/user/hive
# chown hcat:hadoop /mnt/glusterfs/HadoopVol/user/hcat
# chown yarn:hadoop /mnt/glusterfs/HadoopVol/user/yarn
# chown mapred:hadoop /mnt/glusterfs/HadoopVol/user/mapred
```

12.2.5. Deploying and Configuring the HDP 2.0.6 Stack on Red Hat Storage using Ambari Manager

Perform the following steps to deploy and configure HDP stack on Red Hat Storage:

**IMPORTANT**

This section describes how to deploy HDP on Red Hat Storage. Selecting **HDFS** as the storage selection in the HDP 2.0.6.GlusterFS stack is not supported. If you want to deploy HDFS, then you must select the HDP 2.0.6 stack (not HDP 2.0.6.GlusterFS) and follow the instructions of the Hortonworks documentation.


1. Launch a web browser and enter **http://hostname:8080** in the URL by replacing *hostname* with the hostname of your Ambari Management Server.

**NOTE**

If the Ambari Console fails to load in the browser, it is usually because iptables is still running. Stop iptables by opening a terminal window and run **service iptables stop** command.

2. Enter **admin** and **admin** for the username and password.
3. Assign a name to your cluster, such as **MyCluster**.
4. Select the **HDP 2.0.6.GlusterFS Stack** (if not already selected by default) and click **Next**.
5. On the **Install Options** screen:
 1. For **Target Hosts**, add the YARN server and all the nodes in the trusted storage pool.
 2. Select **Perform manual registrations on hosts and do not use SSH** option.

3. Accept any warnings you may see and click **Register and Confirm** button.
4. Click **OK** on **Before you proceed warning** warning. The Ambari Agents have all been installed for you during the `setup_cluster.sh` script.
6. For **Confirm Hosts**, the progress must show as green for all the hosts. Click **Next** and ignore the **Host Check** warning.
7. For **Choose Services**, unselect HDFS and as a minimum select GlusterFS, Ganglia, YARN+MapReduce2 and ZooKeeper.



NOTE

- Do not select the Nagios service, as it is not supported. For more information, see subsection *21.1. Deployment Scenarios* of chapter *21. Administering the Hortonworks Data Platform on Red Hat Storage* in the *Red Hat Storage 3.0 Administration Guide*.
- The use of HBase has not been extensively tested and is not yet supported.
- This section describes how to deploy HDP on Red Hat Storage. Selecting **HDFS** as the storage selection in the HDP 2.0.6.GlusterFS stack is not supported. If users wish to deploy HDFS, then they must select the HDP 2.0.6 stack (not HDP 2.0.6.GlusterFS) and follow the instructions in the Hortonworks documentation.

8. For **Assign Masters**, set all the services to your designated YARN Master Server. For ZooKeeper, select at least 3 separate nodes within your cluster.
9. For **Assign Slaves and Clients**, select all the nodes as **NodeManagers** except the YARN Master Server. You must also ensure to click the **Client** checkbox for each node.
10. On the **Customize Services** screen:
 1. Click **YARN** tab, scroll down to the `yarn.nodemanager.log-dirs` and `yarn.nodemanager.local-dirs` properties and remove any entries that begin with `/mnt/glusterfs/`.
 2. Click **MapReduce2** tab, scroll down to the **Advanced** section, and modify the following property:

Key	Value
<code>yarn.app.mapreduce.am.staging-dir</code>	<code>glusterfs:///user</code>

3. Click **MapReduce2** tab, scroll down to the bottom, and under the custom `mapred-site.xml`, add the following four custom properties and then click on the **Next** button:

Key	Value
<code>mapred.healthChecker.script.path</code>	<code>glusterfs:///mapred/jobstatus</code>

Key	Value
mapred.job.tracker.history.completed.location	glusterfs:///mapred/history/done
mapred.system.dir	glusterfs:///mapred/system
mapreduce.jobtracker.staging.root.dir	glusterfs:///user

- Review other tabs that are highlighted in red. These require you to enter additional information, such as passwords for the respective services.
- Review your configuration and then click **Deploy** button. Once the deployment is complete, it will state that the deployment is 100% complete and the progress bars will be colored in Orange.



NOTE

The deployment process is susceptible to network and bandwidth issues. If the deployment fails, try clicking "Retry" to attempt the deployment again. This often resolves the issue.

- Click **Next** to proceed to the Ambari Dashboard. Select the YARN service on the top left and click **Stop-All**. Do not click **Start-All** until you perform the steps in section [Section 12.2.7, "Configuring the Linux Container Executor"](#).

12.2.6. Enabling Existing Volumes for use with Hadoop



IMPORTANT

This section is mandatory for every volume you intend to use with Hadoop. It is not sufficient to run the `create_vol.sh` script, you must follow the steps listed in this section as well.

If you have an existing Red Hat Storage trusted storage pool with volumes that contain data that you would like to analyze with Hadoop, the volumes need to be configured to support Hadoop workloads. Execute the script given below on every volume that you intend to use with Hadoop. The script provides the necessary configuration parameters for the volume and updates the Hadoop Configuration to make the volume accessible to Hadoop.



NOTE

The supported volume configuration for Hadoop is Distributed Replicated volume with replica count 2.

- Open the terminal window of the server designated to be the Ambari Management Server and navigate to the `/usr/share/rhs-hadoop-install/` directory.
- Run the Hadoop Trusted Storage pool configuration script as given below:

```
# enable_vol.sh [-y] [--hadoop-mgmt-node <node>] [--user <admin-user>] [--pass <admin-password>] [--port <mgmt-port-num>] [--yarn-
```

```
master <node>] [--rhs-node <storage-node>] <volName>
```

For Example;

```
./enable_vol.sh --yarn-master yarn.hdp --rhs-node rhs-1.hdp
HadoopVol
```



NOTE

If `--yarn-master` and/or `--rhs-node` options are omitted then the default of `localhost` (the node from which the script is being executed) is assumed. `--rhs-node` is the hostname of any of the storage nodes in the trusted storage pool. This is required to access the gluster command. Default is `localhost` and it must have gluster CLI access.

12.2.7. Configuring the Linux Container Executor

The Container Executor program used by the YARN framework defines how any **container** is launched and controlled. The Linux Container Executor sets up restricted permissions and the user/group ownership of local files and directories used by the containers such as the shared objects, jars, intermediate files, log files, and so on. Perform the following steps to configure the Linux Container Executor program:

1. In the Ambari console, click **Stop All** in the Services navigation panel. You must wait until all the services are completely stopped.
2. On each server within the Red Hat Storage trusted storage pool:
 1. Open the terminal and navigate to `/usr/share/rhs-hadoop-install/` directory:
 2. Execute the `setup_container_executor.sh` script.
3. On each server within the Red Hat Storage trusted storage pool and the YARN Master server:
 1. Open the terminal and navigate to `/etc/hadoop/conf/` directory.
 2. Replace the contents of `container-executor.cfg` file with the following:

```
yarn.nodemanager.linux-container-executor.group=hadoop
banned.users=yarn
min.user.id=1000
allowed.system.users=tom
```



NOTE

Ensure that there is no additional whitespace at the end of each line and at the end of the file. Also, `tom` is an example user. Hadoop ignores the **`allowed.system.user`** parameter, but we recommend having at least one valid user. You can modify this file on one server and then use Secure Copy (or any another approach) to copy the modified file to the same location on each server.

12.3. ADDING AND REMOVING USERS

Only users that are part of the *Hadoop* group (that was created in the prerequisites section) will be able to submit Hadoop Jobs. This can be relatively easily if you are using LDAP for authentication, but if you are not, you need to run the command given below on each server in the trusted storage pool and the YARN Master Server, for each user you add.

```
# useradd -u 1005 -g hadoop tom
```



NOTE

The UID of 1005 is arbitrary, you can specify the UID of your choice, but it must be both unique and consistent across the trusted storage pools. Also, the Hadoop Container Executor requires that all user IDs be greater than 1000. Thus, 999 is not acceptable as user ID, but 1001 is acceptable.

After adding a user who is part of the *hadoop* group, you need to create a user directory for that user within the default Red Hat Storage Hadoop Volume. The default Red Hat Storage Hadoop Volume is the first volume that was created and enabled for Hadoop and is usually called **HadoopVol1** according to the examples given in installation instructions.

Run the following commands from a server within the Red Hat Storage trusted storage pool (replacing *user_name* with the actual user name) for each user that you are adding. You must run this command only once on the default Volume. If you add subsequent volumes, you do not need to repeat this step.

```
# mkdir /mnt/glusterfs/HadoopVol1/user/<username>
# chown <username>:hadoop /mnt/glusterfs/HadoopVol1/user/<username>
# chmod 0755 /mnt/glusterfs/HadoopVol1/user/<username>
```

Removing Users

To disable a user from submitting Hadoop Jobs, remove the user from the Hadoop group.

12.4. DISABLING A VOLUME FOR USE WITH HADOOP

To keep a volume available but not accessible by Hadoop for analytics, you can disable the volume for use with Hadoop. Perform the following steps to disable the volume:

1. Open the terminal window of the server designated to be the Ambari Management Server and navigate to the `/usr/share/rhs-hadoop-install/` directory.
2. Run the Hadoop cluster configuration script as shown below:


```
disable_vol.sh [-y] [--hadoop-mgmt-node node] [--user admin-user] [-  
-pass admin-password] [--port mgmt-port-num] --yarn-master node  
volName
```

For example,

```
disable_vol.sh --hadoop-mgmt-node mgmt-node --yarn-master yarn-node  
HadoopVol1
```

12.5. VERIFYING THE CONFIGURATION

Prior to submitting any jobs, ensure that the trusted storage pool is running. Launch the Ambari Dashboard (<http://ambari-server-hostname:8080>) and select the YARN service and then click the **Start -All** button.



NOTE

Stopping and starting the services takes some time.

The default volume (usually HadoopVol) must always be running when you are running Hadoop Jobs on other volumes. This is because the user directories for all the deployed Hadoop processes are stored on this volume. For example, if you have created and enabled 3 volumes for use with Hadoop (HadoopVol, MyVolume1, MyVolume2) and you are running a Hadoop Job that reads from MyVolume 1 and writes to MyVolume 2, then HadoopVol must still be running.

To test your trusted storage pool, shell into the YARN Master server and navigate to the `/usr/lib/hadoop/` directory. Then `su` to one of the users you have enabled for Hadoop (such as tom) and submit a Hadoop Job:

```
# su tom bin/hadoop jar /usr/lib/hadoop-mapreduce/hadoop-mapreduce-
examples-2.2.0.2.0.6.0-101.jar teragen 1000 in
```

For more information on using specific components within the Hadoop Ecosystem, see *Chapter 2. Understanding the Hadoop Ecosystem* in the *Hortonworks Data Platform documentation*.

12.6. TROUBLESHOOTING

This section describes the most common troubleshooting scenarios related to Hadoop and Red Hat Storage integration.

Exception stating that “job.jar changed on src file system” or “job.xml changed on src file system”.

This error occurs if the clocks are not synchronized across the trusted storage pool. The time in all the servers must be uniform in the trusted storage pool. It is recommended to set up a NTP (Network Time Protocol) service to keep the bricks' time synchronized, and avoid out-of-time synchronization effects.

For more information on configuring NTP, see https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/6/html/Migration_Planning_Guide/sect-Migration_Guide-Networking-NTP.html

While running a Hadoop job, if `FileNotFoundException` exception is displayed with `jobtoken does not exist` message:

This error occurs when the user IDs (UID) and group IDs (GID) are not consistent across the trusted storage pool. For example, user "tom" has a UID of 1002 on server1, but on server2, the user tom has a UID of 1003. The simplest and recommended approach is to leverage LDAP authentication to resolve this issue. After creating the necessary users and groups on an LDAP server, the servers within the trusted storage pool can be configured to use the LDAP server for authentication. For more information on configuring authentication, see *Chapter 12. Configuring Authentication of Red Hat Enterprise Linux 6 Deployment Guide*.

APPENDIX A. REVISION HISTORY

Revision 3-37	Mon Apr 06 2015	Bhavana Mohan
Incorporated review comments and updated Samba specific Install/Upgrade information throughout the doc.		
Revision 3-36	Mon Mar 30 2015	Divya Muntimadugu
Added support matrix of HDP and Ambari with Red Hat Storage Server.		
Revision 3-30	Wed Mar 25 2015	Bhavana Mohan
Updated the entire document with relevant steps for SMB related Installation and Upgrade scenarios, wherever applicable. Also, added a new chapter, <i>Deploying Samba on Red Hat Storage</i>		
Revision 3-29	Mon Mar 23 2015	Divya Muntimadugu
Updated section <i>Troubleshooting</i> of chapter <i>Deploying the Hortonworks Data Platform 2.1 on Red Hat Storage</i> to fix BZ# 1184392.		
Revision 3-27	Fri Feb 13 2015	Divya Muntimadugu
Bug fix.		
Revision 3-26	Thu Jan 15 2015	Bhavana Mohan
Version for 3.0.3 release.		
Revision 3-25	Mon Nov 24 2014	Divya Muntimadugu
Added Chapter 7. <i>Deploying the Hortonworks Data Platform 2.1 on Red Hat Storage</i> .		
Revision 3-24	Thu Nov 06 2014	Bhavana Mohan
Version for 3.0.2 release.		
Revision 3-23	Wed Oct 15 2014	Bhavana Mohan
Fixed Bugs.		
Revision 3-22	Fri Oct 03 2014	Divya Muntimadugu
Version for 3.0.1 release.		
Revision 3-21	Thu Sep 25 2014	Divya Muntimadugu
Version for 3.0.1 release.		
Revision 3-20	Tue Sep 23 2014	Bhavana Mohan
Version for 3.0 GA release.		