

Gluster File System 3.2.5

Installation Guide



Red Hat Engineering Content Services

Gluster File System 3.2.5 Installation Guide Edition 1

Author Red Hat Engineering Content docfeedback@gluster.com
Services

Copyright © 2011 Red Hat Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at <http://creativecommons.org/licenses/by-sa/3.0/>. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

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

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, JBoss, MetaMatrix, 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.

All other trademarks are the property of their respective owners.

1801 Varsity Drive
Raleigh, NC27606-2072USA
Phone: +1 919 754 3700
Phone: 888 733 4281
Fax: +1 919 754 3701

This guide introduces Gluster File System, describes the system requirements, and provides information on downloading, installing, and upgrading the software in your environment.

Preface	v
1. Audience	v
2. License	v
3. Document Conventions	v
3.1. Typographic Conventions	v
3.2. Pull-quote Conventions	vii
3.3. Notes and Warnings	vii
4. Feedback	viii
1. Introducing Gluster File System	1
2. Preparing to Install GlusterFS Server	3
2.1. Preparation	3
2.2. Checking Minimum Requirements	3
2.2.1. Hardware Requirements	3
2.2.2. Operating System Requirements	3
2.2.3. File System Requirements	4
3. Installing GlusterFS Server	5
3.1. Installing GlusterFS on Red Hat Package Manager (RPM) Distributions	5
3.2. Installing GlusterFS on Debian-based Distributions	6
3.3. Installing GlusterFS from Source	7
3.4. Configuring GlusterFS to work over InfiniBand	8
3.4.1. Configuring GlusterFS for InfiniBand on RPM distribution-based systems	9
3.4.2. Configuring GlusterFS for InfiniBand on Debian distribution-based systems	9
4. Upgrading GlusterFS	11
4.1. Upgrading GlusterFS from v3.0.x or older to v3.2.x	11
4.2. Upgrading GlusterFS from v3.1.x to v3.2.x	12
4.2.1. Using RPMs on RHEL, CentOS, Fedora	12
4.2.2. Using dpkg on Debian, Ubuntu	12
4.2.3. Building from Source	12
5. Uninstalling GlusterFS	15
5.1. Uninstalling GlusterFS on Red Hat Package Manager (RPM) Distributions	15
5.2. Uninstalling GlusterFS on Debian-based Distributions	15
5.3. Uninstalling GlusterFS when Compiled from Source	16

Preface

This guide introduces Gluster File System, describes the system requirements, and provides information on downloading, installing, and upgrading the software in your environment.

1. Audience

This guide is intended for Systems Administrators interested in installing or upgrading GlusterFS.

This guide assumes that you are familiar with the Linux operating system, concepts of File System, and GlusterFS concepts.

2. License

The License information is available at http://www.redhat.com/licenses/rhel_rha_eula.html.

3. Document Conventions

This manual uses several conventions to highlight certain words and phrases and draw attention to specific pieces of information.

In PDF and paper editions, this manual uses typefaces drawn from the *Liberation Fonts*¹ set. The Liberation Fonts set is also used in HTML editions if the set is installed on your system. If not, alternative but equivalent typefaces are displayed. Note: Red Hat Enterprise Linux 5 and later includes the Liberation Fonts set by default.

3.1. Typographic Conventions

Four typographic conventions are used to call attention to specific words and phrases. These conventions, and the circumstances they apply to, are as follows.

Mono-spaced Bold

Used to highlight system input, including shell commands, file names and paths. Also used to highlight keycaps and key combinations. For example:

To see the contents of the file **my_next_bestselling_novel** in your current working directory, enter the **cat my_next_bestselling_novel** command at the shell prompt and press **Enter** to execute the command.

The above includes a file name, a shell command and a keycap, all presented in mono-spaced bold and all distinguishable thanks to context.

Key combinations can be distinguished from keycaps by the hyphen connecting each part of a key combination. For example:

Press **Enter** to execute the command.

Press **Ctrl+Alt+F1** to switch to the first virtual terminal. Press **Ctrl+Alt+F7** to return to your X-Windows session.

¹ <https://fedorahosted.org/liberation-fonts/>

The first paragraph highlights the particular keycap to press. The second highlights two key combinations (each a set of three keycaps with each set pressed simultaneously).

If source code is discussed, class names, methods, functions, variable names and returned values mentioned within a paragraph will be presented as above, in **mono-spaced bold**. For example:

File-related classes include **filesystem** for file systems, **file** for files, and **dir** for directories. Each class has its own associated set of permissions.

Proportional Bold

This denotes words or phrases encountered on a system, including application names; dialog box text; labeled buttons; check-box and radio button labels; menu titles and sub-menu titles. For example:

Choose **System** → **Preferences** → **Mouse** from the main menu bar to launch **Mouse Preferences**. In the **Buttons** tab, click the **Left-handed mouse** check box and click **Close** to switch the primary mouse button from the left to the right (making the mouse suitable for use in the left hand).

To insert a special character into a **gedit** file, choose **Applications** → **Accessories** → **Character Map** from the main menu bar. Next, choose **Search** → **Find...** from the **Character Map** menu bar, type the name of the character in the **Search** field and click **Next**. The character you sought will be highlighted in the **Character Table**. Double-click this highlighted character to place it in the **Text to copy** field and then click the **Copy** button. Now switch back to your document and choose **Edit** → **Paste** from the **gedit** menu bar.

The above text includes application names; system-wide menu names and items; application-specific menu names; and buttons and text found within a GUI interface, all presented in proportional bold and all distinguishable by context.

Mono-spaced Bold Italic or *Proportional Bold Italic*

Whether mono-spaced bold or proportional bold, the addition of italics indicates replaceable or variable text. Italics denotes text you do not input literally or displayed text that changes depending on circumstance. For example:

To connect to a remote machine using ssh, type **ssh *username@domain.name*** at a shell prompt. If the remote machine is **example.com** and your username on that machine is john, type **ssh john@example.com**.

The **mount -o remount *file-system*** command remounts the named file system. For example, to remount the **/home** file system, the command is **mount -o remount /home**.

To see the version of a currently installed package, use the **rpm -q *package*** command. It will return a result as follows: ***package-version-release***.

Note the words in bold italics above — *username*, *domain.name*, *file-system*, *package*, *version* and *release*. Each word is a placeholder, either for text you enter when issuing a command or for text displayed by the system.

Aside from standard usage for presenting the title of a work, italics denotes the first use of a new and important term. For example:

Publican is a *DocBook* publishing system.

3.2. Pull-quote Conventions

Terminal output and source code listings are set off visually from the surrounding text.

Output sent to a terminal is set in **mono-spaced roman** and presented thus:

```
books      Desktop  documentation  drafts  mss    photos  stuff  svn
books_tests Desktop1  downloads      images  notes  scripts svgs
```

Source-code listings are also set in **mono-spaced roman** but add syntax highlighting as follows:

```
package org.jboss.book.jca.ex1;

import javax.naming.InitialContext;

public class ExClient
{
    public static void main(String args[])
        throws Exception
    {
        InitialContext iniCtx = new InitialContext();
        Object          ref    = iniCtx.lookup("EchoBean");
        EchoHome        home   = (EchoHome) ref;
        Echo             echo   = home.create();

        System.out.println("Created Echo");

        System.out.println("Echo.echo('Hello') = " + echo.echo("Hello"));
    }
}
```

3.3. Notes and Warnings

Finally, we use three visual styles to draw attention to information that might otherwise be overlooked.



Note

Notes are tips, shortcuts or alternative approaches to the task at hand. Ignoring a note should have no negative consequences, but you might miss out on a trick that makes your life easier.



Important

Important boxes detail things that are easily missed: configuration changes that only apply to the current session, or services that need restarting before an update will apply. Ignoring a box labeled 'Important' won't cause data loss but may cause irritation and frustration.



Warning

Warnings should not be ignored. Ignoring warnings will most likely cause data loss.

4. Feedback

Red Hat welcomes your comments and suggestions on the quality and usefulness of its documentation. If you find any errors or have any other suggestions, write to us at docfeedback@gluster.com for clarification and provide the chapter, section, and page number, if available.

Red Hat offers a range of resources related to GlusterFS software:

- Discuss technical problems and solutions on the Discussion Forum (<http://community.gluster.org>)
- Get hands-on step-by-step tutorials (http://www.gluster.com/community/documentation/index.php/Main_Page)
- Reach Support (<http://www.gluster.com/services/>)

Introducing Gluster File System

GlusterFS is an open source, clustered file system capable of scaling to several petabytes and handling thousands of clients. GlusterFS can be flexibly combined with commodity physical, virtual, and cloud resources to deliver highly available and performant enterprise storage at a fraction of the cost of traditional solutions.

GlusterFS clusters together storage building blocks over Infiniband RDMA and/or TCP/IP interconnect, aggregating disk and memory resources and managing data in a single global namespace. GlusterFS is based on a stackable user space design, delivering exceptional performance for diverse workloads.




Figure 1.1. Virtualized Cloud Environments

GlusterFS is designed for today's high-performance, virtualized cloud environments. Unlike traditional data centers, cloud environments require multi-tenancy along with the ability to grow or shrink resources on demand. Enterprises can scale capacity, performance, and availability on demand, with no vendor lock-in, across on-premise, public cloud, and hybrid environments.

GlusterFS is in production at thousands of enterprises spanning media, healthcare, government, education, web 2.0, and financial services. The following table lists the commercial offerings and its documentation location:

Product	Documentation Location
Red Hat Storage Software Appliance	http://download.gluster.com/pub/gluster/RHSSA/3.2/Documentation/UG/html/index.html

Chapter 1. Introducing Gluster File System

Product	Documentation Location
Gluster Storage Software Appliance	http://www.gluster.com/community/documentation/index.php/Gluster_3.2_Gluster_Storage_Software_Appliance_User_Guide
Gluster Virtual Storage Appliance	http://www.gluster.com/community/documentation/index.php/Gluster_3.2_Gluster_Virtual_Storage_Appliance_User_Guide
Gluster Virtual Storage Appliance for Amazon Web Services	http://www.gluster.com/community/documentation/index.php/Gluster_3.2_Gluster_Virtual_Storage_Appliance_for_Amazon_Web_Services_Guide

Preparing to Install GlusterFS Server

This section provides an overview of the preparation needed before installing GlusterFS, prerequisites, and the minimum system requirements.

2.1. Preparation

It is recommended that GlusterFS be installed on 64 bit X-86 hardware, preferably in the Red Hat Hardware Certified List. Before installing GlusterFS, you will need to install an Operating System on the hardware in line with the minimum system requirements below.

2.2. Checking Minimum Requirements

Before installing the GlusterFS, verify that your system matches the following minimum requirements:

2.2.1. Hardware Requirements

Storage Servers

The following storage servers are supported:

- Processor: Intel/AMD x86 64-bit
- Disk: 8GB minimum using direct-attached-storage, RAID, Amazon EBS, and FC/Infiniband/iSCSI SAN disk backends using SATA/SAS/FC disks
- Memory: 1GB minimum

Networking Requirements

The following are the supported networks:

- Gigabit Ethernet
- 10 Gigabit Ethernet
- InfiniBand

2.2.2. Operating System Requirements

Recommended Operating System	Architecture	Modes
RHEL 6.1	x86_64	Client & Server

GlusterFS works with other common Linux distribution like CentOS 5.1 or higher, Ubuntu 8.04 or higher, and Fedora 11 or higher, but has not been tested extensively.

Packages

Ensure that the following packages are installed:

- Bison

- Automake/ Autoconf
- Flex
- libtool
- gcc
- Portmapper (for NFS)
- Fuse

2.2.3. File System Requirements

Red Hat recommends XFS when formatting the disk sub-system. XFS supports metadata journaling, which facilitates quicker crash recovery. The XFS file system can also be de-fragmented and enlarged while mounted and active.

Any other POSIX compliant disk file system, such as Ext3, Ext4, ReiserFS may also work, but has not been tested widely.

Installing GlusterFS Server

This section describes how to install the GlusterFS server software in the following environments:

- Installing GlusterFS on Red Hat Package Manager (RPM) Distributions
- Installing GlusterFS on Debian-based Distributions
- Installing GlusterFS from Source



Note

You must install GlusterFS on all servers.

3.1. Installing GlusterFS on Red Hat Package Manager (RPM) Distributions

To install GlusterFS on Red Hat Package Manager (RPM) distributions, such as RHEL and CentOS:

1. Install required prerequisites on the server using the following command:

```
$ sudo yum -y install wget fuse fuse-libs
```

2. To enable the optional Infiniband support ensure that infiniband is setup by using the following command between nodes:

```
$ ibv_srq_pingpong
```

If infiniband is not installed, use the following steps to setup infiniband:

- a. Download OFED-1.5.2 from <http://www.openfabrics.org/downloads/OFED/ofed-1.5.2/>.
 - b. Install OFED-1.5.2 and start IB fabric. For more details, see README file in tarball folder. For information on how to setup infiniband, see <http://pkg-ofed.aliath.debian.org/howto/infiniband-howto-4.html>.
3. Ensure that TCP ports 111, 24007, 24008, 24009 (24009 + number of bricks across all volumes) are open on all Gluster servers. If you will be using NFS, open additional ports 38465 to 38467.

You can use the following chains with iptables:

```
$ iptables -A RH-Firewall-1-INPUT -m state --state NEW -m tcp -p tcp --dport 24007:24047 -j ACCEPT
$ iptables -A RH-Firewall-1-INPUT -m state --state NEW -m tcp -p tcp --dport 111 -j ACCEPT
$ iptables -A RH-Firewall-1-INPUT -m state --state NEW -m udp -p udp --dport 111 -j ACCEPT
$ iptables -A RH-Firewall-1-INPUT -m state --state NEW -m tcp -p tcp --dport 38465:38467 -j ACCEPT
$ service iptables save
$ service iptables restart
```



Note

You need one open port, starting at 24009 for each brick. This example opens enough ports for 20 storage servers and three bricks.

4. Download the latest GlusterFS core and FUSE RPM files to each server in your cluster.

The core package contains GlusterFS, the glusterd daemon, and the Gluster Native Client. The optional FUSE package contains the FUSE module for Gluster native mounting on client systems.

You can download the software at <http://www.gluster.org/download/>.

5. For each RPM file, get the checksum (using the following command) and compare it against the checksum for that file in the md5sum file.

```
$ md5sum RPM_file.rpm
```



Note

On CentOS/RHEL 5.5 & 5.6, install rsync and python-ctypes packages before installing GlusterFS Geo-replication using `$ sudo rpm -ivh` command. You can download these packages from the following locations:

- rsync: http://packages.sw.be/rsync/rsync-3.0.7-1.el5.rfx.x86_64.rpm .
- python-ctypes: http://download.fedora.redhat.com/pub/epel/5/x86_64/python-ctypes-1.0.2-2.el5.x86_64.rpm.

6. Install GlusterFS on all servers using the following commands:

```
$ sudo rpm -Uvh core_RPM_file
$ sudo rpm -Uvh fuse_RPM_file
$ sudo rpm -ivh geo-replication_RPM_file
```

For example:

```
$ sudo rpm -Uvh glusterfs-core--3.2.0-1.fc11.x86_64.rpm
$ sudo rpm -Uvh glusterfs-fuse--3.2.0-1.fc11.x86_64.rpm
$ sudo rpm -ivh glusterfs-geo-replication-3.2.0-1.fc11.x86_64.rpm
```

7. To double check your installation in RHEL, run the following command:

```
$ /usr/sbin/glusterfs -V
```

3.2. Installing GlusterFS on Debian-based Distributions

To install GlusterFS on Debian-based distributions, such as Ubuntu and Debian:

1. Install OpenSSH Server on each node (server) using the following command:

```
$ sudo apt-get install openssh-server wget nfs-common
```

2. Download the latest GlusterFS .deb file and checksum to each server in your cluster.

You can download the software at <http://www.gluster.org/download/>.

3. For each .deb file, get the checksum (using the following command) and compare it against the checksum for that file in the md5sum file.

```
$ md5sum GlusterFS_DEB_file.deb
```

4. Uninstall GlusterFS v3.0 (or an earlier version) from the server using the following command:

```
$ sudo dpkg r glusterfs
```

(Optional) Run `$ sudo dpkg -purge glusterfs` to purge the configuration files.

5. Install GlusterFS on all servers using the following command:

```
$ sudo dpkg -i GlusterFS_DEB_file
```

For example:

```
$ sudo dpkg -i glusterfs-3.2.0.deb
```

6. Ensure that TCP ports 111, 24007, 24008, 24009-(24009 + number of bricks across all volumes) are open on all Gluster servers. If you will be using NFS, open additional ports 38465 to 38467. You can use the following chains with iptables:

```
$ iptables -A INPUT -m state --state NEW -m tcp -p tcp --dport 24007:24047 -j ACCEPT
$ iptables -A INPUT -m state --state NEW -m tcp -p tcp --dport 111 -j ACCEPT
$ iptables -A INPUT -m state --state NEW -m udp -p udp --dport 111 -j ACCEPT
$ iptables -A INPUT -m state --state NEW -m tcp -p tcp --dport 38465:38467 -j ACCEPT
$ service iptables save
$ service iptables restart
```



Note

You need one open port, starting at 24009 for each brick. This example opens enough ports for 20 storage servers and three bricks.

3.3. Installing GlusterFS from Source

To build and install GlusterFS from the source code

1. Create a new directory using the following commands:

```
# mkdir glusterfs
```

```
# cd glusterfs
```

2. Download the source code. You can download the source at <http://www.gluster.org/download/>.

3. Extract the source code using the following command:

```
# tar -xvzf <source file>
```

4. Run the configuration utility using the following command:

```
# ./configure
GlusterFS configure summary
=====
FUSE client: yes
Infiniband verbs: yes
epoll IO multiplex: yes
argp-standalone: no
fusermount: no
readline: yes
georeplication: yes
```

The configuration summary shows the components that will be built with GlusterFS.

5. Build the GlusterFS software using the following commands:

```
# make
```

```
# make install
```

6. Verify that the correct version of GlusterFS is installed, using the following command:

```
# glusterfs --version
```

7. Ensure that TCP ports 111, 24007,24008, 24009-(24009 + number of bricks across all volumes) are open on all Gluster servers. If you will be using NFS, open additional ports 38465-(38465 + number of Gluster servers).

You can use the following chains with iptables:

```
$ iptables -A RH-Firewall-1-INPUT -m state --state NEW -m tcp -p tcp --dport 24007:24011 -
j ACCEPT
$ iptables -A RH-Firewall-1-INPUT -m state --state NEW -m tcp -p tcp --dport 111 -j
ACCEPT
$ iptables -A RH-Firewall-1-INPUT -m state --state NEW -m udp -p udp --dport 111 -j
ACCEPT
$ iptables -A RH-Firewall-1-INPUT -m state --state NEW -m tcp -p tcp --dport 38465:38485 -
j ACCEPT
$ service iptables save
$ service iptables restart
```



Note

You need one open port, starting at 38465 and incrementing sequentially for each Gluster storage server, and one port, starting at 24009 for each bricks. This example opens enough ports for 20 storage servers and three bricks.

3.4. Configuring GlusterFS to work over InfiniBand

You can optionally configure GlusterFS for OpenFabrics verbs RDMA-conformant transports, such as InfiniBand.



Note

Commercial Support for InfiniBand is currently not available, write to info@gluster.com for more information on InfiniBand support.

3.4.1. Configuring GlusterFS for InfiniBand on RPM distribution-based systems

1. Install OFED version 1.5.2 from source using the directions in the download

```
$ wget http://www.openfabrics.org/downloads/OFED/ofed-1.5.2/
```

2. Install GlusterFS. For more information, see [Section 3.1, “Installing GlusterFS on Red Hat Package Manager \(RPM\) Distributions”](#).

3. Using SSH, download the RDMA RPM file to each server in your cluster.

The RDMA packages contain the OpenFabrics verbs RDMA module for Infiniband and 10GigE networks.

Download the RPM files from the following location:

```
http://download.gluster.com/pub/gluster/glusterfs/3.2
```

4. Install the RDMA module on the server using the following command:

```
$ sudo rpm -Uvh <rdma_RPM_file>
```

For example:

```
$ sudo rpm -Uvh glusterfs-rdma-3.2.0alpha-1.x86_64.rpm
```

3.4.2. Configuring GlusterFS for InfiniBand on Debian distribution-based systems

1. Install OFED version 1.5.2 from source using the directions in the download

```
$ wget http://www.openfabrics.org/downloads/OFED/ofed-1.5.2/  
OFED-1.5.2.tgz
```

2. Download the latest GlusterFS rdma.deb file and checksum to each server in your cluster. You can download the software at <http://www.gluster.org/download/>.

3. For each .deb file, get the checksum (using the following command) and compare it against the checksum for that file in the md5sum file.

```
$ md5sum GlusterFS_RDMA_DEB_file.deb
```

4. Uninstall GlusterFS v3.0 (or an earlier version) from the server using the following command:

```
$ sudo dpkg -r glusterfs
```

5. Install GlusterFS on all servers using the following command:

```
$ sudo dpkg -i GlusterFS_RDMA_DEB_file
```

For example:

```
$ sudo dpkg -i glusterfs_3.2.0-1_amd64_with_rdma.deb
```

Upgrading GlusterFS

This document describes how to upgrade your environment from Gluster v3.0.x or lower to Gluster v3.2.x and from Gluster v3.1.x to Gluster v3.2.x. Red Hat recommends that you back up your data before upgrading to Gluster v3.2.x.

4.1. Upgrading GlusterFS from v3.0.x or older to v3.2.x

Upgrading all glusterfs and glusterd instances to v3.2.x enables you to take advantage of new capabilities like geo-replication, directory quota, and other new features available in GlusterFS v3.2.x.

Existing GlusterFS 2.0.x and 3.0.x users must install GlusterFS v3.2.x on all servers and clients when upgrading.

1. Uninstall the currently installed version of GlusterFS using the following command:

```
# rpm -e glusterfs-common glusterfs-server glusterfs-client
```

Use the equivalent command for Debian-based distributions.

2. Install GlusterFS v3.2.x.

For detailed instructions on installing GlusterFS, see [Chapter 3, Installing GlusterFS Server](#).

3. Examine the volume file to determine the current volume configuration.

You can do this by displaying the first few lines of the volume file and searching for the glusterfs-volgen entry.

For example, enter the following line to display the current volume file:

```
# head VOLNAME.vol
```

The command displays output similar to the following

```
# ....
# ....
# 'cmd line: glusterfs-volgen -n VOLNAME -r 1 node1:/data/export1 node2:/data/export1
node3:/data/export1 node4:/data/export1'
# ....
....
```

In this example, the volume is configured as a distributed replicated with export paths defined as /data/export1 on servers node1, node2, node3, and node4.

4. Create GlusterFS v3.2.x volumes corresponding to volumes in your current GlusterFS v3.0.x environment. See http://download.gluster.com/pub/gluster/glusterfs/3.2/Documentation/AG/html/chap-Administration_Guide-Setting_Volumes.html.
5. Make sure to mount the volume from only one client (not nfs mount) and also all the bricks are up and running. Then traverse the whole volume using the following command:

```
# find /mount/glusterfs >/dev/null
```



Note

This step ensures that the data from 3.0.x versions will become compatible with 3.1.x or 3.2.x versions.

6. Mount the other clients using either NFS or the Gluster Native Client.

For more information, see http://download.gluster.com/pub/gluster/glusterfs/3.2/Documentation/AG/html/chap-Administration_Guide-GlusterFS_Client.html.

Congratulations! You have successfully upgraded from GlusterFS v3.0.x to GlusterFS v3.2.x.

4.2. Upgrading GlusterFS from v3.1.x to v3.2.x

Red Hat recommends that you back up your existing configuration and volume files prior to installing GlusterFS v3.2.x.

Upgrading all glusterfs and glusterd instances to v3.2.x enables you to take advantage of new capabilities like geo-replication, directory quota, and other new features available in GlusterFS v3.2.x.

To upgrade from GlusterFS v3.1.x to GlusterFS v3.2.x

In an environment with replicated bricks it is recommended that you upgrade a single storage server, confirm stability, and then upgrade the replica.

Use the same installation method for the upgrade as the original glusterfs installation.

4.2.1. Using RPMs on RHEL, CentOS, Fedora

1. Download the 3.2.5 RPMs from <http://download.gluster.com/pub/gluster/glusterfs/3.2/>.
2. Run rpm using the following command:

```
# rpm -U glusterfs*
```

4.2.2. Using dpkg on Debian, Ubuntu

1. Download the 3.2.5 packages from <http://download.gluster.com/pub/gluster/glusterfs/3.2/>.
2. Run dpkg using the following command:

```
# dpkg -i glusterfs*
```

4.2.3. Building from Source

1. Download the 3.2.5 source code from <http://download.gluster.com/pub/gluster/glusterfs/3.2/3.2.5/glusterfs-3.2.5.tar.gz>
2. Unpack and install GlusterFS using the following commands:

```
# gunzip glusterfs-3.2.5.tar.gz
```

```
# tar xvf glusterfs-3.2.5.tar
```

```
# cd glusterfs-3.2.5
```

```
# ./configure
```

```
# make
```

```
# make install
```

3. Stop GlusterFS using the following commands, this step will disconnect Gluster Native clients.

```
# killall glusterfsd
```

```
# killall glusterfs
```

```
# killall glusterd
```

4. Start GlusterFS using the following command:

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


Uninstalling GlusterFS

This section describes how to uninstall the GlusterFS software in the following environments:

- Uninstalling GlusterFS on Red Hat Package Manager (RPM) Distributions
- Uninstalling GlusterFS on Debian-based Distributions
- Uninstalling GlusterFS when Compiled from Source



Note

You must uninstall GlusterFS on all servers.

5.1. Uninstalling GlusterFS on Red Hat Package Manager (RPM) Distributions

To uninstall GlusterFS on Red Hat Package Manager (RPM) distributions, such as RHEL and CentOS:

1. List the installed glusterfs packages using the command:

```
# rpm -qa | grep glusterfs
```

The following files will be listed:

- glusterfs-fuse-3.2.x-1.fc11.x86_64 .
- glusterfs-core-3.2.x-1.fc11.x86_64 .
- glusterfs-geo-replication-3.2.x-1.fc11.x86_64.rpm

2. Uninstall these rpms by running the following commands:

```
# rpm -e glusterfs-fuse
```

```
# rpm -e glusterfs-core
```

```
# rpm -e glusterfs-geo-replication
```

5.2. Uninstalling GlusterFS on Debian-based Distributions

To uninstall GlusterFS on Debian-based distributions, such as Ubuntu and Debian:

1. Remove the glusterfs packages using the following command:

```
# dpkg -r glusterfs
```

2. Remove the glusterfs configuration files using the following command:

```
# dpkg --purge glusterfs
```

5.3. Uninstalling GlusterFS when Compiled from Source

To uninstall GlusterFS when compiled from source:

1. Navigate to the directory where GlusterFS source was compiled.
2. Uninstall GlusterFS using the following command:

```
# make uninstall
```

3. (Optional) Cleanup object files using the following command:

```
# make clean
```