

# Red Hat Enterprise Linux 5

## Global File System 2

Red Hat Global File System 2



# Red Hat Enterprise Linux 5 Global File System 2

## Red Hat Global File System 2

### Edition 7

Copyright © 2012 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, NC 27606-2072 USA  
Phone: +1 919 754 3700  
Phone: 888 733 4281  
Fax: +1 919 754 3701

This book provides information about configuring and maintaining Red Hat GFS2 (Red Hat Global File System 2) for Red Hat Enterprise Linux 5.

---

---

<b>Introduction</b>	<b>v</b>
1. Audience .....	v
2. Related Documentation .....	v
3. Feedback .....	vi
4. Document Conventions .....	vi
4.1. Typographic Conventions .....	vi
4.2. Pull-quote Conventions .....	viii
4.3. Notes and Warnings .....	viii
<b>1. GFS2 Overview</b>	<b>1</b>
1.1. Before Setting Up GFS2 .....	2
1.2. Differences between GFS and GFS2 .....	3
1.2.1. GFS2 Command Names .....	3
1.2.2. Additional Differences Between GFS and GFS2 .....	4
1.2.3. GFS2 Performance Improvements .....	5
<b>2. Getting Started</b>	<b>7</b>
2.1. Prerequisite Tasks .....	7
2.2. Initial Setup Tasks .....	7
<b>3. Managing GFS2</b>	<b>11</b>
3.1. Making a File System .....	11
3.2. Mounting a File System .....	14
3.3. Unmounting a File System .....	17
3.4. Special Considerations when Mounting GFS2 File Systems .....	17
3.5. GFS2 Quota Management .....	18
3.5.1. Setting Quotas .....	18
3.5.2. Displaying Quota Limits and Usage .....	19
3.5.3. Synchronizing Quotas .....	21
3.5.4. Enabling/Disabling Quota Enforcement .....	22
3.5.5. Enabling Quota Accounting .....	22
3.6. Growing a File System .....	23
3.7. Adding Journals to a File System .....	24
3.8. Data Journaling .....	26
3.9. Configuring <b>atime</b> Updates .....	27
3.10. Suspending Activity on a File System .....	28
3.11. Repairing a File System .....	28
3.12. Bind Mounts and Context-Dependent Path Names .....	31
3.13. Bind Mounts and File System Mount Order .....	32
3.14. The GFS2 Withdraw Function .....	34
<b>A. Converting a File System from GFS to GFS2</b>	<b>37</b>
<b>B. Revision History</b>	<b>39</b>
<b>Index</b>	<b>41</b>



---

# Introduction

This book provides information about configuring and maintaining Red Hat GFS2 (Red Hat Global File System 2). Red Hat GFS2 can be run in Red Hat Enterprise Linux 5.3 and later. For the Red Hat Enterprise Linux 5.5 release and later, Red Hat does not support the use of GFS2 as a single-node file system. For information about Red Hat Cluster Suite see *Red Hat Cluster Suite Overview* and *Configuring and Managing a Red Hat Cluster*.

HTML and PDF versions of all the official Red Hat Enterprise Linux manuals and release notes are available online at <http://docs.redhat.com/docs/en-US/index.html>.

## 1. Audience

This book is intended primarily for Linux system administrators who are familiar with the following activities:

- Linux system administration procedures, including kernel configuration
- Installation and configuration of shared storage networks, such as Fibre Channel SANs

## 2. Related Documentation

For more information about using Red Hat Enterprise Linux, refer to the following resources:

- *Red Hat Enterprise Linux Installation Guide* — Provides information regarding installation of Red Hat Enterprise Linux.
- *Red Hat Enterprise Linux Deployment Guide* — Provides information regarding the deployment, configuration and administration of Red Hat Enterprise Linux 5.

For more information about Red Hat Cluster Suite, refer to the following resources:

- *Red Hat Cluster Suite Overview* — Provides a high level overview of the Red Hat Cluster Suite.
- *Configuring and Managing a Red Hat Cluster* — Provides information about installing, configuring and managing Red Hat Cluster components.
- *Logical Volume Manager Administration* — Provides a description of the Logical Volume Manager (LVM), including information on running LVM in a clustered environment.
- *Global File System: Configuration and Administration* — Provides information about installing, configuring, and maintaining Red Hat GFS (Red Hat Global File System).
- *Using Device-Mapper Multipath* — Provides information about using the Device-Mapper Multipath feature of Red Hat Enterprise Linux.
- *Using GNBD with Global File System* — Provides an overview on using Global Network Block Device (GNBD) with Red Hat GFS.
- *Linux Virtual Server Administration* — Provides information on configuring high-performance systems and services with the Linux Virtual Server (LVS).
- *Red Hat Cluster Suite Release Notes* — Provides information about the current release of Red Hat Cluster Suite.

Red Hat Cluster Suite documentation and other Red Hat documents are available in HTML, PDF, and RPM versions on the Red Hat Enterprise Linux Documentation CD and online at <http://www.redhat.com/docs/>.

### 3. Feedback

If you spot a typo, or if you have thought of a way to make this manual better, we would love to hear from you. Please submit a report in Bugzilla (<http://bugzilla.redhat.com/bugzilla/>) against the component **Documentation-cluster**.

Be sure to mention the manual's identifier:

```
rh-gfs2(EN)-5 (2012-2-16T15:15)
```

By mentioning this manual's identifier, we know exactly which version of the guide you have.

If you have a suggestion for improving the documentation, try to be as specific as possible. If you have found an error, please include the section number and some of the surrounding text so we can find it easily.

### 4. 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*<sup>1</sup> 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.

#### 4.1. 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:

---

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

Press **Enter** to execute the command.

Press **Ctrl+Alt+F2** to switch to the first virtual terminal. Press **Ctrl+Alt+F1** to return to your X-Windows session.

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.

## Introduction

---

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.

## 4.2. 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"));
    }
}
```

## 4.3. 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' will not cause data loss but may cause irritation and frustration.





## Warning

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



# GFS2 Overview

The Red Hat GFS2 file system is a native file system that interfaces directly with the Linux kernel file system interface (VFS layer). When implemented as a cluster file system, GFS2 employs distributed metadata and multiple journals. Red Hat supports the use of GFS2 file systems only as implemented in Red Hat Cluster Suite.



## Note

Although a GFS2 file system can be implemented in a standalone system or as part of a cluster configuration, for the Red Hat Enterprise Linux 5.5 release and later Red Hat does not support the use of GFS2 as a single-node file system. Red Hat does support a number of high-performance single node file systems which are optimized for single node and thus have generally lower overhead than a cluster file system. Red Hat recommends using these file systems in preference to GFS2 in cases where only a single node needs to mount the file system.

Red Hat will continue to support single-node GFS2 file systems for existing customers.



## Note

Red Hat does not support using GFS2 for cluster file system deployments greater than 16 nodes.

GFS2 is based on a 64-bit architecture, which can theoretically accommodate an 8 EB file system. However, the current supported maximum size of a GFS2 file system for 64-bit hardware is 100 TB. The current supported maximum size of a GFS2 file system for 32-bit hardware for Red Hat Enterprise Linux Release 5.3 and later is 16 TB. If your system requires larger GFS2 file systems, contact your Red Hat service representative.

When determining the size of your file system, you should consider your recovery needs. Running the **fsck.gfs2** command on a very large file system can take a long time and consume a large amount of memory. Additionally, in the event of a disk or disk-subsystem failure, recovery time is limited by the speed of your backup media. For information on the amount of memory the **fsck.gfs2** command requires, see [Section 3.11, “Repairing a File System”](#).

When configured in a Red Hat Cluster Suite, Red Hat GFS2 nodes can be configured and managed with Red Hat Cluster Suite configuration and management tools. Red Hat GFS2 then provides data sharing among GFS2 nodes in a Red Hat cluster, with a single, consistent view of the file system name space across the GFS2 nodes. This allows processes on different nodes to share GFS2 files in the same way that processes on the same node can share files on a local file system, with no discernible difference. For information about Red Hat Cluster Suite refer to *Configuring and Managing a Red Hat Cluster*.

While a GFS2 file system may be used outside of LVM, Red Hat supports only GFS2 file systems that are created on a CLVM logical volume. CLVM is a cluster-wide implementation of LVM, enabled by the CLVM daemon **clvmd**, which manages LVM logical volumes in a Red Hat Cluster Suite cluster. The daemon makes it possible to use LVM2 to manage logical volumes across a cluster, allowing all nodes

in the cluster to share the logical volumes. For information on the LVM volume manager, see *Logical Volume Manager Administration*

The `gfs2.ko` kernel module implements the GFS2 file system and is loaded on GFS2 cluster nodes.



### Note

When you configure a GFS2 file system as a cluster file system, you must ensure that all nodes in the cluster have access to the shared storage. Asymmetric cluster configurations in which some nodes have access to the shared storage and others do not are not supported. This does not require that all nodes actually mount the GFS2 file system itself.

This chapter provides some basic, abbreviated information as background to help you understand GFS2. It contains the following sections:

- [Section 1.1, “Before Setting Up GFS2”](#)
- [Section 1.2, “Differences between GFS and GFS2”](#)

## 1.1. Before Setting Up GFS2

Before you install and set up GFS2, note the following key characteristics of your GFS2 file systems:

### GFS2 nodes

Determine which nodes in the Red Hat Cluster Suite will mount the GFS2 file systems.

### Number of file systems

Determine how many GFS2 file systems to create initially. (More file systems can be added later.)

### File system name

Determine a unique name for each file system. The name must be unique for all **lock\_dlm** file systems over the cluster, and for all file systems (**lock\_dlm** and **lock\_nolock**) on each local node. Each file system name is required in the form of a parameter variable. For example, this book uses file system names **mydata1** and **mydata2** in some example procedures.

### Journals

Determine the number of journals for your GFS2 file systems. One journal is required for each node that mounts a GFS2 file system. GFS2 allows you to add journals dynamically at a later point as additional servers mount a file system. For information on adding journals to a GFS2 file system, see [Section 3.7, “Adding Journals to a File System”](#).

### GNBD server nodes

If you are using GNBD, determine how many GNBD server nodes are needed. Note the hostname and IP address of each GNBD server node for setting up GNBD clients later. For information on using GNBD with GFS2, see the *Using GNBD with Global File System* document.

### Storage devices and partitions

Determine the storage devices and partitions to be used for creating logical volumes (via CLVM) in the file systems.



## Note

You may see performance problems with GFS2 when many create and delete operations are issued from more than one node in the same directory at the same time. If this causes performance problems in your system, you should localize file creation and deletions by a node to directories specific to that node as much as possible.

## 1.2. Differences between GFS and GFS2

This section lists the improvements and changes that GFS2 offers over GFS.

Migrating from GFS to GFS2 requires that you convert your GFS file systems to GFS2 with the **gfs2\_convert** utility. For information on the **gfs2\_convert** utility, see [Appendix A, Converting a File System from GFS to GFS2](#).

### 1.2.1. GFS2 Command Names

In general, the functionality of GFS2 is identical to GFS. The names of the file system commands, however, specify GFS2 instead of GFS. [Table 1.1, “GFS and GFS2 Commands”](#) shows the equivalent GFS and GFS2 commands.

Table 1.1. GFS and GFS2 Commands

GFS Command	GFS2 Command	Description
<b>mount</b>	<b>mount</b>	Mount a file system. The system can determine whether the file system is a GFS or GFS2 file system type. For information on the GFS2 mount options see the <code>gfs2_mount(8)</code> man page.
<b>umount</b>	<b>umount</b>	Unmount a file system.
<b>fsck</b> <b>gfs_fsck</b>	<b>fsck</b> <b>fsck.gfs2</b>	Check and repair an unmounted file system.
<b>gfs_grow</b>	<b>gfs2_grow</b>	Grow a mounted file system.
<b>gfs_jadd</b>	<b>gfs2_jadd</b>	Add a journal to a mounted file system.
<b>gfs_mkfs</b> <b>mkfs -t gfs</b>	<b>mkfs.gfs2</b> <b>mkfs -t gfs2</b>	Create a file system on a storage device.
<b>gfs_quota</b>	<b>gfs2_quota</b>	Manage quotas on a mounted file system.
<b>gfs_tool</b>	<b>gfs2_tool</b>	Configure, tune, or gather information about a file system.
<b>gfs_edit</b>	<b>gfs2_edit</b>	Display, print, or edit file system internal structures. The <b>gfs2_edit</b> command can be used for GFS file systems as well as GFS2 file system.
<b>gfs_tool</b> <b>setflag</b> <b>jdata/</b> <b>inherit_jdata</b>	<b>chattr +j</b> (preferred)	Enable journaling on a file or directory.
<b>setfacl/</b> <b>getfacl</b>	<b>setfacl/</b> <b>getfacl</b>	Set or get file access control list for a file or directory.

GFS Command	GFS2 Command	Description
<b>setfattr/ getfattr</b>	<b>setfattr/ getfattr</b>	set or get the extended attributes of a file.

For a full listing of the supported options for the GFS2 file system commands, see the man pages for those commands.

### 1.2.2. Additional Differences Between GFS and GFS2

This section summarizes the additional differences in GFS and GFS2 administration that are not described in [Section 1.2.1, “GFS2 Command Names”](#).

#### Context-Dependent Path Names

GFS2 file systems do not provide support for context-dependent path names, which allow you to create symbolic links that point to variable destination files or directories. For this functionality in GFS2, you can use the **bind** option of the **mount** command. For information on managing pathnames in GFS2, see [Section 3.12, “Bind Mounts and Context-Dependent Path Names”](#).

#### gfs2.ko Module

The kernel module that implements the GFS file system is **gfs.ko**. The kernel module that implements the GFS2 file system is **gfs2.ko**.

#### Enabling Quota Enforcement in GFS2

In GFS2 file systems, quota enforcement is disabled by default and must be explicitly enabled. To enable and disable quotas for GFS2 file systems, you use the **quota=on|off|account** option for the **mount** command. For information on enabling and disabling quota enforcement, see [Section 3.5.4, “Enabling/Disabling Quota Enforcement”](#).

#### Data Journaling

GFS2 file systems support the use of the **chattr** command to set and clear the **j** flag on a file or directory. Setting the **+j** flag on a file enables data journaling on that file. Setting the **+j** flag on a directory means "inherit jdata", which indicates that all files and directories subsequently created in that directory are journaled. Using the **chattr** command is the preferred way to enable and disable data journaling on a file.

#### Adding Journals Dynamically

In GFS file systems, journals are embedded metadata that exists outside of the file system, making it necessary to extend the size of the logical volume that contains the file system before adding journals. In GFS2 file systems, journals are plain (though hidden) files. This means that for GFS2 file systems, journals can be dynamically added as additional servers mount a file system, as long as space remains on the file system for the additional journals. For information on adding journals to a GFS2 file system, see [Section 3.7, “Adding Journals to a File System”](#).

#### atime\_quantum parameter removed

The GFS2 file system does not support the **atime\_quantum** tunable parameter, which can be used by the GFS file system to specify how often **atime** updates occur. In its place GFS2 supports the **relatime** and **noatime** mount options. The **relatime** mount option is recommended to achieve similar behavior to setting the **atime\_quantum** parameter in GFS.

## The `data=` option of the `mount` command

When mounting GFS2 file systems, you can specify the `data=ordered` or `data=writeback` option of the `mount`. When `data=ordered` is set, the user data modified by a transaction is flushed to the disk before the transaction is committed to disk. This should prevent the user from seeing uninitialized blocks in a file after a crash. When `data=writeback` is set, the user data is written to the disk at any time after it is dirtied. This does not provide the same consistency guarantee as `ordered` mode, but it should be slightly faster for some workloads. The default is `ordered` mode.

## The `gfs2_tool` command

The `gfs2_tool` command supports a different set of options for GFS2 than the `gfs_tool` command supports for GFS:

- The `gfs2_tool` command supports a `journals` parameter that prints out information about the currently configured journals, including how many journals the file system contains.
- The `gfs2_tool` command does not support the `counters` flag, which the `gfs_tool` command uses to display GFS statistics.
- The `gfs2_tool` command does not support the `inherit_jdata` flag. To flag a directory as "inherit jdata", you can set the `jdata` flag on the directory or you can use the `chattr` command to set the `+j` flag on the directory. Using the `chattr` command is the preferred way to enable and disable data journaling on a file.

## The `gfs2_edit` command

The `gfs2_edit` command supports a different set of options for GFS2 than the `gfs_edit` command supports for GFS. For information on the specific options each version of the command supports, see the `gfs2_edit` and `gfs_edit` man pages.

### 1.2.3. GFS2 Performance Improvements

There are many features of GFS2 file systems that do not result in a difference in the user interface from GFS file systems but which improve file system performance.

A GFS2 file system provides improved file system performance in the following ways:

- Better performance for heavy usage in a single directory
- Faster synchronous I/O operations
- Faster cached reads (no locking overhead)
- Faster direct I/O with preallocated files (provided I/O size is reasonably large, such as 4M blocks)
- Faster I/O operations in general
- Faster Execution of the `df` command, because of faster `statfs` calls
- Improved `atime` mode to reduce the number of write I/O operations generated by `atime` when compared with GFS

GFS2 file systems provide broader and more mainstream support in the following ways:

- GFS2 is part of the upstream kernel (integrated into 2.6.19).
- GFS2 supports the following features.

- SELinux extended attributes
- the **lsattr()** and **chattr()** attribute settings via standard **ioctl()** calls
- nanosecond timestamps

A GFS2 file system provides the following improvements to the internal efficiency of the file system:

- GFS2 uses less kernel memory.
- GFS2 requires no metadata generation numbers.

Allocating GFS2 metadata does not require reads. Copies of metadata blocks in multiple journals are managed by revoking blocks from the journal before lock release.

- GFS2 includes a much simpler log manager that knows nothing about unlinked inodes or quota changes.
- The **gfs2\_grow** and **gfs2\_jadd** commands use locking to prevent multiple instances running at the same time.
- The ACL code has been simplified for calls like **creat()** and **mkdir()**.
- Unlinked inodes, quota changes, and **statfs** changes are recovered without remounting the journal.



# Getting Started

This chapter describes procedures for initial setup of GFS2 and contains the following sections:

- [Section 2.1, “Prerequisite Tasks”](#)
- [Section 2.2, “Initial Setup Tasks”](#)

## 2.1. Prerequisite Tasks

You should complete the following tasks before setting up Red Hat GFS2:

- Make sure that you have noted the key characteristics of the GFS2 nodes (refer to [Section 1.1, “Before Setting Up GFS2”](#)).
- Make sure that the clocks on the GFS2 nodes are synchronized. It is recommended that you use the Network Time Protocol (NTP) software provided with your Red Hat Enterprise Linux distribution.



### Note

The system clocks in GFS2 nodes must be within a few minutes of each other to prevent unnecessary inode time-stamp updating. Unnecessary inode time-stamp updating severely impacts cluster performance.

- In order to use GFS2 in a clustered environment, you must configure your system to use the Clustered Logical Volume Manager (CLVM), a set of clustering extensions to the LVM Logical Volume Manager. In order to use CLVM, the Red Hat Cluster Suite software, including the `clvmd` daemon, must be running. For information on using CLVM, see *Logical Volume Manager Administration*. For information on installing and administering Red Hat Cluster Suite, see *Cluster Administration*.

## 2.2. Initial Setup Tasks

Initial GFS2 setup consists of the following tasks:

1. Setting up logical volumes.
2. Making a GFS2 files system.
3. Mounting file systems.

Follow these steps to set up GFS2 initially.

1. Using LVM, create a logical volume for each Red Hat GFS2 file system.



### Note

You can use `init.d` scripts included with Red Hat Cluster Suite to automate activating and deactivating logical volumes. For more information about `init.d` scripts, refer to *Configuring and Managing a Red Hat Cluster*.

2. Create GFS2 file systems on logical volumes created in Step 1. Choose a unique name for each file system. For more information about creating a GFS2 file system, refer to [Section 3.1, “Making a File System”](#).

You can use either of the following formats to create a clustered GFS2 file system:

```
mkfs.gfs2 -p lock_dlm -t ClusterName:FSName -j NumberJournals BlockDevice
```

```
mkfs -t gfs2 -p lock_dlm -t LockTableName -j NumberJournals BlockDevice
```

You can use either of the following formats to create a local GFS2 file system:

```
mkfs.gfs2 -p lock_nolock -j NumberJournals BlockDevice
```

```
mkfs -t gfs2 -p lock_nolock -j NumberJournals BlockDevice
```

For more information on creating a GFS2 file system, see [Section 3.1, “Making a File System”](#).

3. At each node, mount the GFS2 file systems. For more information about mounting a GFS2 file system, see [Section 3.2, “Mounting a File System”](#).

Command usage:

```
mount BlockDevice MountPoint
```

```
mount -o acl BlockDevice MountPoint
```

The `-o acl` mount option allows manipulating file ACLs. If a file system is mounted without the `-o acl` mount option, users are allowed to view ACLs (with `getfacl`), but are not allowed to set them (with `setfacl`).

**Note**

You can use **init.d** scripts included with Red Hat Cluster Suite to automate mounting and unmounting GFS2 file systems. For more information about **init.d** scripts, refer to *Configuring and Managing a Red Hat Cluster*.



# Managing GFS2

This chapter describes the tasks and commands for managing GFS2 and consists of the following sections:

- [Section 3.1, “Making a File System”](#)
- [Section 3.2, “Mounting a File System”](#)
- [Section 3.3, “Unmounting a File System”](#)
- [Section 3.5, “GFS2 Quota Management”](#)
- [Section 3.6, “Growing a File System”](#)
- [Section 3.7, “Adding Journals to a File System”](#)
- [Section 3.8, “Data Journaling”](#)
- [Section 3.9, “Configuring \*atime\* Updates”](#)
- [Section 3.10, “Suspending Activity on a File System”](#)
- [Section 3.11, “Repairing a File System”](#)
- [Section 3.12, “Bind Mounts and Context-Dependent Path Names”](#)
- [Section 3.13, “Bind Mounts and File System Mount Order”](#)
- [Section 3.14, “The GFS2 Withdraw Function”](#)

## 3.1. Making a File System

You create a GFS2 file system with the `mkfs.gfs2` command. You can also use the `mkfs` command with the `-t gfs2` option specified. A file system is created on an activated LVM volume. The following information is required to run the `mkfs.gfs2` command:

- Lock protocol/module name (the lock protocol for a cluster is `lock_dlm`)
- Cluster name (when running as part of a cluster configuration)
- Number of journals (one journal required for each node that may be mounting the file system)

When creating a GFS2 file system, you can use the `mkfs.gfs2` command directly, or you can use the `mkfs` command with the `-t` parameter specifying a file system of type `gfs2`, followed by the `gfs2` file system options.



### Note

Once you have created a GFS2 file system with the `mkfs.gfs2` command, you cannot decrease the size of the file system. You can, however, increase the size of an existing file system with the `gfs2_grow` command, as described in [Section 3.6, “Growing a File System”](#).

### Usage

When creating a clustered GFS2 file system, you can use either of the following formats:

```
mkfs.gfs2 -p LockProtoName -t LockTableName -j NumberJournals BlockDevice
```

```
mkfs -t gfs2 -p LockProtoName -t LockTableName -j NumberJournals BlockDevice
```

When creating a local GFS2 file system, you can use either of the following formats:



#### Note

For the Red Hat Enterprise Linux 5.5 release and later Red Hat does not support the use of GFS2 as a single-node file system. Red Hat will continue to support single-node GFS2 file systems for existing customers.

```
mkfs.gfs2 -p LockProtoName -j NumberJournals BlockDevice
```

```
mkfs -t gfs2 -p LockProtoName -j NumberJournals BlockDevice
```



#### Warning

Make sure that you are very familiar with using the *LockProtoName* and *LockTableName* parameters. Improper use of the *LockProtoName* and *LockTableName* parameters may cause file system or lock space corruption.

#### *LockProtoName*

Specifies the name of the locking protocol to use. The lock protocol for a cluster is **lock\_dlm**.

#### *LockTableName*

This parameter is specified for GFS2 file system in a cluster configuration. It has two parts separated by a colon (no spaces) as follows: *ClusterName:FSName*

- *ClusterName*, the name of the Red Hat cluster for which the GFS2 file system is being created.
- *FSName*, the file system name, can be 1 to 16 characters long. The name must be unique for all **lock\_dlm** file systems over the cluster, and for all file systems (**lock\_dlm** and **lock\_noLOCK**) on each local node.

#### *Number*

Specifies the number of journals to be created by the **mkfs.gfs2** command. One journal is required for each node that mounts the file system. For GFS2 file systems, more journals can be

added later without growing the file system, as described in [Section 3.7, “Adding Journals to a File System”](#).

### *BlockDevice*

Specifies a logical or physical volume.

## Examples

In these examples, **lock\_dlm** is the locking protocol that the file system uses, since this is a clustered file system. The cluster name is **alpha**, and the file system name is **mydata1**. The file system contains eight journals and is created on **/dev/vg01/lvol0**.

```
mkfs.gfs2 -p lock_dlm -t alpha:mydata1 -j 8 /dev/vg01/lvol0
```

```
mkfs -t gfs2 -p lock_dlm -t alpha:mydata1 -j 8 /dev/vg01/lvol0
```

In these examples, a second **lock\_dlm** file system is made, which can be used in cluster **alpha**. The file system name is **mydata2**. The file system contains eight journals and is created on **/dev/vg01/lvol1**.

```
mkfs.gfs2 -p lock_dlm -t alpha:mydata2 -j 8 /dev/vg01/lvol1
```

```
mkfs -t gfs2 -p lock_dlm -t alpha:mydata2 -j 8 /dev/vg01/lvol1
```

## Complete Options

[Table 3.1, “Command Options: mkfs.gfs2”](#) describes the **mkfs.gfs2** command options (flags and parameters).

Table 3.1. Command Options: **mkfs.gfs2**

Flag	Parameter	Description
<b>-c</b>	<i>Megabytes</i>	Sets the initial size of each journal's quota change file to <i>Megabytes</i> .
<b>-D</b>		Enables debugging output.
<b>-h</b>		Help. Displays available options.
<b>-J</b>	<i>MegaBytes</i>	Specifies the size of the journal in megabytes. Default journal size is 128 megabytes. The minimum size is 8 megabytes. Larger journals improve performance, although they use more memory than smaller journals.
<b>-j</b>	<i>Number</i>	Specifies the number of journals to be created by the <b>mkfs.gfs2</b> command. One journal is required for each node that mounts the file system. If this option is not specified, one journal will be created. For GFS2 file systems, you can add additional journals at a later time without growing the file system.

Flag	Parameter	Description
<b>-o</b>		Prevents the <code>mkfs.gfs2</code> command from asking for confirmation before writing the file system.
<b>-p</b>	<i>LockProtoName</i>	Specifies the name of the locking protocol to use. Recognized locking protocols include: <b>lock_dlm</b> — The standard locking module, required for a clustered file system. <b>lock_nolock</b> — Used when GFS2 is acting as a local file system (one node only).
<b>-q</b>		Quiet. Do not display anything.
<b>-r</b>	<i>MegaBytes</i>	Specifies the size of the resource groups in megabytes. The minimum resource group size is 32 MB. The maximum resource group size is 2048 MB. A large resource group size may increase performance on very large file systems. If this is not specified, <code>mkfs.gfs2</code> chooses the resource group size based on the size of the file system: average size file systems will have 256 MB resource groups, and bigger file systems will have bigger RGs for better performance.
<b>-t</b>	<i>LockTableName</i>	A unique identifier that specifies the lock table field when you use the <b>lock_dlm</b> protocol; the <b>lock_nolock</b> protocol does not use this parameter. This parameter has two parts separated by a colon (no spaces) as follows: <i>ClusterName:FSName</i> . <i>ClusterName</i> is the name of the Red Hat cluster for which the GFS2 file system is being created; only members of this cluster are permitted to use this file system. The cluster name is set in the <code>/etc/cluster/cluster.conf</code> file via the <b>Cluster Configuration Tool</b> and displayed at the <b>Cluster Status Tool</b> in the Red Hat Cluster Suite cluster management GUI. <i>FSName</i> , the file system name, can be 1 to 16 characters in length, and the name must be unique among all file systems in the cluster.
<b>-u</b>	<i>MegaBytes</i>	Specifies the initial size of each journal's unlinked tag file.
<b>-v</b>		Displays command version information.

## 3.2. Mounting a File System

Before you can mount a GFS2 file system, the file system must exist (refer to [Section 3.1, “Making a File System”](#)), the volume where the file system exists must be activated, and the supporting clustering and locking systems must be started (refer to [Configuring and Managing a Red Hat Cluster](#)). After those requirements have been met, you can mount the GFS2 file system as you would any Linux file system.



To manipulate file ACLs, you must mount the file system with the **-o acl** mount option. If a file system is mounted without the **-o acl** mount option, users are allowed to view ACLs (with **getfacl**), but are not allowed to set them (with **setfacl**).

## Usage

### Mounting Without ACL Manipulation

```
mount BlockDevice MountPoint
```

### Mounting With ACL Manipulation

```
mount -o acl BlockDevice MountPoint
```

#### **-o acl**

GFS2-specific option to allow manipulating file ACLs.

#### *BlockDevice*

Specifies the block device where the GFS2 file system resides.

#### *MountPoint*

Specifies the directory where the GFS2 file system should be mounted.

## Example

In this example, the GFS2 file system on **/dev/vg01/lvol0** is mounted on the **/mygfs2** directory.

```
mount /dev/vg01/lvol0 /mygfs2
```

## Complete Usage

```
mount BlockDevice MountPoint -o option
```

The **-o option** argument consists of GFS2-specific options (refer to [Table 3.2, “GFS2-Specific Mount Options”](#)) or acceptable standard Linux **mount -o** options, or a combination of both. Multiple *option* parameters are separated by a comma and no spaces.



### Note

The **mount** command is a Linux system command. In addition to using GFS2-specific options described in this section, you can use other, standard, **mount** command options (for example, **-r**). For information about other Linux **mount** command options, see the Linux **mount** man page.

[Table 3.2, “GFS2-Specific Mount Options”](#) describes the available GFS2-specific **-o option** values that can be passed to GFS2 at mount time.



**Note**

This table includes descriptions of options that are used with local file systems only. For the Red Hat Enterprise Linux 5.5 release and later, Red Hat does not support the use of GFS2 as a single-node file system. Red Hat will continue to support single-node GFS2 file systems for existing customers.

Table 3.2. GFS2-Specific Mount Options

Option	Description
<b>acl</b>	Allows manipulating file ACLs. If a file system is mounted without the <b>acl</b> mount option, users are allowed to view ACLs (with <b>getfacl</b> ), but are not allowed to set them (with <b>setfacl</b> ).
<b>data=[ordered writeback]</b>	When <b>data=ordered</b> is set, the user data modified by a transaction is flushed to the disk before the transaction is committed to disk. This should prevent the user from seeing uninitialized blocks in a file after a crash. When <b>data=writeback</b> mode is set, the user data is written to the disk at any time after it is dirtied; this does not provide the same consistency guarantee as <b>ordered</b> mode, but it should be slightly faster for some workloads. The default value is <b>ordered</b> mode.
<b>ignore_local_fs</b> <b>Caution:</b> This option should <i>not</i> be used when GFS2 file systems are shared.	Forces GFS2 to treat the file system as a multihost file system. By default, using <b>lock_nolock</b> automatically turns on the <b>localflocks</b> flags.
<b>localflocks</b> <b>Caution:</b> This option should not be used when GFS2 file systems are shared.	Tells GFS2 to let the VFS (virtual file system) layer do all flock and fcntl. The <b>localflocks</b> flag is automatically turned on by <b>lock_nolock</b> . Note that the <b>localflocks</b> mount option affects only advisory <b>fcntl()/POSIX</b> locks and <b>flock</b> locks that are issued by applications. The internal locking that ensures coherency of data across the cluster by means of GFS2's <b>glock</b> abstraction is separate from and not affected by the <b>localflocks</b> setting. If you are unsure whether an application uses <b>fcntl()/POSIX</b> locks and thus requires that you mount your file system with the <b>localflocks</b> , you can use the <b>strace</b> utility to print out the system calls that are made during a test run of the application. Look for <b>fcntl</b> calls that have <b>F_GETLK</b> , <b>F_SETLK</b> , or <b>F_SETLKW</b> as the <b>cmd</b> argument. Note that GFS2 does not currently support either leases or mandatory locking.
<b>lockproto=LockModuleName</b>	Allows the user to specify which locking protocol to use with the file system. If <b>LockModuleName</b> is not

Option	Description
	specified, the locking protocol name is read from the file system superblock.
<b>locktable=</b> <i>LockTableName</i>	Allows the user to specify which locking table to use with the file system.
<b>quota=[off/account/on]</b>	Turns quotas on or off for a file system. Setting the quotas to be in the <b>account</b> state causes the per UID/GID usage statistics to be correctly maintained by the file system; limit and warn values are ignored. The default value is <b>off</b> .
<b>errors=panic withdraw</b>	When <b>errors=panic</b> is specified, file system errors will cause a kernel panic. The default behavior, which is the same as specifying <b>errors=withdraw</b> , is for the system to withdraw from the file system and make it inaccessible until the next reboot; in some cases the system may remain running. For information on the GFS2 withdraw function, see <a href="#">Section 3.14, “The GFS2 Withdraw Function”</a> .

### 3.3. Unmounting a File System

The GFS2 file system can be unmounted the same way as any Linux file system — by using the **umount** command.



#### Note

The **umount** command is a Linux system command. Information about this command can be found in the Linux **umount** command man pages.

### Usage

```
umount MountPoint
```

#### *MountPoint*

Specifies the directory where the GFS2 file system is currently mounted.

### 3.4. Special Considerations when Mounting GFS2 File Systems

GFS2 file systems that have been mounted manually rather than automatically through an entry in the **fstab** file will not be known to the system when file systems are unmounted at system shutdown. As a result, the GFS2 script will not unmount the GFS2 file system. After the GFS2 shutdown script is run, the standard shutdown process kills off all remaining user processes, including the cluster infrastructure, and tries to unmount the file system. This unmount will fail without the cluster infrastructure and the system will hang.

To prevent the system from hanging when the GFS2 file systems are unmounted, you should do one of the following:

- Always use an entry in the **fstab** file to mount the GFS2 file system.
- If a GFS2 file system has been mounted manually with the **mount** command, be sure to unmount the file system manually with the **umount** command before rebooting or shutting down the system.

If your file system hangs while it is being unmounted during system shutdown under these circumstances, perform a hardware reboot. It is unlikely that any data will be lost since the file system is synced earlier in the shutdown process.

### 3.5. GFS2 Quota Management

File-system quotas are used to limit the amount of file system space a user or group can use. A user or group does not have a quota limit until one is set. GFS2 keeps track of the space used by each user and group even when there are no limits in place. GFS2 updates quota information in a transactional way so system crashes do not require quota usages to be reconstructed.

To prevent a performance slowdown, a GFS2 node synchronizes updates to the quota file only periodically. The "fuzzy" quota accounting can allow users or groups to slightly exceed the set limit. To minimize this, GFS2 dynamically reduces the synchronization period as a "hard" quota limit is approached.

GFS2 uses its **gfs2\_quota** command to manage quotas. Other Linux quota facilities cannot be used with GFS2.

#### 3.5.1. Setting Quotas

Two quota settings are available for each user ID (UID) or group ID (GID): a *hard limit* and a *warn limit*.

A hard limit is the amount of space that can be used. The file system will not let the user or group use more than that amount of disk space. A hard limit value of *zero* means that no limit is enforced.

A warn limit is usually a value less than the hard limit. The file system will notify the user or group when the warn limit is reached to warn them of the amount of space they are using. A warn limit value of *zero* means that no limit is enforced.

Limits are set using the **gfs2\_quota** command. The command only needs to be run on a single node where GFS2 is mounted.

By default, quota enforcement is not set on GFS2 file systems. To enable quota accounting, use the **quota=** of the **mount** command when mounting the GFS2 file system, as described in [Section 3.5.4, "Enabling/Disabling Quota Enforcement"](#).

### Usage

#### Setting Quotas, Hard Limit

```
gfs2_quota limit -u User -l Size -f MountPoint
```

```
gfs2_quota limit -g Group -l Size -f MountPoint
```

### Setting Quotas, Warn Limit

```
gfs2_quota warn -u User -l Size -f MountPoint
```

```
gfs2_quota warn -g Group -l Size -f MountPoint
```

#### *User*

A user ID to limit or warn. It can be either a user name from the password file or the UID number.

#### *Group*

A group ID to limit or warn. It can be either a group name from the group file or the GID number.

#### *Size*

Specifies the new value to limit or warn. By default, the value is in units of megabytes. The additional **-k**, **-s** and **-b** flags change the units to kilobytes, sectors, and file system blocks, respectively.

#### *MountPoint*

Specifies the GFS2 file system to which the actions apply.

## Examples

This example sets the hard limit for user *Bert* to 1024 megabytes (1 gigabyte) on file system */mygfs2*.

```
gfs2_quota limit -u Bert -l 1024 -f /mygfs2
```

This example sets the warn limit for group ID 21 to 50 kilobytes on file system */mygfs2*.

```
gfs2_quota warn -g 21 -l 50 -k -f /mygfs2
```

## 3.5.2. Displaying Quota Limits and Usage

Quota limits and current usage can be displayed for a specific user or group using the **gfs2\_quota get** command. The entire contents of the quota file can also be displayed using the **gfs2\_quota list** command, in which case all IDs with a non-zero hard limit, warn limit, or value are listed.

### Usage

#### Displaying Quota Limits for a User

```
gfs2_quota get -u User -f MountPoint
```

#### Displaying Quota Limits for a Group

## Chapter 3. Managing GFS2

---

```
gfs2_quota get -g Group -f MountPoint
```

### Displaying Entire Quota File

```
gfs2_quota list -f MountPoint
```

#### *User*

A user ID to display information about a specific user. It can be either a user name from the password file or the UID number.

#### *Group*

A group ID to display information about a specific group. It can be either a group name from the group file or the GID number.

#### *MountPoint*

Specifies the GFS2 file system to which the actions apply.

## Command Output

GFS2 quota information from the **gfs2\_quota** command is displayed as follows:

```
user User: limit:LimitSize warn:WarnSize value:Value  
group Group: limit:LimitSize warn:WarnSize value:Value
```

The *LimitSize*, *WarnSize*, and *Value* numbers (values) are in units of megabytes by default. Adding the **-k**, **-s**, or **-b** flags to the command line change the units to kilobytes, sectors, or file system blocks, respectively.

#### *User*

A user name or ID to which the data is associated.

#### *Group*

A group name or ID to which the data is associated.

#### *LimitSize*

The hard limit set for the user or group. This value is zero if no limit has been set.

#### *Value*

The actual amount of disk space used by the user or group.

## Comments

When displaying quota information, the **gfs2\_quota** command does not resolve UIDs and GIDs into names if the **-n** option is added to the command line.

Space allocated to GFS2's hidden files can be left out of displayed values for the root UID and GID by adding the **-d** option to the command line. This is useful when trying to match the numbers from **gfs2\_quota** with the results of a **du** command.

## Examples

This example displays quota information for all users and groups that have a limit set or are using any disk space on file system **/mygfs2**.

```
gfs2_quota list -f /mygfs2
```

This example displays quota information in sectors for group **users** on file system **/mygfs2**.

```
gfs2_quota get -g users -f /mygfs2 -s
```

### 3.5.3. Synchronizing Quotas

GFS2 stores all quota information in its own internal file on disk. A GFS2 node does not update this quota file for every file system write; rather, it updates the quota file once every 60 seconds. This is necessary to avoid contention among nodes writing to the quota file, which would cause a slowdown in performance.

As a user or group approaches their quota limit, GFS2 dynamically reduces the time between its quota-file updates to prevent the limit from being exceeded. The normal time period between quota synchronizations is a tunable parameter, **quota\_quantum**, and can be changed using the **gfs2\_tool** command. By default, the time period is 60 seconds. Also, the **quota\_quantum** parameter must be set on each node and each time the file system is mounted. (Changes to the **quota\_quantum** parameter are not persistent across unmounts.)

You can use the **gfs2\_quota sync** command to synchronize the quota information from a node to the on-disk quota file between the automatic updates performed by GFS2.

## Usage

### Synchronizing Quota Information

```
gfs2_quota sync -f MountPoint
```

*MountPoint*

Specifies the GFS2 file system to which the actions apply.

### Tuning the Time Between Synchronizations

```
gfs2_tool settune MountPoint quota_quantum Seconds
```

*MountPoint*

Specifies the GFS2 file system to which the actions apply.

*Seconds*

Specifies the new time period between regular quota-file synchronizations by GFS2. Smaller values may increase contention and slow down performance.

## Examples

This example synchronizes the quota information from the node it is run on to file system **/mygfs2**.

```
gfs2_quota sync -f /mygfs2
```

This example changes the default time period between regular quota-file updates to one hour (3600 seconds) for file system **/mygfs2** on a single node.

```
gfs2_tool settune /mygfs2 quota_quantum 3600
```

### 3.5.4. Enabling/Disabling Quota Enforcement

In GFS2 file systems, quota enforcement is disabled by default. To enable quota enforcement for a file system, mount the file system with the **quota=on** option specified.

#### Usage

```
mount -o quota=on BlockDevice MountPoint
```

To mount a file system with quota enforcement disabled, mount the file system with the **quota=off** option specified. This is the default setting.

```
mount -o quota=off BlockDevice MountPoint
```

#### **-o quota={on|off}**

Specifies that quota enforcement is enabled or disabled when the file system is mounted.

#### *BlockDevice*

Specifies the block device where the GFS2 file system resides.

#### *MountPoint*

Specifies the directory where the GFS2 file system should be mounted.

#### Examples

In this example, the GFS2 file system on **/dev/vg01/lvol0** is mounted on the **/mygfs2** directory with quota enforcement enabled.

```
mount -o quota=on /dev/vg01/lvol0 /mygfs2
```

### 3.5.5. Enabling Quota Accounting

It is possible to keep track of disk usage and maintain quota accounting for every user and group without enforcing the limit and warn values. To do this, mount the file system with the **quota=account** option specified.

#### Usage

```
mount -o quota=account BlockDevice MountPoint
```



**-o quota=account**

Specifies that user and group usage statistics are maintained by the file system, even though the quota limits are not enforced.

*BlockDevice*

Specifies the block device where the GFS2 file system resides.

*MountPoint*

Specifies the directory where the GFS2 file system should be mounted.

**Example**

In this example, the GFS2 file system on `/dev/vg01/lvo10` is mounted on the `/mygfs2` directory with quota accounting enabled.

```
mount -o quota=account /dev/vg01/lvo10 /mygfs2
```

## 3.6. Growing a File System

The `gfs2_grow` command is used to expand a GFS2 file system after the device where the file system resides has been expanded. Running a `gfs2_grow` command on an existing GFS2 file system fills all spare space between the current end of the file system and the end of the device with a newly initialized GFS2 file system extension. When the fill operation is completed, the resource index for the file system is updated. All nodes in the cluster can then use the extra storage space that has been added.

The `gfs2_grow` command must be run on a mounted file system, but only needs to be run on one node in a cluster. All the other nodes sense that the expansion has occurred and automatically start using the new space.

**Note**

Once you have created a GFS2 file system with the `mkfs.gfs2` command, you cannot decrease the size of the file system.

## Usage

```
gfs2_grow MountPoint
```

*MountPoint*

Specifies the GFS2 file system to which the actions apply.

**Comments**

Before running the `gfs2_grow` command:

- Back up important data on the file system.

- Determine the volume that is used by the file system to be expanded by running a **df** *MountPoint* command.
- Expand the underlying cluster volume with LVM. For information on administering LVM volumes, see *Logical Volume Manager Administration*.

After running the **gfs2\_grow** command, run a **df** command to check that the new space is now available in the file system.

### Examples

In this example, the file system on the **/mygfs2fs** directory is expanded.

```
[root@dash-01 ~]# gfs2_grow /mygfs2fs
FS: Mount Point: /mygfs2fs
FS: Device:      /dev/mapper/gfs2testvg-gfs2testlv
FS: Size:        524288 (0x80000)
FS: RG size:     65533 (0xffffd)
DEV: Size:       655360 (0xa0000)
The file system grew by 512MB.
gfs2_grow complete.
```

### Complete Usage

```
gfs2_grow [Options] {MountPoint | Device} [MountPoint | Device]
```

#### *MountPoint*

Specifies the directory where the GFS2 file system is mounted.

#### *Device*

Specifies the device node of the file system.

[Table 3.3, “GFS2-specific Options Available While Expanding A File System”](#) describes the GFS2-specific options that can be used while expanding a GFS2 file system.

Table 3.3. GFS2-specific Options Available While Expanding A File System

Option	Description
<b>-h</b>	Help. Displays a short usage message.
<b>-q</b>	Quiet. Turns down the verbosity level.
<b>-r MegaBytes</b>	Specifies the size of the new resource group. The default size is 256MB.
<b>-T</b>	Test. Do all calculations, but do not write any data to the disk and do not expand the file system.
<b>-V</b>	Displays command version information.

## 3.7. Adding Journals to a File System

The **gfs2\_jadd** command is used to add journals to a GFS2 file system. You can add journals to a GFS2 file system dynamically at any point without expanding the underlying logical volume. The **gfs2\_jadd** command must be run on a mounted file system, but it needs to be run on only one node in the cluster. All the other nodes sense that the expansion has occurred.



### Note

If a GFS2 file system is full, the **gfs2\_jadd** will fail, even if the logical volume containing the file system has been extended and is larger than the file system. This is because in a GFS2 file system, journals are plain files rather than embedded metadata, so simply extending the underlying logical volume will not provide space for the journals.

Before adding journals to a GFS file system, you can use the **journals** option of the **gfs2\_tool** to find out how many journals the GFS2 file system currently contains. The following example displays the number and size of the journals in the file system mounted at **/mnt/gfs2**.

```
[root@roth-01 ../cluster/gfs2]# gfs2_tool journals /mnt/gfs2
journal2 - 128MB
journal1 - 128MB
journal0 - 128MB
3 journal(s) found.
```

## Usage

```
gfs2_jadd -j Number MountPoint
```

### *Number*

Specifies the number of new journals to be added.

### *MountPoint*

Specifies the directory where the GFS2 file system is mounted.

## Examples

In this example, one journal is added to the file system on the **/mygfs2** directory.

```
gfs2_jadd -j1 /mygfs2
```

In this example, two journals are added to the file system on the **/mygfs2** directory.

```
gfs2_jadd -j2 /mygfs2
```

## Complete Usage

```
gfs2_jadd [Options] {MountPoint | Device} [MountPoint | Device]
```

### MountPoint

Specifies the directory where the GFS2 file system is mounted.

### Device

Specifies the device node of the file system.

Table 3.4, “GFS2-specific Options Available When Adding Journals” describes the GFS2-specific options that can be used when adding journals to a GFS2 file system.

Table 3.4. GFS2-specific Options Available When Adding Journals

Flag	Parameter	Description
-h		Help. Displays short usage message.
-J	<i>MegaBytes</i>	Specifies the size of the new journals in megabytes. Default journal size is 128 megabytes. The minimum size is 32 megabytes. To add journals of different sizes to the file system, the <b>gfs2_jadd</b> command must be run for each size journal. The size specified is rounded down so that it is a multiple of the journal-segment size that was specified when the file system was created.
-j	<i>Number</i>	Specifies the number of new journals to be added by the <b>gfs2_jadd</b> command. The default value is 1.
-q		Quiet. Turns down the verbosity level.
-v		Displays command version information.

## 3.8. Data Journaling

Ordinarily, GFS2 writes only metadata to its journal. File contents are subsequently written to disk by the kernel's periodic sync that flushes file system buffers. An **fsync()** call on a file causes the file's data to be written to disk immediately. The call returns when the disk reports that all data is safely written.

Data journaling can result in a reduced **fsync()** time for very small files because the file data is written to the journal in addition to the metadata. This advantage rapidly reduces as the file size increases. Writing to medium and larger files will be much slower with data journaling turned on.

Applications that rely on **fsync()** to sync file data may see improved performance by using data journaling. Data journaling can be enabled automatically for any GFS2 files created in a flagged directory (and all its subdirectories). Existing files with zero length can also have data journaling turned on or off.

Enabling data journaling on a directory sets the directory to "inherit jdata", which indicates that all files and directories subsequently created in that directory are journaled. You can enable and disable data journaling on a file with the **chattr** command.

The following commands enable data journaling on the `/mnt/gfs2/gfs2_dir/newfile` file and then check whether the flag has been set properly.

```
[root@roth-01 ~]# chattr +j /mnt/gfs2/gfs2_dir/newfile
[root@roth-01 ~]# lsattr /mnt/gfs2/gfs2_dir
```

```
-----j--- /mnt/gfs2/gfs2_dir/newfile
```

The following commands disable data journaling on the `/mnt/gfs2/gfs2_dir/newfile` file and then check whether the flag has been set properly.

```
[root@roth-01 ~]# chattr -j /mnt/gfs2/gfs2_dir/newfile
[root@roth-01 ~]# lsattr /mnt/gfs2/gfs2_dir
-----j--- /mnt/gfs2/gfs2_dir/newfile
```

You can also use the **chattr** command to set the **j** flag on a directory. When you set this flag for a directory, all files and directories subsequently created in that directory are journaled. The following set of commands sets the **j** flag on the **gfs2\_dir** directory, then checks whether the flag has been set properly. After this, the commands create a new file called **newfile** in the `/mnt/gfs2/gfs2_dir` directory and then check whether the **j** flag has been set for the file. Since the **j** flag is set for the directory, then **newfile** should also have journaling enabled.

```
[root@roth-01 ~]# chattr -j /mnt/gfs2/gfs2_dir
[root@roth-01 ~]# lsattr /mnt/gfs2
-----j--- /mnt/gfs2/gfs2_dir
[root@roth-01 ~]# touch /mnt/gfs2/gfs2_dir/newfile
[root@roth-01 ~]# lsattr /mnt/gfs2/gfs2_dir
-----j--- /mnt/gfs2/gfs2_dir/newfile
```

## 3.9. Configuring **atime** Updates

Each file inode and directory inode has three time stamps associated with it:

- **ctime** — The last time the inode status was changed
- **mtime** — The last time the file (or directory) data was modified
- **atime** — The last time the file (or directory) data was accessed

If **atime** updates are enabled as they are by default on GFS2 and other Linux file systems then every time a file is read, its inode needs to be updated.

Because few applications use the information provided by **atime**, those updates can require a significant amount of unnecessary write traffic and file locking traffic. That traffic can degrade performance; therefore, it may be preferable to turn off or reduce the frequency of **atime** updates.

To reduce the effects of **atime** updating, you can mount with the **noatime** option, which disables **atime** updates on that file system. The **noatime** Linux mount option can be specified when the file system is mounted, which disables **atime** updates on that file system.

### Usage

```
mount BlockDevice MountPoint -o noatime
```

#### *BlockDevice*

Specifies the block device where the GFS2 file system resides.

### *MountPoint*

Specifies the directory where the GFS2 file system should be mounted.

### Example

In this example, the GFS2 file system resides on the `/dev/vg01/lvol10` and is mounted on directory `/mygfs2` with `atime` updates turned off.

```
mount /dev/vg01/lvol10 /mygfs2 -o noatime
```

## 3.10. Suspending Activity on a File System

You can suspend write activity to a file system by using the `gfs2_tool freeze` command. Suspending write activity allows hardware-based device snapshots to be used to capture the file system in a consistent state. The `gfs2_tool unfreeze` command ends the suspension.

### Usage

#### Start Suspension

```
gfs2_tool freeze MountPoint
```

#### End Suspension

```
gfs2_tool unfreeze MountPoint
```

### *MountPoint*

Specifies the file system.

### Examples

This example suspends writes to file system `/mygfs2`.

```
gfs2_tool freeze /mygfs2
```

This example ends suspension of writes to file system `/mygfs2`.

```
gfs2_tool unfreeze /mygfs2
```

## 3.11. Repairing a File System

When nodes fail with the file system mounted, file system journaling allows fast recovery. However, if a storage device loses power or is physically disconnected, file system corruption may occur. (Journaling cannot be used to recover from storage subsystem failures.) When that type of corruption occurs, you can recover the GFS2 file system by using the `fsck.gfs2` command.

**Important**

The **fsck.gfs2** command must be run only on a file system that is unmounted from all nodes.

**Important**

You should not check a GFS2 file system at boot time with the **fsck.gfs2** command. The **fsck.gfs2** command can not determine at boot time whether the file system is mounted by another node in the cluster. You should run the **fsck.gfs2** command manually only after the system boots.

To ensure that the **fsck.gfs2** command does not run on a GFS2 file system at boot time, modify the **/etc/fstab** file so that the final two columns for a GFS2 file system mount point show "0 0" rather than "1 1" (or any other numbers), as in the following example:

```
/dev/VG12/lv_svr_home /svr_home gfs2 defaults,noatime,nodiratime,noquota
0 0
```

**Note**

If you have previous experience using the **gfs\_fsck** command on GFS file systems, note that the **fsck.gfs2** command differs from some earlier releases of **gfs\_fsck** in the following ways:

- Pressing **Ctrl+C** while running the **fsck.gfs2** interrupts processing and displays a prompt asking whether you would like to abort the command, skip the rest of the current pass, or continue processing.
- You can increase the level of verbosity by using the **-v** flag. Adding a second **-v** flag increases the level again.
- You can decrease the level of verbosity by using the **-q** flag. Adding a second **-q** flag decreases the level again.
- The **-n** option opens a file system as read-only and answers **no** to any queries automatically. The option provides a way of trying the command to reveal errors without actually allowing the **fsck.gfs2** command to take effect.

Refer to the **fsck.gfs2** man page for additional information about other command options.

memory you will need to run the **fsck.gfs2** command on your file system, determine how many blocks the file system contains and multiply that number by 5/8.

For example, to determine approximately how much memory is required to run the **fsck.gfs2** command on a GFS2 file system that is 16TB with a block size of 4K, first determine how many blocks of memory the file system contains by dividing 16Tb by 4K:

```
17592186044416 / 4096 = 4294967296
```

Since this file system contains 4294967296 blocks, multiply that number by 5/8 to determine how many bytes of memory are required:

```
4294967296 * 5/8 = 2684354560
```

This file system requires approximately 2.6GB of free memory to run the **fsck.gfs2** command. Note that if the block size was 1K, running the **fsck.gfs2** command would require four times the memory, or approximately 11GB.

## Usage

```
fsck.gfs2 -y BlockDevice
```

### -y

The **-y** flag causes all questions to be answered with **yes**. With the **-y** flag specified, the **fsck.gfs2** command does not prompt you for an answer before making changes.

### *BlockDevice*

Specifies the block device where the GFS2 file system resides.

## Example

In this example, the GFS2 file system residing on block device **/dev/testvol/testlv** is repaired. All queries to repair are automatically answered with **yes**.

```
[root@dash-01 ~]# fsck.gfs2 -y /dev/testvg/testlv
Initializing fsck
Validating Resource Group index.
Level 1 RG check.
(level 1 passed)
Clearing journals (this may take a while)...
Journals cleared.
Starting pass1
Pass1 complete
Starting pass1b
Pass1b complete
Starting pass1c
Pass1c complete
Starting pass2
Pass2 complete
Starting pass3
Pass3 complete
Starting pass4
Pass4 complete
Starting pass5
Pass5 complete
Writing changes to disk
```



```
fscck.gfs2 complete
```

## 3.12. Bind Mounts and Context-Dependent Path Names

GFS2 file systems do not provide support for Context-Dependent Path Names (CDPNs), which allow you to create symbolic links that point to variable destination files or directories. For this functionality in GFS2, you can use the **bind** option of the **mount** command.

The **bind** option of the **mount** command allows you to remount part of a file hierarchy at a different location while it is still available at the original location. The format of this command is as follows.

```
mount --bind olddir newdir
```

After executing this command, the contents of the *olddir* directory are available at two locations: *olddir* and *newdir*. You can also use this option to make an individual file available at two locations.

For example, after executing the following commands the contents of **/root/tmp** will be identical to the contents of the previously mounted **/var/log** directory.

```
[root@mencryfa ~]# cd ~root
[root@mencryfa ~]# mkdir ./tmp
[root@mencryfa ~]# mount --bind /var/log /tmp
```

Alternately, you can use an entry in the **/etc/fstab** file to achieve the same results at mount time. The following **/etc/fstab** entry will result in the contents of **/root/tmp** being identical to the contents of the **/var/log** directory.

```
/var/log          /root/tmp        none    bind          0 0
```

After you have mounted the file system, you can use the **mount** command to see that the file system has been mounted, as in the following example.

```
[root@mencryfa ~]# mount | grep /tmp
/var/log on /root/tmp type none (rw,bind)
```

With a file system that supports Context-Dependent Path Names, you might have defined the **/bin** directory as a Context-Dependent Path Name that would resolve to one of the following paths, depending on the system architecture.

```
/usr/i386-bin
/usr/x86_64-bin
/usr/ppc64-bin
```

You can achieve this same functionality by creating an empty **/bin** directory. Then, using a script or an entry in the **/etc/fstab** file, you can mount each of the individual architecture directories onto the **/bin** directory with a **mount -bind** command. For example, you can use the following command as a line in a script.

```
mount --bind /usr/i386-bin /bin
```

Alternately, you can use the following entry in the `/etc/fstab` file.

```
/usr/i386-bin          /bin          none          bind          0 0
```

A bind mount can provide greater flexibility than a Context-Dependent Path Name, since you can use this feature to mount different directories according to any criteria you define (such as the value of `%fill` for the file system). Context-Dependent Path Names are more limited in what they can encompass. Note, however, that you will need to write your own script to mount according to a criteria such as the value of `%fill`.



### Warning

When you mount a file system with the `bind` option and the original file system was mounted `rw`, the new file system will also be mounted `rw` even if you use the `ro` flag; the `ro` flag is silently ignored. In this case, the new file system might be marked as `ro` in the `/proc/mounts` directory, which may be misleading.

## 3.13. Bind Mounts and File System Mount Order

When you use the `bind` option of the `mount` command, you must be sure that the file systems are mounted in the correct order. In the following example, the `/var/log` directory must be mounted before executing the bind mount on the `/tmp` directory:

```
# mount --bind /var/log /tmp
```

The ordering of file system mounts is determined as follows:

- In general, file system mount order is determined by the order in which the file systems appear in the `fstab` file. The exceptions to this ordering are file systems mounted with the `_netdev` flag or file systems that have their own `init` scripts.
- A file system with its own `init` script is mounted later in the initialization process, after the file systems in the `fstab` file.
- File systems mounted with the `_netdev` flag are mounted when the network has been enabled on the system.

If your configuration requires that you create a bind mount on which to mount a GFS2 file system, you can order your `fstab` file as follows:

1. Mount local file systems that are required for the bind mount.
2. Bind mount the directory on which to mount the GFS2 file system.
3. Mount the GFS2 file system.

If your configuration requires that you bind mount a local directory or file system onto a GFS2 file system, listing the file systems in the correct order in the `fstab` file will not mount the file systems

correctly since the GFS2 file system will not be mounted until the GFS2 **init** script is run. In this case, you should write an **init** script to execute the bind mount so that the bind mount will not take place until after the GFS2 file system is mounted.

The following script is an example of a custom **init** script. This script performs a bind mount of two directories onto two directories of a GFS2 file system. In this example, there is an existing GFS2 mount point at **/mnt/gfs2a**, which is mounted when the GFS2 **init** script runs, after cluster startup.

In this example script, the values of the **chkconfig** statement indicate the following:

- 345 indicates the run levels that the script will be started in
- 29 is the start priority, which in this case indicates that the script will run at startup time after the GFS2 **init** script, which has a start priority of 26
- 73 is the stop priority, which in this case indicates that the script will be stopped during shutdown before the GFS2 script, which has a stop priority of 74

The start and stop values indicate that you can manually perform the indicated action by executing a **service start** and a **service stop** command. For example, if the script is named **fredwilma**, then you can execute **service fredwilma start**.

This script should be put in the **/etc/init.d** directory with the same permissions as the other scripts in that directory. You can then execute a **chkconfig on** command to link the script to the indicated run levels. For example, if the script is named **fredwilma**, then you can execute **chkconfig fredwilma on**.

```
#!/bin/bash
#
# chkconfig: 345 29 73
# description: mount/unmount my custom bind mounts onto a gfs2 subdirectory
#
### BEGIN INIT INFO
# Provides:
### END INIT INFO

. /etc/init.d/functions
case "$1" in
  start)
    # In this example, fred and wilma want their home directories
    # bind-mounted over the gfs2 directory /mnt/gfs2a, which has
    # been mounted as /mnt/gfs2a
    mkdir -p /mnt/gfs2a/home/fred &> /dev/null
    mkdir -p /mnt/gfs2a/home/wilma &> /dev/null
    /bin/mount --bind /mnt/gfs2a/home/fred /home/fred
    /bin/mount --bind /mnt/gfs2a/home/wilma /home/wilma
    ;;

  stop)
    /bin/umount /mnt/gfs2a/home/fred
    /bin/umount /mnt/gfs2a/home/wilma
    ;;

  status)
    ;;

  restart)
    $0 stop
    $0 start
    ;;
)
```

```
reload)
    $0 start
    ;;
*)
    echo $"Usage: $0 {start|stop|restart|reload|status}"
    exit 1
esac

exit 0
```

### 3.14. The GFS2 Withdraw Function

The GFS2 *withdraw* function is a data integrity feature of GFS2 file systems in a cluster. If the GFS2 kernel module detects an inconsistency in a GFS2 file system following an I/O operation, the file system becomes unavailable to the cluster. The I/O operation stops and the system waits for further I/O operations to stop with an error, preventing further damage. When this occurs, you can stop any other services or applications manually, after which you can reboot and remount the GFS2 file system to replay the journals. If the problem persists, you can unmount the file system from all nodes in the cluster and perform file system recovery with the **fsck.gfs2** command. The GFS withdraw function is less severe than a kernel panic, which would cause another node to fence the node.

If your system is configured with the **gfs2** startup script enabled and the GFS2 file system is included in the **/etc/fstab** file, the GFS2 file system will be remounted when you reboot. If the GFS2 file system withdrew because of perceived file system corruption, it is recommended that you run the **fsck.gfs2** command before remounting the file system. In this case, in order to prevent your file system from remounting at boot time, you can perform the following procedure:

1. Temporarily disable the startup script on the affected node with the following command:

```
# chkconfig gfs2 off
```

2. Reboot the affected node, starting the cluster software. The GFS2 file system will not be mounted.
3. Unmount the file system from every node in the cluster.
4. Run the **fsck.gfs2** on the file system from one node only to ensure there is no file system corruption.
5. Re-enable the startup script on the affected node by running the following command:

```
# chkconfig gfs2 on
```

6. Remount the GFS2 file system from all nodes in the cluster.

An example of an inconsistency that would yield a GFS2 withdraw is an incorrect block count. When the GFS kernel deletes a file from a file system, it systematically removes all the data and metadata blocks associated with that file. When it is done, it checks the block count. If the block count is not one (meaning all that is left is the disk inode itself), that indicates a file system inconsistency since the block count did not match the list of blocks found.

You can override the GFS2 withdraw function by mounting the file system with the **-o errors=panic** option specified. When this option is specified, any errors that would normally

cause the system to withdraw cause the system to panic instead. This stops the node's cluster communications, which causes the node to be fenced.

Internally, the GFS2 withdraw function works by having the kernel send a message to the **gfs\_control**d daemon requesting withdraw. The **gfs\_control**d daemon runs the **dmsetup** program to place the device mapper error target underneath the filesystem preventing further access to the block device. It then tells the kernel that this has been completed. This is the reason for the GFS2 support requirement to always use a CLVM device under GFS2, since otherwise it is not possible to insert a device mapper target.

The purpose of the device mapper error target is to ensure that all future I/O operations will result in an I/O error that will allow the filesystem to be unmounted in an orderly fashion. As a result, when the withdraw occurs, it is normal to see a number of I/O errors from the device mapper device reported in the system logs.

Occasionally, the withdraw may fail if it is not possible for the **dmsetup** program to insert the error target as requested. This can happen if there is a shortage of memory at the point of the withdraw and memory cannot be reclaimed due to the problem that triggered the withdraw in the first place.

A withdraw does not always mean that there is an error in GFS2. Sometimes the withdraw function can be triggered by device I/O errors relating to the underlying block device. It is highly recommended to check the logs to see if that is the case if a withdraw occurs.



# Appendix A. Converting a File System from GFS to GFS2

Use the `gfs2_convert` command. To convert a GFS file system to a GFS2 file system.



## Warning

Before converting the GFS file system, you must back up the file system, since the conversion process is irreversible and any errors encountered during the conversion can result in the abrupt termination of the program and consequently an unusable file system.

Before converting the GFS file system, you must use the `gfs_fsck` command to check the file system and fix any errors.

If the conversion from GFS to GFS2 is interrupted by a power failure or any other issue, restart the conversion tool. Do not attempt to execute the `fsck.gfs2` command on the file system until the conversion is complete.



## Context-Dependent Path Names

GFS2 file systems do not provide support for Context-Dependent Path Names (CDPNs), which allow you to create symbolic links that point to variable destination files or directories. To achieve the same functionality as CDPNs in GFS2 file systems, you can use the `bind` option of the `mount` command.

The `gfs2_convert` command identifies CDPNs and replaces them with empty directories with the same name. In order to configure bind mounts to replace the CDPNs, however, you need to know the full paths of the link targets of the CDPNs you are replacing. Before converting your file system, you can use the `find` command to identify the links.

The following command lists the symlinks that point to a `hostname` CDPN:

```
[root@smoke-01 gfs]# find /mnt/gfs -lname @hostname
/mnt/gfs/log
```

Similarly, you can execute the `find` command for other CDPNs (`mach`, `os`, `sys`, `uid`, `gid`, `jid`). Note that since CDPN names can be of the form `@hostname` or `{hostname}`, you will need to run the `find` command for each variant.

For more information on bind mounts and context-dependent pathnames in GFS2, see [Section 3.12, “Bind Mounts and Context-Dependent Path Names”](#).

When converting full or nearly full file systems, it is possible that there will not be enough space available to fit all the GFS2 file system data structures. In such cases, the size of all the journals is reduced uniformly such that everything fits in the available space.

## Appendix A. Converting a File System from GFS to GFS2

---

1. Make a backup of your existing GFS file system.
2. Unmount the GFS file system from all nodes in the cluster.
3. Execute the **gfs\_fsck** command on the GFS file system to ensure there is no file system corruption.
4. Execute **gfs2\_convert gfsfilesystem**. The system will display warnings and confirmation questions before converting *gfsfilesystem* to GFS2.

The following example converts a GFS filesystem on block device **/dev/shell\_vg/500g** to a GFS2 filesystem.

```
[root@shell-01 ~]# /root/cluster/gfs2/convert/gfs2_convert /dev/shell_vg/500g
gfs2_convert version 2 (built May 10 2010 10:05:40)
Copyright (C) Red Hat, Inc. 2004-2006 All rights reserved.

Examining file system.....
This program will convert a gfs1 filesystem to a gfs2 filesystem.
WARNING: This can't be undone. It is strongly advised that you:

    1. Back up your entire filesystem first.
    2. Run gfs_fsck first to ensure filesystem integrity.
    3. Make sure the filesystem is NOT mounted from any node.
    4. Make sure you have the latest software versions.
Convert /dev/shell_vg/500g from GFS1 to GFS2? (y/n)y
Converting resource groups.....
Converting inodes.
24208 inodes from 1862 rgs converted.
Fixing file and directory information.
18 cdpn symlinks moved to empty directories.
Converting journals.
Converting journal space to rg space.
Writing journal #1...done.
Writing journal #2...done.
Writing journal #3...done.
Writing journal #4...done.
Building GFS2 file system structures.
Removing obsolete GFS1 file system structures.
Committing changes to disk.
/dev/shell_vg/500g: filesystem converted successfully to gfs2.
```



---

# Appendix B. Revision History

**Revision 6.0-3 Thu Feb 16 2012**

Steven Levine [slevine@redhat.com](mailto:slevine@redhat.com)

Release for GA of Red Hat Enterprise Linux 5.8

**Revision 6.0-2 Thu Dec 15 2011**

Steven Levine [slevine@redhat.com](mailto:slevine@redhat.com)

Beta release of Red Hat Enterprise Linux 5.8

**Revision 6.0-1 Thu Nov 10 2011**

Steven Levine [slevine@redhat.com](mailto:slevine@redhat.com)

Resolves: #758841

Notes CLVM requirement for clustered environment.

Resolves: #736156

Adds note warning not to check a GFS2 file system at boot time.

Resolves: #750328

Corrects small typographical errors.

**Revision 5.0-1 Thu Jul 21 2011**

Steven Levine [slevine@redhat.com](mailto:slevine@redhat.com)

Resolves: #664848

Updates documentation of GFS to GFS2 conversion process to clarify procedure for replacing Context-Dependent Path Names.

Resolves: #676131

Clarifies section on the withdraw function.

Resolves: #654490

Corrects minor typos.

**Revision 4.0-1 Thu Dec 23 2010**

Steven Levine [slevine@redhat.com](mailto:slevine@redhat.com)

Resolves: #661519

Corrects maximum file system size numbers.

Resolves: #588374

Updates gfs2\_convert documentation appendix.

Resolves: #606224

Removes relatime documentation.

**Revision 3.0-2 Tue Aug 3 2010**

Steven Levine [slevine@redhat.com](mailto:slevine@redhat.com)

Resolves: #562251

Adds information about the **locallocks** mount option and when it may be required.

**Revision 3.0-1 Thu Mar 11 2010**

Steven Levine [slevine@redhat.com](mailto:slevine@redhat.com)

Resolves: #546687

Documents features for the Red Hat Enterprise Linux 5.5 release.

## Appendix B. Revision History

---

Resolves: #568179

Adds note clarifying support policy for single-node system.

Resolves: #562199

Adds note clarifying 16-node limitation.

Resolves: #515348

Documents new -o errors mount option.

Resolves: #573750

Documents memory requirements for fsck.gfs2.

Resolves: #574470

Clarifies issue of gfs2 requiring CLVM for Red Hat support.

### Revision 2.0-1 Tue Aug 18 2009

Steven Levine [slevine@redhat.com](mailto:slevine@redhat.com)

Resolves: #515807

Adds note clarifying that you cannot reduce the size of an existing file system.

Resolves: #480002

Adds caveat about unmounting a file system manually if you mounted it manually, adds section on bind mounting a non-GFS2 file system to a GFS2 file system, adds sample custom init script.

Resolves: #458604

Adds section on GFS2 withdraw function.

Resolves: #498292

Clarifies documentation on adding journals.

### Revision 1.0-1 Thu Jan 29 2009

---

# Index

## A

- acl mount option, 14
- adding journals to a file system, 24
- atime, configuring updates, 27
  - mounting with noatime , 27
- audience, v

## B

- bind mount
  - mount order, 32
- bind mounts, 31

## C

- configuration, before, 2
- configuration, initial,
  - prerequisite tasks, 7

## D

- data journaling, 26

## F

- feedback, vi, vi
- file system
  - adding journals, 24
  - atime, configuring updates, 27
    - mounting with noatime , 27
  - bind mounts, 31
  - context-dependent path names (CDPNs), 31
  - data journaling, 26
  - growing, 23
  - making, 11
  - mount order, 32
  - mounting, 14, 17
  - quota management, 18
    - displaying quota limits, 19
    - enabling quota accounting, 22
    - enabling/disabling quota enforcement, 22
    - setting quotas, 18
    - synchronizing quotas, 21
  - repairing, 28
  - suspending activity, 28
  - unmounting, 17, 17
- fsck.gfs2 command, 28

## G

- GFS2
  - atime, configuring updates, 27
    - mounting with noatime , 27
  - managing,
  - quota management, 18

- displaying quota limits, 19
- enabling quota accounting, 22
- enabling/disabling quota enforcement, 22
- setting quotas, 18
- synchronizing quotas, 21
- withdraw function, 34

- GFS2 file system maximum size,
- GFS2-specific options for adding journals table, 26
- GFS2-specific options for expanding file systems table, 24
- gfs2\_grow command, 23
- gfs2\_jadd command, 24
- gfs2\_quota command, 18
- growing a file system, 23

## I

- initial tasks
  - setup, initial, 7
- introduction,
  - audience, v

## M

- making a file system, 11
- managing GFS2,
- maximum size, GFS2 file system,
- mkfs command, 11
- mkfs.gfs2 command options table, 13
- mount command, 14
- mount table, 16
- mounting a file system, 14, 17

## O

- overview,
  - configuration, before, 2

## P

- path names, context-dependent (CDPNs), 31
- preface (see introduction)
- prerequisite tasks
  - configuration, initial, 7

## Q

- quota management, 18
  - displaying quota limits, 19
  - enabling quota accounting, 22
  - enabling/disabling quota enforcement, 22
  - setting quotas, 18
  - synchronizing quotas, 21
- quota= mount option, 18
- quota\_quantum tunable parameter, 21

### R

repairing a file system, 28

### S

setup, initial

    initial tasks, 7

suspending activity on a file system, 28

system hang at unmount, 18

### T

tables

    GFS2-specific options for adding journals, 26

    GFS2-specific options for expanding file systems, 24

    mkfs.gfs2 command options, 13

    mount options, 16

### U

umount command, 17

unmount, system hang, 18

unmounting a file system, 17, 17

### W

withdraw function, GFS2, 34