

Red Hat Enterprise Linux 6

Global File System 2

Red Hat Global File System 2



Red Hat Enterprise Linux 6 Global File System 2

Red Hat Global File System 2

Edition 7

Copyright © 2011 Red Hat, Inc. and others.

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 6.

Introduction	v
1. Audience	v
2. Related Documentation	v
3. We Need Feedback!	v
4. Document Conventions	vi
4.1. Typographic Conventions	vi
4.2. Pull-quote Conventions	vii
4.3. Notes and Warnings	viii
1. GFS2 Overview	1
1.1. New and Changed Features	2
1.1.1. New and Changed Features for Red Hat Enterprise Linux 6.0	2
1.1.2. New and Changed Features for Red Hat Enterprise Linux 6.1	2
1.1.3. New and Changed Features for Red Hat Enterprise Linux 6.2	3
1.2. Before Setting Up GFS2	3
1.3. Differences between GFS and GFS2	4
1.3.1. GFS2 Command Names	4
1.3.2. Additional Differences Between GFS and GFS2	5
1.3.3. GFS2 Performance Improvements	6
1.4. GFS2 Node Locking	7
1.4.1. Performance Tuning With GFS2	8
1.4.2. Troubleshooting GFS2 Performance with the GFS2 Lock Dump	9
2. Getting Started	13
2.1. Prerequisite Tasks	13
2.2. Initial Setup Tasks	13
2.3. Deploying a GFS2 Cluster	14
3. Managing GFS2	15
3.1. Making a File System	15
3.2. Mounting a File System	18
3.3. Unmounting a File System	22
3.4. Special Considerations when Mounting GFS2 File Systems	22
3.5. GFS2 Quota Management	23
3.5.1. Configuring Disk Quotas	23
3.5.2. Managing Disk Quotas	26
3.5.3. Keeping Quotas Accurate	27
3.5.4. Synchronizing Quotas with the quotasync Command	27
3.5.5. References	28
3.6. Growing a File System	29
3.7. Adding Journals to a File System	30
3.8. Data Journaling	32
3.9. Configuring atime Updates	33
3.9.1. Mount with relatime	33
3.9.2. Mount with noatime	34
3.10. Suspending Activity on a File System	34
3.11. Repairing a File System	35
3.12. Bind Mounts and Context-Dependent Path Names	37
3.13. Bind Mounts and File System Mount Order	38
3.14. The GFS2 Withdraw Function	40
4. Diagnosing and Correcting Problems with GFS2 File Systems	43
4.1. GFS2 File System Shows Slow Performance	43
4.2. Setting Up NFS Over GFS2	43
4.3. GFS2 File System Hangs and Requires Reboot of One Node	44
4.4. GFS2 File System Hangs and Requires Reboot of All Nodes	44

4.5. GFS2 File System Does Not Mount on Newly-Added Cluster Node	45
4.6. Space Indicated as Used in Empty File System	45
A. GFS2 Quota Management with the gfs2_quota Command	47
A.1. Setting Quotas with the gfs2_quota command	47
A.2. Displaying Quota Limits and Usage with the gfs2_quota Command	48
A.3. Synchronizing Quotas with the gfs2_quota Command	50
A.4. Enabling/Disabling Quota Enforcement	51
A.5. Enabling Quota Accounting	51
B. Converting a File System from GFS to GFS2	53
C. GFS2 tracepoints and the debugfs glocks File	55
C.1. GFS2 tracepoint Types	55
C.2. Tracepoints	55
C.3. Glocks	56
C.4. The glock debugfs Interface	57
C.5. Glock Holders	60
C.6. Glock tracepoints	61
C.7. Bmap tracepoints	61
C.8. Log tracepoints	61
C.9. References	62
D. Revision History	63
Index	65

Introduction

This book provides information about configuring and maintaining Red Hat GFS2 (Red Hat Global File System 2), which is included in the Resilient Storage Add-On.

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:

- *Installation Guide* — Documents relevant information regarding the installation of Red Hat Enterprise Linux 6.
- *Deployment Guide* — Documents relevant information regarding the deployment, configuration and administration of Red Hat Enterprise Linux 6.
- *Storage Administration Guide* — Provides instructions on how to effectively manage storage devices and file systems on Red Hat Enterprise Linux 6.

For more information about the High Availability Add-On and the Resilient Storage Add-On for Red Hat Enterprise Linux 6, refer to the following resources:

- *High Availability Add-On Overview* — Provides a high-level overview of the Red Hat High Availability Add-On.
- *Cluster Administration* — Provides information about installing, configuring and managing the High Availability Add-On.
- *Logical Volume Manager Administration* — Provides a description of the Logical Volume Manager (LVM), including information on running LVM in a clustered environment.
- *DM Multipath* — Provides information about using the Device-Mapper Multipath feature of Red Hat Enterprise Linux.
- *Load Balancer Administration* — Provides information on configuring high-performance systems and services with the Load Balancer Add-On, a set of integrated software components that provide Linux Virtual Servers (LVS) for balancing IP load across a set of real servers.
- *Release Notes* — Provides information about the current release of Red Hat products.

High Availability Add-On 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://docs.redhat.com/docs/en-US/index.html>.

3. We Need Feedback!

If you find a typographical error in this manual, 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/> against the product **Red Hat Enterprise Linux 6** and the component **doc-Global_File_System_2**. When submitting a bug report, be sure to mention the manual's identifier:

```
rh-gfs2(EN)-6 (2011-11-14T15:15)
```

If you have a suggestion for improving the documentation, try to be as specific as possible when describing it. 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*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.

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:

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:

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

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.

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:

Introduction

```
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 included in the Resilient Storage Add-On. It 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 the High Availability Add-On.



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 6 release 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 mounting snapshots of cluster file systems (for example, for backup purposes).



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 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 cluster, Red Hat GFS2 nodes can be configured and managed with High Availability Add-On configuration and management tools. Red Hat GFS2 then provides data sharing among GFS2 nodes in a 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 the High Availability Add-On 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 included in the Resilient Storage Add-On. It is a cluster-wide implementation of LVM, enabled by the CLVM daemon **clvmd**, which manages LVM logical volumes in a 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, “New and Changed Features”](#)
- [Section 1.2, “Before Setting Up GFS2”](#)
- [Section 1.3, “Differences between GFS and GFS2”](#)
- [Section 1.4, “GFS2 Node Locking”](#)

1.1. New and Changed Features

This section lists new and changed features of the GFS2 file system and the GFS2 documentation that are included with the initial and subsequent releases of Red Hat Enterprise Linux 6.

1.1.1. New and Changed Features for Red Hat Enterprise Linux 6.0

Red Hat Enterprise Linux 6.0 includes the following documentation and feature updates and changes.

- For the Red Hat Enterprise Linux 6 release, Red Hat does not support the use of GFS2 as a single-node file system.
- For the Red Hat Enterprise Linux 6 release, the **gfs2_convert** command to upgrade from a GFS to a GFS2 file system has been enhanced. For information on this command, see [Appendix B, *Converting a File System from GFS to GFS2*](#).
- The Red Hat Enterprise Linux 6 release supports the **discard**, **nodiscard**, **barrier**, **nobarrier**, **quota_quantum**, **statfs_quantum**, and **statfs_percent** mount options. For information about mounting a GFS2 file system, see [Section 3.2, “Mounting a File System”](#).
- The Red Hat Enterprise Linux 6 version of this document contains a new section, [Section 1.4, “GFS2 Node Locking”](#). This section describes some of the internals of GFS2 file systems.

1.1.2. New and Changed Features for Red Hat Enterprise Linux 6.1

Red Hat Enterprise Linux 6.1 includes the following documentation and feature updates and changes.

- As of the Red Hat Enterprise Linux 6.1 release, GFS2 supports the standard Linux quota facilities. GFS2 quota management is documented in [Section 3.5, “GFS2 Quota Management”](#).

For earlier releases of Red Hat Enterprise Linux, GFS2 required the **gfs2_quota** command to manage quotas. Documentation for the **gfs2_quota** is now provided in [Appendix A, *GFS2 Quota Management with the gfs2_quota Command*](#).

- This document now contains a new chapter, [Chapter 4, Diagnosing and Correcting Problems with GFS2 File Systems](#).
- Small technical corrections and clarifications have been made throughout the document.

1.1.3. New and Changed Features for Red Hat Enterprise Linux 6.2

Red Hat Enterprise Linux 6.2 includes the following documentation and feature updates and changes.

- As of the Red Hat Enterprise Linux 6.2 release, GFS2 supports the **tunegfs2** command, which replaces some of the features of the **gfs2_tool** command. For further information, refer to the **tunegfs2** man page.

The following sections have been updated to provide administrative procedures that do not require the use of the **gfs2_tool** command:

- [Section 3.5.4, “Synchronizing Quotas with the **quotasync** Command”](#). and [Section A.3, “Synchronizing Quotas with the **gfs2_quota** Command”](#) now describe how to change the **quota_quantum** parameter from its default value of 60 seconds by using the **quota_quantum=** mount option.
- [Section 3.10, “Suspending Activity on a File System”](#) now describes how to suspend write activity to a file system using the **dmsetup suspend** command.
- This document includes a new appendix, [Appendix C, GFS2 tracepoints and the **debugfs** glocks File](#). This appendix describes the **glock debugfs** interface and the GFS2 tracepoints. It is intended for advanced users who are familiar with file system internals who would like to learn more about the design of GFS2 and how to debug GFS2-specific issues.

1.2. 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 cluster 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. 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”](#).

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.

1.3. 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 B, Converting a File System from GFS to GFS2](#).

1.3.1. 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 and functionality.

Table 1.1. GFS and GFS2 Commands

GFS Command	GFS2 Command	Description
mount	mount	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.
umount	umount	Unmount a file system.
fsck gfs_fsck	fsck fsck.gfs2	Check and repair an unmounted file system.
gfs_grow	gfs2_grow	Grow a mounted file system.
gfs_jadd	gfs2_jadd	Add a journal to a mounted file system.
gfs_mkfs mkfs -t gfs	mkfs.gfs2 mkfs -t gfs2	Create a file system on a storage device.
gfs_quota	gfs2_quota	Manage quotas on a mounted file system. As of the Red Hat Enterprise Linux 6.1 release, GFS2 supports the standard Linux quota facilities. For further information on quota management in GFS2, refer to Section 3.5, “GFS2 Quota Management” .
gfs_tool	tunegfs2 mount parameters dmsetup suspend	Configure, tune, or gather information about a file system. The tunegfs2 command is supported as of the Red Hat Enterprise Linux 6.2 release. There is also a gfs2_tool command.
gfs_edit	gfs2_edit	Display, print, or edit file system internal structures. The gfs2_edit command can be used for GFS file systems as well as GFS2 file system.

GFS Command	GFS2 Command	Description
gfs_tool setflag jdata/ inherit_jdata	chattr +j (preferred)	Enable journaling on a file or directory.
setfacl/ getfacl	setfacl/ getfacl	Set or get file access control list for a file or directory.
setfattr/ getfattr	setfattr/ getfattr	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.

1.3.2. Additional Differences Between GFS and GFS2

This section summarizes the additional differences in GFS and GFS2 administration that are not described in [Section 1.3.1, "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 bind mounts and context-dependent 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. For information on enabling and disabling quota enforcement, see [Section 3.5, "GFS2 Quota Management"](#).

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.



Note

As of the Red Hat Enterprise Linux 6.2 release, GFS2 supports the **tunegfs2** command, which replaces some of the features of the **gfs2_tool** command. For further information, refer to the **tunegfs2(8)** man page. The **set tune** and **get tune** functions of the **gfs2_tool** command have been replaced by command line options of the **mount** command, which allows them to be set by means of the **fstab** file when required.

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.

1.3.3. 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.

1.4. GFS2 Node Locking

In order to get the best performance from a GFS2 file system, it is very important to understand some of the basic theory of its operation. A single node file system is implemented alongside a cache, the purpose of which is to eliminate latency of disk accesses when using frequently requested data. In Linux the page cache (and historically the buffer cache) provide this caching function.

With GFS2, each node has its own page cache which may contain some portion of the on-disk data. GFS2 uses a locking mechanism called *glocks* (pronounced gee-locks) to maintain the integrity of the cache between nodes. The glock subsystem provides a cache management function which is implemented using the *distributed lock manager* (DLM) as the underlying communication layer.

The glocks provide protection for the cache on a per-inode basis, so there is one lock per inode which is used for controlling the caching layer. If that glock is granted in shared mode (DLM lock mode: PR) then the data under that glock may be cached upon one or more nodes at the same time, so that all the nodes may have local access to the data.

If the glock is granted in exclusive mode (DLM lock mode: EX) then only a single node may cache the data under that glock. This mode is used by all operations which modify the data (such as the **write** system call).

If another node requests a glock which cannot be granted immediately, then the DLM sends a message to the node or nodes which currently hold the glocks blocking the new request to ask them to drop their locks. Dropping glocks can be (by the standards of most file system operations) a long process. Dropping a shared glock requires only that the cache be invalidated, which is relatively quick and proportional to the amount of cached data.

Dropping an exclusive glock requires a log flush, and writing back any changed data to disk, followed by the invalidation as per the shared glock.

The difference between a single node file system and GFS2 then, is that a single node file system has a single cache and GFS2 has a separate cache on each node. In both cases, latency to access to cached data is of a similar order of magnitude, but the latency to access uncached data is much greater in GFS2 if another node has previously cached that same data.



Note

Due to the way in which GFS2's caching is implemented the best performance is obtained when either of the following takes place:

- An inode is used in a read only fashion across all nodes.
- An inode is written or modified from a single node only.

Note that inserting and removing entries from a directory during file creation and deletion counts as writing to the directory inode.

It is possible to break this rule provided that it is broken relatively infrequently. Ignoring this rule too often will result in a severe performance penalty.

If you **mmap()** a file on GFS2 with a read/write mapping, but only read from it, this only counts as a read. On GFS though, it counts as a write, so GFS2 is much more scalable with **mmap()** I/O.

If you do not set the **noatime mount** parameter, then reads will also result in writes to update the file timestamps. We recommend that all GFS2 users should mount with **noatime** unless they have a specific requirement for **atime**.

1.4.1. Performance Tuning With GFS2

It is usually possible to alter the way in which a troublesome application stores its data in order to gain a considerable performance advantage.

A typical example of a troublesome application is an email server. These are often laid out with a spool directory containing files for each user (**mbox**), or with a directory for each user containing a file for each message (**maildir**). When requests arrive over IMAP, the ideal arrangement is to give each

user an affinity to a particular node. That way their requests to view and delete email messages will tend to be served from the cache on that one node. Obviously if that node fails, then the session can be restarted on a different node.

When mail arrives via SMTP, then again the individual nodes can be set up so as to pass a certain user's mail to a particular node by default. If the default node is not up, then the message can be saved directly into the user's mail spool by the receiving node. Again this design is intended to keep particular sets of files cached on just one node in the normal case, but to allow direct access in the case of node failure.

This setup allows the best use of GFS2's page cache and also makes failures transparent to the application, whether **imap** or **smtp**.

Backup is often another tricky area. Again, if it is possible it is greatly preferable to back up the working set of each node directly from the node which is caching that particular set of inodes. If you have a backup script which runs at a regular point in time, and that seems to coincide with a spike in the response time of an application running on GFS2, then there is a good chance that the cluster may not be making the most efficient use of the page cache.

Obviously, if you are in the (enviable) position of being able to stop the application in order to perform a backup, then this won't be a problem. On the other hand, if a backup is run from just one node, then after it has completed a large portion of the file system will be cached on that node, with a performance penalty for subsequent accesses from other nodes. This can be mitigated to a certain extent by dropping the VFS page cache on the backup node after the backup has completed with following command:

```
echo -n 3 >/proc/sys/vm/drop_caches
```

However this is not as good a solution as taking care to ensure the working set on each node is either shared, mostly read only across the cluster, or accessed largely from a single node.

1.4.2. Troubleshooting GFS2 Performance with the GFS2 Lock Dump

If your cluster performance is suffering because of inefficient use of GFS2 caching, you may see large and increasing I/O wait times. You can make use of GFS2's lock dump information to determine the cause of the problem.

This section provides an overview of the GFS2 lock dump. For a more complete description of the GFS2 lock dump, see [Appendix C, GFS2 tracepoints and the debugfs glocks File](#).

The GFS2 lock dump information can be gathered from the **debugfs** file which can be found at the following path name, assuming that **debugfs** is mounted on **/sys/kernel/debug/**:

```
/sys/kernel/debug/gfs2/fsname/glocks
```

The content of the file is a series of lines. Each line starting with G: represents one glock, and the following lines, indented by a single space, represent an item of information relating to the glock immediately before them in the file.

The best way to use the **debugfs** file is to use the **cat** command to take a copy of the complete content of the file (it might take a long time if you have a large amount of RAM and a lot of cached

inodes) while the application is experiencing problems, and then looking through the resulting data at a later date.



Tip

It can be useful to make two copies of the **debugfs** file, one a few seconds or even a minute or two after the other. By comparing the holder information in the two traces relating to the same glock number, you can tell whether the workload is making progress (that is, it is just slow) or whether it has become stuck (which is always a bug and should be reported to Red Hat support immediately).

Lines in the **debugfs** file starting with H: (holders) represent lock requests either granted or waiting to be granted. The flags field on the holders line f: shows which: The 'W' flag refers to a waiting request, the 'H' flag refers to a granted request. The glocks which have large numbers of waiting requests are likely to be those which are experiencing particular contention.

[Table 1.2, “Glock flags”](#) shows the meanings of the different glock flags and [Table 1.3, “Glock holder flags”](#) shows the meanings of the different glock holder flags.

Table 1.2. Glock flags

Flag	Name	Meaning
d	Pending demote	A deferred (remote) demote request
D	Demote	A demote request (local or remote)
f	Log flush	The log needs to be committed before releasing this glock
F	Frozen	Replies from remote nodes ignored - recovery is in progress
i	Invalidate in progress	In the process of invalidating pages under this glock
l	Initial	Set when DLM lock is associated with this glock
l	Locked	The glock is in the process of changing state
p	Demote in progress	The glock is in the process of responding to a demote request
r	Reply pending	Reply received from remote node is awaiting processing
y	Dirty	Data needs flushing to disk before releasing this glock

Table 1.3. Glock holder flags

Flag	Name	Meaning
a	Async	Do not wait for glock result (will poll for result later)
A	Any	Any compatible lock mode is acceptable
c	No cache	When unlocked, demote DLM lock immediately
e	No expire	Ignore subsequent lock cancel requests
E	exact	Must have exact lock mode
F	First	Set when holder is the first to be granted for this lock
H	Holder	Indicates that requested lock is granted

Flag	Name	Meaning
p	Priority	Enqueue holder at the head of the queue
t	Try	A "try" lock
T	Try 1CB	A "try" lock that sends a callback
W	Wait	Set while waiting for request to complete

Having identified a glock which is causing a problem, the next step is to find out which inode it relates to. The glock number (n: on the G: line) indicates this. It is of the form *typenumber* and if *type* is 2, then the glock is an inode glock and the *number* is an inode number. To track down the inode, you can then run **find -inum number** where *number* is the inode number converted from the hex format in the glocks file into decimal.



Warning

If you run the **find** on a file system when it is experiencing lock contention, you are likely to make the problem worse. It is a good idea to stop the application before running the **find** when you are looking for contended inodes.

Table 1.4, “Glock types” shows the meanings of the different glock types.

Table 1.4. Glock types

Type number	Lock type	Use
1	Trans	Transaction lock
2	Inode	Inode metadata and data
3	Rgrp	Resource group metadata
4	Meta	The superblock
5	lopen	Inode last closer detection
6	Flock	flock(2) syscall
8	Quota	Quota operations
9	Journal	Journal mutex

If the glock that was identified was of a different type, then it is most likely to be of type 3: (resource group). If you see significant numbers of processes waiting for other types of glock under normal loads, then please report this to Red Hat support.

If you do see a number of waiting requests queued on a resource group lock there may be a number of reason for this. One is that there are a large number of nodes compared to the number of resource groups in the file system. Another is that the file system may be very nearly full (requiring, on average, longer searches for free blocks). The situation in both cases can be improved by adding more storage and using the **gfs2_grow** command to expand the file system.

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”](#)
- [Section 2.3, “Deploying a GFS2 Cluster”](#)

2.1. Prerequisite Tasks

Before setting up Red Hat GFS2, make sure that you have noted the key characteristics of the GFS2 nodes (refer to [Section 1.2, “Before Setting Up GFS2”](#)). Also, 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.

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
```

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 the Red Hat High Availability Add-On to automate mounting and unmounting GFS2 file systems.

2.3. Deploying a GFS2 Cluster

Deploying a cluster file system is not a “drop in” replacement for a single node deployment. We recommend that you allow a period of around 8-12 weeks of testing on new installations in order to test the system and ensure that it is working at the required performance level. During this period any performance or functional issues can be worked out and any queries should be directed to the Red Hat support team. We also recommend that customers considering deploying clusters have their configurations reviewed by Red Hat support before deployment to avoid any possible support issues later on.

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 6 release, Red Hat does not support the use of GFS2 as a single-node file system.

```
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 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 example, **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
-c	<i>Megabytes</i>	Sets the initial size of each journal's quota change file to <i>Megabytes</i> .
-D		Enables debugging output.
-h		Help. Displays available options.
-J	<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.
-j	<i>Number</i>	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. 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.
-O		Prevents the mkfs.gfs2 command from asking for confirmation before writing the file system.

Flag	Parameter	Description
-p	<i>LockProtoName</i>	Specifies the name of the locking protocol to use. Recognized locking protocols include: lock_dlm — The standard locking module, required for a clustered file system. lock_no_lock — Used when GFS2 is acting as a local file system (one node only).
-q		Quiet. Do not display anything.
-r	<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.
-t	<i>LockTableName</i>	A unique identifier that specifies the lock table field when you use the lock_dlm protocol; the lock_no_lock 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 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 Cluster Configuration Tool and displayed at the Cluster Status Tool 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.
-u	<i>MegaBytes</i>	Specifies the initial size of each journal's unlinked tag file.
-v		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.



Note

Attempting to mount a GFS2 file system when the Cluster Manager (**cman**) has not been started produces the following error message:

```
[root@gfs-a24c-01 ~]# mount -t gfs2 -o noatime /dev/mapper/mpathap1 /mnt
gfs_controld join connect error: Connection refused
error mounting lockproto lock_dlm
```

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/lvol10** is mounted on the **/mygfs2** directory.

```
mount /dev/vg01/lvol10 /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. Note, however, that for the Red Hat Enterprise Linux 6 release, 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 mounting snapshots of cluster file systems (for example, for backup purposes).

Table 3.2. GFS2-Specific Mount Options

Option	Description
acl	Allows manipulating file ACLs. If a file system is mounted without the acl mount option, users are allowed to view ACLs (with getfacl), but are not allowed to set them (with setfacl).
data=[ordered writeback]	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 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 ordered mode, but it should be slightly faster for some workloads. The default value is ordered mode.
ignore_local_fs Caution: 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 lock_nolock automatically turns on the localflocks flag.
localflocks Caution: 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 localflocks flag is automatically turned on by lock_nolock .
lockproto=LockModuleName	Allows the user to specify which locking protocol to use with the file system. If <i>LockModuleName</i> is not

Option	Description
	specified, the locking protocol name is read from the file system superblock.
locktable=LockTableName	Allows the user to specify which locking table to use with the file system.
quota=[off/account/on]	Turns quotas on or off for a file system. Setting the quotas to be in the account 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 off .
errors=panic withdraw	When errors=panic is specified, file system errors will cause a kernel panic. The default behavior, which is the same as specifying errors=withdraw , 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 Section 3.14, "The GFS2 Withdraw Function" .
discard/nodiscard	Causes GFS2 to generate "discard" I/O requests for blocks that have been freed. These can be used by suitable hardware to implement thin provisioning and similar schemes.
barrier/nobarrier	Causes GFS2 to send I/O barriers when flushing the journal. The default value is on . This option is automatically turned off if the underlying device does not support I/O barriers. Use of I/O barriers with GFS2 is highly recommended at all times unless the block device is designed so that it cannot lose its write cache content (for example, if it is on a UPS or it does not have a write cache).
quota_quantum=secs	Sets the number of seconds for which a change in the quota information may sit on one node before being written to the quota file. This is the preferred way to set this parameter. The value is an integer number of seconds greater than zero. The default is 60 seconds. Shorter settings result in faster updates of the lazy quota information and less likelihood of someone exceeding their quota. Longer settings make file system operations involving quotas faster and more efficient.
statfs_quantum=secs	Setting statfs_quantum to 0 is the preferred way to set the slow version of statfs . The default value is 30 secs which sets the maximum time period before statfs changes will be synced to the master statfs file. This can be adjusted to allow for faster, less accurate statfs values or slower more accurate values. When this option is set to 0, statfs will always report the true values.
statfs_percent=value	Provides a bound on the maximum percentage change in the statfs information on a local basis before

Option	Description
	it is synced back to the master <code>statfs</code> file, even if the time period has not expired. If the setting of <code>statfs_quantum</code> is 0, then this setting is ignored.

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. When a GFS2 file system is mounted with the **quota=on** or **quota=account** option, 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.



Note

As of the Red Hat Enterprise Linux 6.1 release, GFS2 supports the standard Linux quota facilities. In order to use this you will need to install the **quota** RPM. This is the preferred way to administer quotas on GFS2 and should be used for all new deployments of GFS2 using quotas. This section documents GFS2 quota management using these facilities.

For earlier releases of Red Hat Enterprise Linux, GFS2 required the **gfs2_quota** command to manage quotas. For information on using the **gfs2_quota** command, see [Appendix A, GFS2 Quota Management with the gfs2_quota Command](#).

3.5.1. Configuring Disk Quotas

To implement disk quotas, use the following steps:

1. Set up quotas in enforcement or accounting mode.
2. Initialize the quota database file with current block usage information.
3. Assign quota policies. (In accounting mode, these policies are not enforced.)

Each of these steps is discussed in detail in the following sections.

3.5.1.1. Setting Up Quotas in Enforcement or Accounting Mode

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

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

To mount a file system with quotas enabled, mount the file system with the **quota=on** option specified.



Chapter 3. Managing GFS2

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

To mount a file system with quota accounting maintained, even though the quota limits are not enforced, mount the file system with the **quota=account** option specified.

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

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

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

quota={on|off|account}

on - Specifies that quotas are enabled when the file system is mounted.

off - Specifies that quotas are disabled when the file system is mounted.

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.

Examples

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

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

In this example, the GFS2 file system on **/dev/vg01/lvol0** is mounted on the **/mygfs2** directory with quota accounting maintained, but not enforced.

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

3.5.1.2. Creating the Quota Database Files

After each quota-enabled file system is mounted, the system is capable of working with disk quotas. However, the file system itself is not yet ready to support quotas. The next step is to run the **quotacheck** command.

The **quotacheck** command examines quota-enabled file systems and builds a table of the current disk usage per file system. The table is then used to update the operating system's copy of disk usage. In addition, the file system's disk quota files are updated.

To create the quota files on the file system, use the **-u** and the **-g** options of the **quotacheck** command; both of these options must be specified for user and group quotas to be initialized. For example, if quotas are enabled for the **/home** file system, create the files in the **/home** directory:

```
quotacheck -ug /home
```

3.5.1.3. Assigning Quotas per User

The last step is assigning the disk quotas with the **edquota** command. Note that if you have mounted your file system in accounting mode (with the **quota=account** option specified), the quotas are not enforced.

To configure the quota for a user, as root in a shell prompt, execute the command:

```
edquota username
```

Perform this step for each user who needs a quota. For example, if a quota is enabled in **/etc/fstab** for the **/home** partition (**/dev/Vo1Group00/LogVo102** in the example below) and the command **edquota testuser** is executed, the following is shown in the editor configured as the default for the system:

```
Disk quotas for user testuser (uid 501):
Filesystem      blocks    soft    hard    inodes    soft    hard
/dev/Vo1Group00/LogVo102 440436      0      0
```



Note

The text editor defined by the EDITOR environment variable is used by **edquota**. To change the editor, set the EDITOR environment variable in your **~/ .bash_profile** file to the full path of the editor of your choice.

The first column is the name of the file system that has a quota enabled for it. The second column shows how many blocks the user is currently using. The next two columns are used to set soft and hard block limits for the user on the file system.

The soft block limit defines the maximum amount of disk space that can be used.

The hard block limit is the absolute maximum amount of disk space that a user or group can use. Once this limit is reached, no further disk space can be used.

The GFS2 file system does not maintain quotas for inodes, so these columns do not apply to GFS2 file systems and will be blank.

If any of the values are set to 0, that limit is not set. In the text editor, change the desired limits. For example:

```
Disk quotas for user testuser (uid 501):
```

```
Filesystem      blocks      soft      hard      inodes      soft      hard
/dev/VolGroup00/LogVol102 440436    500000    550000
```

To verify that the quota for the user has been set, use the command:

```
quota testuser
```

3.5.1.4. Assigning Quotas per Group

Quotas can also be assigned on a per-group basis. Note that if you have mounted your file system in accounting mode (with the **account=on** option specified), the quotas are not enforced.

To set a group quota for the **devel** group (the group must exist prior to setting the group quota), use the following command:

```
edquota -g devel
```

This command displays the existing quota for the group in the text editor:

```
Disk quotas for group devel (gid 505):
Filesystem      blocks      soft      hard      inodes      soft      hard
/dev/VolGroup00/LogVol102 440400      0          0
```

The GFS2 file system does not maintain quotas for inodes, so these columns do not apply to GFS2 file systems and will be blank. Modify the limits, then save the file.

To verify that the group quota has been set, use the following command:

```
quota -g devel
```

3.5.2. Managing Disk Quotas

If quotas are implemented, they need some maintenance — mostly in the form of watching to see if the quotas are exceeded and making sure the quotas are accurate.

Of course, if users repeatedly exceed their quotas or consistently reach their soft limits, a system administrator has a few choices to make depending on what type of users they are and how much disk space impacts their work. The administrator can either help the user determine how to use less disk space or increase the user's disk quota.

You can create a disk usage report by running the **repquota** utility. For example, the command **repquota /home** produces this output:

```
*** Report for user quotas on device /dev/mapper/VolGroup00-LogVol102
Block grace time: 7days; Inode grace time: 7days
  Block limits  File limits
User  used soft hard grace used soft hard grace
-----
```

root	--	36	0	0	4	0	0
kristin	--	540	0	0	125	0	0
testuser	--	440400	500000	550000	37418	0	0

To view the disk usage report for all (option **-a**) quota-enabled file systems, use the command:

```
repquota -a
```

While the report is easy to read, a few points should be explained. The `--` displayed after each user is a quick way to determine whether the block limits have been exceeded. If the block soft limit is exceeded, a `+` appears in place of the the first `-` in the output. The second `-` indicates the inode limit, but GFS2 file systems do not support inode limits so that character will remain as `-`. GFS2 file systems do not support a grace period, so the **grace** column will remain blank.

Note that the **repquota** command is not supported over NFS, irrespective of the underlying file system.

3.5.3. Keeping Quotas Accurate

If you enable quotas on your file system after a period of time when you have been running with quotas disabled, you should run the **quotacheck** command to create, check, and repair quota files. Additionally, you may want to run the **quotacheck** if you think your quota files may not be accurate, as may occur when a file system is not unmounted cleanly after a system crash.

For more information about the **quotacheck** command, see the **quotacheck** man page.



Note

Run **quotacheck** when the file system is relatively idle on all nodes because disk activity may affect the computed quota values.

3.5.4. Synchronizing Quotas with the **quotasync** Command

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, by default 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**. You can change this from its default value of 60 seconds using the **quota_quantum=** mount option, as described in [Table 3.2, “GFS2-Specific Mount Options”](#). 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 update the **quota_quantum** value with the **mount -o remount**.

You can use the **quotasync** 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

```
quotasync [-ug] -a|mntpnt...
```

u

Sync the user quota files.

g

Sync the group quota files

a

Sync all file systems that are currently quota-enabled and support sync. When *-a* is absent, a file system mountpoint should be specified.

mntpnt

Specifies the GFS2 file system to which the actions apply.

Tuning the Time Between Synchronizations

```
mount -o quota_quantum=secs,remount BlockDevice MountPoint
```

MountPoint

Specifies the GFS2 file system to which the actions apply.

secs

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 all the cached dirty quotas from the node it is run on to the ondisk quota file for the file system **/mnt/mygfs2**.

```
# quotasync -ug /mnt/mygfs2
```

This example changes the default time period between regular quota-file updates to one hour (3600 seconds) for file system **/mnt/mygfs2** when remounting that file system on logical volume **/dev/volgroup/logical_volume**.

```
# mount -o quota_quantum=3600,remount /dev/volgroup/logical_volume /mnt/mygfs2
```

3.5.5. References

For more information on disk quotas, refer to the **man** pages of the following commands:

- **quotacheck**
- **edquota**

- `repquota`
- `quota`

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

Flag	Parameter	Description
-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 gfs2_jadd 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 gfs2_jadd 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
----- /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.

Two methods of reducing the effects of **atime** updating are available:

- Mount with **relatime** (relative atime), which updates the **atime** if the previous **atime** update is older than the **mtime** or **ctime** update.
- Mount with **noatime**, which disables **atime** updates on that file system.

3.9.1. Mount with **relatime**

The **relatime** (relative atime) Linux mount option can be specified when the file system is mounted. This specifies that the **atime** is updated if the previous **atime** update is older than the **mtime** or **ctime** update.

Usage

```
mount BlockDevice MountPoint -o relatime
```

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`. The `atime` updates take place only if the previous `atime` update is older than the `mtime` or `ctime` update.

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

3.9.2. Mount with `noatime`

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 `dmsetup suspend` command. Suspending write activity allows hardware-based device snapshots to be used to capture the file system in a consistent state. The `dmsetup resume` command ends the suspension.

Usage

Start Suspension

```
dmsetup suspend MountPoint
```

End Suspension

```
dmsetup resume MountPoint
```

MountPoint

Specifies the file system.

Examples

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

```
# dmsetup suspend /mygfs2
```

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

```
# dmsetup resume /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.

Running the `fsck.gfs2` command requires system memory above and beyond the memory used for the operating system and kernel. Each block of memory in the GFS2 file system itself requires approximately five bits of additional memory, or 5/8 of a byte. So to estimate how many bytes of 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
fsck.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 /root/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@menscryfa ~]# 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_controld** daemon requesting withdraw. The **gfs_controld** daemon runs the **dmsetup** program to place the device mapper error target underneath the file system 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 file system 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.

Diagnosing and Correcting Problems with GFS2 File Systems

This chapter provides information about some common GFS2 issues and how to address them.

4.1. GFS2 File System Shows Slow Performance

You may find that your GFS2 file system shows slower performance than an ext3 file system. GFS2 performance may be affected by a number of influences and in certain use cases. Information that addresses GFS2 performance issues is found throughout this document.

4.2. Setting Up NFS Over GFS2

Due to the added complexity of the GFS2 locking subsystem and its clustered nature, setting up NFS over GFS2 requires taking many precautions and careful configuration. This section describes the caveats you should take into account when configuring an NFS service over a GFS2 file system.



Warning

If the GFS2 file system is NFS exported, and NFS client applications use POSIX locks, then you must mount the file system with the **locallocks** option. The intended effect of this is to force POSIX locks from each server to be local: i.e., non-clustered, independent of each other. (A number of problems exist if GFS2 attempts to implement POSIX locks from NFS across the nodes of a cluster.) For applications running on NFS clients, localized POSIX locks means that two clients can hold the same lock concurrently if the two clients are mounting from different servers. If all clients mount NFS from one server, then the problem of separate servers granting the same locks independently goes away.

In addition to the locking considerations, you should take the following into account when configuring an NFS service over a GFS2 file system.

- Red Hat supports only Red Hat High Availability Add-On configurations using NFSv3 with locking in an active/passive configuration with the following characteristics:
 - The backend file system is a GFS2 file system running on a 2 to 16 node cluster.
 - An NFSv3 server is defined as a service exporting the entire GFS2 file system from a single cluster node at a time.
 - The NFS server can fail over from one cluster node to another (active/passive configuration).
 - No access to the GFS2 file system is allowed *except* through the NFS server. This includes both local GFS2 file system access as well as access through Samba or Clustered Samba.
 - There is no NFS quota support on the system.

This configuration provides HA for the file system and reduces system downtime since a failed node does not result in the requirement to execute the **fsck** command when failing the NFS server from one node to another.

- The **fsid=** NFS option is mandatory for NFS exports of GFS2.

- If problems arise with your cluster (for example, the cluster becomes inquorate and fencing is not successful), the clustered logical volumes and the GFS2 file system will be frozen and no access is possible until the cluster is quorate. You should consider this possibility when determining whether a simple failover solution such as the one defined in this procedure is the most appropriate for your system.

4.3. GFS2 File System Hangs and Requires Reboot of One Node

If your GFS2 file system hangs and does not return commands run against it, but rebooting one specific node returns the system to normal, this may be indicative of a locking problem or bug. Should this occur, gather the following data:

- The gfs2 lock dump for the file system on each node:

```
cat /sys/kernel/debug/gfs2/fsname/glocks >glocks.fsname.nodename
```

- The DLM lock dump for the file system on each node: You can get this information with the **dlm_tool**:

```
dlm_tool lockdebug -sv lname.
```

In this command, *lname* is the lockspace name used by DLM for the file system in question. You can find this value in the output from the **group_tool** command.

- The output from the **sysrq -t** command.
- The contents of the **/var/log/messages** file.

Once you have gathered that data, you can open a ticket with Red Hat Support and provide the data you have collected.

4.4. GFS2 File System Hangs and Requires Reboot of All Nodes

If your GFS2 file system hangs and does not return commands run against it, requiring that you reboot all nodes in the cluster before using it, check for the following issues.

- You may have had a failed fence. GFS2 file systems will freeze to ensure data integrity in the event of a failed fence. Check the messages logs to see if there are any failed fences at the time of the hang. Ensure that fencing is configured correctly.
- The GFS2 file system may have withdrawn. Check through the messages logs for the word **withdraw** and check for any messages and calltraces from GFS2 indicating that the file system has been withdrawn. A withdraw is indicative of file system corruption, a storage failure, or a bug. Unmount the file system, update the **gfs2-utils** package, and execute the **fsck** command on the file system to return it to service. Open a support ticket with Red Hat Support. Inform them you experienced a GFS2 withdraw and provide sosreports with logs.

For information on the GFS2 withdraw function, see [Section 3.14, “The GFS2 Withdraw Function”](#).

- This error may be indicative of a locking problem or bug. Gather data during one of these occurrences and open a support ticket with Red Hat Support, as described in [Section 4.3, “GFS2 File System Hangs and Requires Reboot of One Node”](#).

4.5. GFS2 File System Does Not Mount on Newly-Added Cluster Node

If you add a new node to a cluster and you find that you cannot mount your GFS2 file system on that node, you may have fewer journals on the GFS2 file system than you have nodes attempting to access the GFS2 file system. You must have one journal per GFS2 host you intend to mount the file system on (with the exception of GFS2 file systems mounted with the **spectator** mount option set, since these do not require a journal). You can add journals to a GFS2 file system with the **gfs2_jadd** command, as described in [Section 3.7, “Adding Journals to a File System”](#).

4.6. Space Indicated as Used in Empty File System

If you have an empty GFS2 file system, the **df** command will show that there is space being taken up. This is because GFS2 file system journals consume space (number of journals * journal size) on disk. If you created a GFS2 file system with a large number of journals or specified a large journal size then you will see (number of journals * journal size) as already in use when you execute the **df**. Even if you did not specify a large number of journals or large journals, small GFS2 file systems (in the 1GB or less range) will show a large amount of space as being in use with the default GFS2 journal size.

Appendix A. GFS2 Quota Management with the `gfs2_quota` Command

As of the Red Hat Enterprise Linux 6.1 release, GFS2 supports the standard Linux quota facilities. In order to use this you will need to install the **quota** RPM. This is the preferred way to administer quotas on GFS2 and should be used for all new deployments of GFS2 using quotas. For information on using the standard Linux quota facilities, see [Section 3.5, “GFS2 Quota Management”](#).

For earlier releases of Red Hat Enterprise Linux, GFS2 required the **gfs2_quota** command to manage quotas. This appendix documents the use of the **gfs2_quota** command for managing GFS2 file system quotas.

A.1. Setting Quotas with the `gfs2_quota` command

Two quota settings are available for each user ID (UID) or group ID (GID): a *hard limit* and a *soft 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 soft limit is usually a value less than the hard limit. The file system will notify the user or group when the soft limit is reached to warn them of the amount of space they are using. A soft limit value of zero means that no limit is enforced.

You can set limits 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 A.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 soft limit for group ID 21 to 50 kilobytes on file system `/mygfs2`.

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

A.2. Displaying Quota Limits and Usage with the `gfs2_quota` Command

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, soft 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

```
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
```

A.3. Synchronizing Quotas with the `gfs2_quota` Command

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, by default 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`. You can change this from its default value of 60 seconds using the `quota_quantum=` mount option, as described in [Table 3.2, “GFS2-Specific Mount Options”](#). 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 update the `quota_quantum` value with the `mount -o remount`.

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

```
mount -o quota_quantum=secs,remount BlockDevice MountPoint
```

MountPoint

Specifies the GFS2 file system to which the actions apply.

secs

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 `/mnt/mygfs2` when remounting that file system on logical volume `/dev/volgroup/logical_volume`.

```
# mount -o quota_quantum=3600,remount /dev/volgroup/logical_volume /mnt/mygfs2
```

A.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/lvol10** is mounted on the **/mygfs2** directory with quota enforcement enabled.

```
# mount -o quota=on /dev/vg01/lvol10 /mygfs2
```

A.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.

Appendix A. GFS2 Quota Management with the `gfs2_quota` Command

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 with quota accounting enabled.

```
# mount -o quota=account /dev/vg01/lvol0 /mygfs2
```

Appendix B. Converting a File System from GFS to GFS2

Since the Red Hat Enterprise Linux 6 release does not support GFS file systems, you must upgrade any existing GFS file systems to GFS2 file systems with the **gfs2_convert** command. Note that you must perform this conversion procedure on a Red Hat Enterprise Linux 5 system before upgrading to Red Hat Enterprise Linux 6.



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”](#).

Appendix B. Converting a File System from GFS to GFS2

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.

Use the following procedure to convert a GFS file system to a GFS2 file system.

1. On a Red Hat Enterprise Linux system, 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.
5. Upgrade to Red Hat Enterprise Linux 6.

The following example converts a GFS file system on block device **/dev/shell_vg/500g** to a GFS2 file system.

```
[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 C. GFS2 tracepoints and the debugfs glocks File

This appendix describes both the glock **debugfs** interface and the GFS2 tracepoints. It is intended for advanced users who are familiar with file system internals who would like to learn more about the design of GFS2 and how to debug GFS2-specific issues.

C.1. GFS2 tracepoint Types

There are currently three types of GFS2 tracepoints: *glock* (pronounced "gee-lock") tracepoints, *bmap* tracepoints and *log* tracepoints. These can be used to monitor a running GFS2 file system and give additional information to that which can be obtained with the debugging options supported in previous releases of Red Hat Enterprise Linux. Tracepoints are particularly useful when a problem, such as a hang or performance issue, is reproducible and thus the tracepoint output can be obtained during the problematic operation. In GFS2, glocks are the primary cache control mechanism and they are the key to understanding the performance of the core of GFS2. The bmap (block map) tracepoints can be used to monitor block allocations and block mapping (lookup of already allocated blocks in the on-disk metadata tree) as they happen and check for any issues relating to locality of access. The log tracepoints keep track of the data being written to and released from the journal and can provide useful information on that part of GFS2.

The tracepoints are designed to be as generic as possible. This should mean that it will not be necessary to change the API during the course of Red Hat Enterprise Linux 6. On the other hand, users of this interface should be aware that this is a debugging interface and not part of the normal Red Hat Enterprise Linux 6 API set, and as such Red Hat makes no guarantees that changes in the GFS2 tracepoints interface will not occur.

Tracepoints are a generic feature of Red Hat Enterprise Linux 6 and their scope goes well beyond GFS2. In particular they are used to implement the **blktrace** infrastructure and the **blktrace** tracepoints can be used in combination with those of GFS2 to gain a fuller picture of the system performance. Due to the level at which the tracepoints operate, they can produce large volumes of data in a very short period of time. They are designed to put a minimum load on the system when they are enabled, but it is inevitable that they will have some effect. Filtering events via a variety of means can help reduce the volume of data and help focus on obtaining just the information which is useful for understanding any particular situation.

C.2. Tracepoints

The tracepoints can be found under `/sys/kernel/debug/tracing/` directory assuming that **debugfs** is mounted in the standard place at the `/sys/kernel/debug` directory. The **events** subdirectory contains all the tracing events that may be specified and, provided the **gfs2** module is loaded, there will be a **gfs2** subdirectory containing further subdirectories, one for each GFS2 event. The contents of the `/sys/kernel/debug/tracing/events/gfs2` directory should look roughly like the following:

```
[root@chywoon gfs2]# ls
enable          gfs2_bmap          gfs2_glock_queue    gfs2_log_flush
filter          gfs2_demote_rq    gfs2_glock_state_change gfs2_pin
gfs2_block_alloc gfs2_glock_put    gfs2_log_blocks     gfs2_promote
```

To enable all the GFS2 tracepoints, run the following command:

```
[root@chywoon gfs2]# echo -n 1 >/sys/kernel/debug/tracing/events/gfs2/enable
```

To enable a specific tracepoint, there is an **enable** file in each of the individual event subdirectories. The same is true of the **filter** file which can be used to set an event filter for each event or set of events. The meaning of the individual events is explained in more detail below.

The output from the tracepoints is available in ASCII or binary format. This appendix does not currently cover the binary interface. The ASCII interface is available in two ways. To list the current content of the ring buffer, you can run the following command:

```
[root@chywoon gfs2]# cat /sys/kernel/debug/tracing/trace
```

This interface is useful in cases where you are using a long-running process for a certain period of time and, after some event, want to look back at the latest captured information in the buffer. An alternative interface, `/sys/kernel/debug/tracing/trace_pipe`, can be used when all the output is required. Events are read from this file as they occur; there is no historical information available via this interface. The format of the output is the same from both interfaces and is described for each of the GFS2 events in the later sections of this appendix.

A utility called **trace-cmd** is available for reading tracepoint data. For more information on this utility, refer to the link in [Section C.9, “References”](#). The **trace-cmd** utility can be used in a similar way to the **strace** utility, for example to run a command while gathering trace data from various sources.

C.3. Glocks

To understand GFS2, the most important concept to understand, and the one which sets it aside from other file systems, is the concept of glocks. In terms of the source code, a glock is a data structure that brings together the DLM and caching into a single state machine. Each glock has a 1:1 relationship with a single DLM lock, and provides caching for that lock state so that repetitive operations carried out from a single node of the file system do not have to repeatedly call the DLM, and thus they help avoid unnecessary network traffic. There are two broad categories of glocks, those which cache metadata and those which do not. The inode glocks and the resource group glocks both cache metadata, other types of glocks do not cache metadata. The inode glock is also involved in the caching of data in addition to metadata and has the most complex logic of all glocks.

Table C.1. Glock Modes and DLM Lock Modes

Glock mode	DLM lock mode	Notes
UN	IV/NL	Unlocked (no DLM lock associated with glock or NL lock depending on I flag)
SH	PR	Shared (protected read) lock
EX	EX	Exclusive lock
DF	CW	Deferred (concurrent write) used for Direct I/O and file system freeze

Glocks remain in memory until either they are unlocked (at the request of another node or at the request of the VM) and there are no local users. At that point they are removed from the glock hash table and freed. When a glock is created, the DLM lock is not associated with the glock immediately. The DLM lock becomes associated with the glock upon the first request to the DLM, and if this request is successful then the 'I' (initial) flag will be set on the glock. [Table C.4, “Glock flags”](#) shows the meanings of the different glock flags. Once the DLM has been associated with the glock, the DLM lock will always remain at least at NL (Null) lock mode until the glock is to be freed. A demotion of the DLM lock from NL to unlocked is always the last operation in the life of a glock.



Note

This particular aspect of DLM lock behavior has changed since Red Hat Enterprise Linux 5, which does sometimes unlock the DLM locks attached to glocks completely, and thus Red Hat Enterprise Linux 5 has a different mechanism to ensure that LVBs (lock value blocks) are preserved where required. The new scheme that Red Hat Enterprise Linux 6 uses was made possible due to the merging of the `lock_dlm` lock module (not to be confused with the DLM itself) into GFS2.

Each glock can have a number of "holders" associated with it, each of which represents one lock request from the higher layers. System calls relating to GFS2 queue and dequeue holders from the glock to protect the critical section of code.

The glock state machine is based on a workqueue. For performance reasons, tasklets would be preferable; however, in the current implementation we need to submit I/O from that context which prohibits their use.



Note

Workqueues have their own tracepoints which can be used in combination with the GFS2 tracepoints if desired

[Table C.2, "Glock Modes and Data Types"](#) shows what state may be cached under each of the glock modes and whether that cached state may be dirty. This applies to both inode and resource group locks, although there is no data component for the resource group locks, only metadata.

Table C.2. Glock Modes and Data Types

Glock mode	Cache Data	Cache Metadata	Dirty Data	Dirty Metadata
UN	No	No	No	No
SH	Yes	Yes	No	No
DF	No	Yes	No	No
EX	Yes	Yes	Yes	Yes

C.4. The glock debugfs Interface

The glock `debugfs` interface allows the visualization of the internal state of the glocks and the holders and it also includes some summary details of the objects being locked in some cases. Each line of the file either begins G: with no indentation (which refers to the glock itself) or it begins with a different letter, indented with a single space, and refers to the structures associated with the glock immediately above it in the file (H: is a holder, I: an inode, and R: a resource group). Here is an example of what the content of this file might look like:

```
G: s:SH n:5/75320 f:I t:SH d:EX/0 a:0 r:3
  H: s:SH f:EH e:0 p:4466 [postmark] gfs2_inode_lookup+0x14e/0x260 [gfs2]
  G: s:EX n:3/258028 f:yI t:EX d:EX/0 a:3 r:4
```

Appendix C. GFS2 tracepoints and the debugfs glocks File

```
H: s:EX f:tH e:0 p:4466 [postmark] gfs2_inplace_reserve_i+0x177/0x780 [gfs2]
R: n:258028 f:05 b:22256/22256 i:16800
G: s:EX n:2/219916 f:yfI t:EX d:EX/0 a:0 r:3
I: n:75661/219916 t:8 f:0x10 d:0x00000000 s:7522/7522
G: s:SH n:5/127205 f:I t:SH d:EX/0 a:0 r:3
H: s:SH f:EH e:0 p:4466 [postmark] gfs2_inode_lookup+0x14e/0x260 [gfs2]
G: s:EX n:2/50382 f:yfI t:EX d:EX/0 a:0 r:2
G: s:SH n:5/302519 f:I t:SH d:EX/0 a:0 r:3
H: s:SH f:EH e:0 p:4466 [postmark] gfs2_inode_lookup+0x14e/0x260 [gfs2]
G: s:SH n:5/313874 f:I t:SH d:EX/0 a:0 r:3
H: s:SH f:EH e:0 p:4466 [postmark] gfs2_inode_lookup+0x14e/0x260 [gfs2]
G: s:SH n:5/271916 f:I t:SH d:EX/0 a:0 r:3
H: s:SH f:EH e:0 p:4466 [postmark] gfs2_inode_lookup+0x14e/0x260 [gfs2]
G: s:SH n:5/312732 f:I t:SH d:EX/0 a:0 r:3
H: s:SH f:EH e:0 p:4466 [postmark] gfs2_inode_lookup+0x14e/0x260 [gfs2]
```

The above example is a series of excerpts (from an approximately 18MB file) generated by the command `cat /sys/kernel/debug/gfs2/unity:myfs/glocks >my.lock` during a run of the postmark benchmark on a single node GFS2 file system. The glocks in the figure have been selected in order to show some of the more interesting features of the glock dumps.

The glock states are either EX (exclusive), DF (deferred), SH (shared) or UN (unlocked). These states correspond directly with DLM lock modes except for UN which may represent either the DLM null lock state, or that GFS2 does not hold a DLM lock (depending on the I flag as explained above). The s: field of the glock indicates the current state of the lock and the same field in the holder indicates the requested mode. If the lock is granted, the holder will have the H bit set in its flags (f: field). Otherwise, it will have the W wait bit set.

The n: field (number) indicates the number associated with each item. For glocks, that is the type number followed by the glock number so that in the above example, the first glock is n:5/75320; that is, an **ioopen** glock which relates to inode 75320. In the case of inode and **ioopen** glocks, the glock number is always identical to the inode's disk block number.



Note

The glock numbers (n: field) in the debugfs glocks file are in hexadecimal, whereas the tracepoints output lists them in decimal. This is for historical reasons; glock numbers were always written in hex, but decimal was chosen for the tracepoints so that the numbers could easily be compared with the other tracepoint output (from **blktrace** for example) and with output from **stat(1)**.

The full listing of all the flags for both the holder and the glock are set out in [Table C.4, “Glock flags”](#) and [Table C.5, “Glock holder flags”](#). The content of lock value blocks is not currently available via the glock **debugfs** interface.

[Table C.3, “Glock types”](#) shows the meanings of the different glock types.

Table C.3. Glock types

Type number	Lock type	Use
1	trans	Transaction lock
2	inode	Inode metadata and data
3	rgrp	Resource group metadata

Type number	Lock type	Use
4	meta	The superblock
5	iopen	Inode last closer detection
6	flock	flock(2) syscall
8	quota	Quota operations
9	journal	Journal mutex

One of the more important glock flags is the I (locked) flag. This is the bit lock that is used to arbitrate access to the glock state when a state change is to be performed. It is set when the state machine is about to send a remote lock request via the DLM, and only cleared when the complete operation has been performed. Sometimes this can mean that more than one lock request will have been sent, with various invalidations occurring between times.

Table C.4, "Glock flags" shows the meanings of the different glock flags.

Table C.4. Glock flags

Flag	Name	Meaning
d	Pending demote	A deferred (remote) demote request
D	Demote	A demote request (local or remote)
f	Log flush	The log needs to be committed before releasing this glock
F	Frozen	Replies from remote nodes ignored - recovery is in progress.
i	Invalidate in progress	In the process of invalidating pages under this glock
I	Initial	Set when DLM lock is associated with this glock
l	Locked	The glock is in the process of changing state
p	Demote in progress	The glock is in the process of responding to a demote request
r	Reply pending	Reply received from remote node is awaiting processing
y	Dirty	Data needs flushing to disk before releasing this glock

When a remote callback is received from a node that wants to get a lock in a mode that conflicts with that being held on the local node, then one or other of the two flags D (demote) or d (demote pending) is set. In order to prevent starvation conditions when there is contention on a particular lock, each lock is assigned a minimum hold time. A node which has not yet had the lock for the minimum hold time is allowed to retain that lock until the time interval has expired.

If the time interval has expired, then the D (demote) flag will be set and the state required will be recorded. In that case the next time there are no granted locks on the holders queue, the lock will be demoted. If the time interval has not expired, then the d (demote pending) flag is set instead. This also schedules the state machine to clear d (demote pending) and set D (demote) when the minimum hold time has expired.

The I (initial) flag is set when the glock has been assigned a DLM lock. This happens when the glock is first used and the I flag will then remain set until the glock is finally freed (which the DLM lock is unlocked).

C.5. Glock Holders

Table C.5, “Glock holder flags” shows the meanings of the different glock holder flags.

Table C.5. Glock holder flags

Flag	Name	Meaning
a	Async	Do not wait for glock result (will poll for result later)
A	Any	Any compatible lock mode is acceptable
c	No cache	When unlocked, demote DLM lock immediately
e	No expire	Ignore subsequent lock cancel requests
E	Exact	Must have exact lock mode
F	First	Set when holder is the first to be granted for this lock
H	Holder	Indicates that requested lock is granted
p	Priority	Enqueue holder at the head of the queue
t	Try	A "try" lock
T	Try 1CB	A "try" lock that sends a callback
W	Wait	Set while waiting for request to complete

The most important holder flags are H (holder) and W (wait) as mentioned earlier, since they are set on granted lock requests and queued lock requests respectively. The ordering of the holders in the list is important. If there are any granted holders, they will always be at the head of the queue, followed by any queued holders.

If there are no granted holders, then the first holder in the list will be the one that triggers the next state change. Since demote requests are always considered higher priority than requests from the file system, that might not always directly result in a change to the state requested.

The glock subsystem supports two kinds of "try" lock. These are useful both because they allow the taking of locks out of the normal order (with suitable back-off and retry) and because they can be used to help avoid resources in use by other nodes. The normal t (try) lock is basically just what its name indicates; it is a "try" lock that does not do anything special. The T (**try 1CB**) lock, on the other hand, is identical to the t lock except that the DLM will send a single callback to current incompatible lock holders. One use of the T (**try 1CB**) lock is with the **iopen** locks, which are used to arbitrate among the nodes when an inode's **i_nlink** count is zero, and determine which of the nodes will be responsible for deallocating the inode. The **iopen** glock is normally held in the shared state, but when the **i_nlink** count becomes zero and **->delete_inode()** is called, it will request an exclusive lock with T (**try 1CB**) set. It will continue to deallocate the inode if the lock is granted. If the lock is not granted it will result in the node(s) which were preventing the grant of the lock marking their glock(s) with the D (demote) flag, which is checked at **->drop_inode()** time in order to ensure that the deallocation is not forgotten.

This means that inodes that have zero link count but are still open will be deallocated by the node on which the final **close()** occurs. Also, at the same time as the inode's link count is decremented to zero the inode is marked as being in the special state of having zero link count but still in use in the resource group bitmap. This functions like the ext3 file system's orphan list in that it allows any subsequent reader of the bitmap to know that there is potentially space that might be reclaimed, and to attempt to reclaim it.

C.6. Glock tracepoints

The tracepoints are also designed to be able to confirm the correctness of the cache control by combining them with the `blktrace` output and with knowledge of the on-disk layout. It is then possible to check that any given I/O has been issued and completed under the correct lock, and that no races are present.

The `gfs2_glock_state_change` tracepoint is the most important one to understand. It tracks every state change of the glock from initial creation right through to the final demotion which ends with `gfs2_glock_put` and the final NL to unlocked transition. The I (locked) glock flag is always set before a state change occurs and will not be cleared until after it has finished. There are never any granted holders (the H glock holder flag) during a state change. If there are any queued holders, they will always be in the W (waiting) state. When the state change is complete then the holders may be granted which is the final operation before the I glock flag is cleared.

The `gfs2_demote_rq` tracepoint keeps track of demote requests, both local and remote. Assuming that there is enough memory on the node, the local demote requests will rarely be seen, and most often they will be created by `umount` or by occasional memory reclaim. The number of remote demote requests is a measure of the contention between nodes for a particular inode or resource group.

When a holder is granted a lock, `gfs2_promote` is called, this occurs as the final stages of a state change or when a lock is requested which can be granted immediately due to the glock state already caching a lock of a suitable mode. If the holder is the first one to be granted for this glock, then the f (first) flag is set on that holder. This is currently used only by resource groups.

C.7. Bmap tracepoints

Block mapping is a task central to any file system. GFS2 uses a traditional bitmap-based system with two bits per block. The main purpose of the tracepoints in this subsystem is to allow monitoring of the time taken to allocate and map blocks.

The `gfs2_bmap` tracepoint is called twice for each bmap operation: once at the start to display the bmap request, and once at the end to display the result. This makes it easy to match the requests and results together and measure the time taken to map blocks in different parts of the file system, different file offsets, or even of different files. It is also possible to see what the average extent sizes being returned are in comparison to those being requested.

To keep track of allocated blocks, `gfs2_block_alloc` is called not only on allocations, but also on freeing of blocks. Since the allocations are all referenced according to the inode for which the block is intended, this can be used to track which physical blocks belong to which files in a live file system. This is particularly useful when combined with `blktrace`, which will show problematic I/O patterns that may then be referred back to the relevant inodes using the mapping gained via this tracepoint.

C.8. Log tracepoints

The tracepoints in this subsystem track blocks being added to and removed from the journal (`gfs2_pin`), as well as the time taken to commit the transactions to the log (`gfs2_log_flush`). This can be very useful when trying to debug journaling performance issues.

The `gfs2_log_blocks` tracepoint keeps track of the reserved blocks in the log, which can help show if the log is too small for the workload, for example.

The `gfs2_ail_flush` tracepoint (Red Hat Enterprise Linux 6.2 and later) is similar to the `gfs2_log_flush` tracepoint in that it keeps track of the start and end of flushes of the AIL list. The AIL list contains buffers which have been through the log, but have not yet been written back in place

and this is periodically flushed in order to release more log space for use by the filesystem, or when a process requests a sync or fsync.

C.9. References

For more information about tracepoints and the GFS2 **glocks** file, refer to the following resources:

- This appendix has been partially adapted from a paper delivered by Steve Whitehouse at Linux Symposium 2009, which can be found at <http://git.kernel.org/?p=linux/kernel/git/torvalds/linux-2.6.git;a=blob;f=Documentation/filesystems/gfs2-glocks.txt;h=0494f78d87e40c225eb1dc1a1489acd891210761;hb=HEAD>¹.
- For information on glock internal locking rules, see <http://git.kernel.org/?p=linux/kernel/git/torvalds/linux-2.6.git;a=blob;f=Documentation/filesystems/gfs2-glocks.txt;h=0494f78d87e40c225eb1dc1a1489acd891210761;hb=HEAD>².
- For information on event tracing, see <http://git.kernel.org/?p=linux/kernel/git/torvalds/linux-2.6.git;a=blob;f=Documentation/trace/events.txt;h=09bd8e9029892e4e1d48078de4d076e24eff3dd2;hb=HEAD>³.
- For information on the **trace-cmd** utility, see <http://lwn.net/Articles/341902/>⁴.

Appendix D. Revision History

Revision 3.0-2 Thu Dec 1 2011

Steven Levine slevine@redhat.com

Release for GA of Red Hat Enterprise Linux 6.2

Revision 3.0-1 Mon Sep 19 2011

Steven Levine slevine@redhat.com

Initial revision for Red Hat Enterprise Linux 6.2 Beta release

Resolves: #704179

Documents support for the **tunegfs2** command.

Resolves: #712390

Adds new appendix on GFS2 tracepoints.

Resolves: #705961

Resolves minor typographical errors.

Revision 2.0-1 Thu May 19 2011

Steven Levine slevine@redhat.com

Initial release for Red Hat Enterprise Linux 6.1

Resolves: #549838

Documents support for standard Linux quota facilities in Red Hat Enterprise Linux 6.1.

Resolves: #608750

Clarifies description of GFS2 withdraw function.

Resolves: #660364

Corrects maximum GFS2 file system size information.

Resolves: #687874

Adds new chapter on GFS2 troubleshooting.

Resolves: #664848

Adds information on finding Context-Dependent Path Names before converting from GFS to GFS2.

Revision 1.0-1 Wed Nov 15 2010

Steven Levine slevine@redhat.com

Initial release for Red Hat Enterprise Linux 6

Index

A

- acl mount option, 19
- adding journals to a file system, 30
- atime, configuring updates, 33
 - mounting with noatime , 34
 - mounting with relatime , 33
- audience, v

B

- bind mount
 - mount order, 38
- bind mounts, 37

C

- configuration, before, 3
- configuration, initial,
 - prerequisite tasks, 13

D

- data journaling, 32
- debugfs,
 - debugfs file, 9
- disk quotas
 - additional resources, 28
 - assigning per group, 26
 - assigning per user, 25
 - enabling, 23
 - creating quota files, 24
 - quotacheck, running, 24
 - hard limit, 25
 - management of, 26
 - quotacheck command, using to check, 27
 - reporting, 26
 - soft limit, 25

F

- features, new and changed, 2
- feedback
 - contact information for this manual, v
- file system
 - adding journals, 30
 - atime, configuring updates, 33
 - mounting with noatime , 34
 - mounting with relatime , 33
 - bind mounts, 37
 - context-dependent path names (CDPNs), 37
 - data journaling, 32
 - growing, 29
 - making, 15
 - mount order, 38

- mounting, 18, 22
- quota management, 23, 23,
 - displaying quota limits, 48
 - enabling quota accounting, 51
 - enabling/disabling quota enforcement, 51
 - setting quotas, 47
 - synchronizing quotas, 27, 50
- repairing, 35
- suspending activity, 34
- unmounting, 22, 22
- fsck.gfs2 command, 35

G

- GFS2
 - atime, configuring updates, 33
 - mounting with noatime , 34
 - mounting with relatime , 33
 - managing,
 - quota management, 23, 23,
 - displaying quota limits, 48
 - enabling quota accounting, 51
 - enabling/disabling quota enforcement, 51
 - setting quotas, 47
 - synchronizing quotas, 27, 50
 - withdraw function, 40
 - GFS2 file system maximum size,
 - GFS2-specific options for adding journals table, 31
 - GFS2-specific options for expanding file systems table, 30
 - gfs2_grow command, 29
 - gfs2_jadd command, 30
 - gfs2_quota command,
 - glock,
 - glock flags, 10, 59
 - glock holder flags, 10, 60
 - glock types, 11, 58
 - growing a file system, 29

I

- initial tasks
 - setup, initial, 13
- introduction,
 - audience, v

M

- making a file system, 15
- managing GFS2,
- maximum size, GFS2 file system,
- mkfs command, 15
- mkfs.gfs2 command options table, 17
- mount command, 18
- mount table, 20

mounting a file system, 18, 22

N

node locking, 7

O

overview,

- configuration, before, 3

- features, new and changed, 2

P

path names, context-dependent (CDPNs), 37

performance tuning, 8

preface (see introduction)

prerequisite tasks

- configuration, initial, 13

Q

quota management, 23, 23,

- displaying quota limits, 48

- enabling quota accounting, 51

- enabling/disabling quota enforcement, 51

- setting quotas, 47

- synchronizing quotas, 27, 50

quota= mount option, 47

quotacheck , 24

quotacheck command

- checking quota accuracy with, 27

quota_quantum tunable parameter, 27, 50

R

repairing a file system, 35

S

setup, initial

- initial tasks, 13

suspending activity on a file system, 34

system hang at unmount, 22

T

tables

- GFS2-specific options for adding journals, 31

- GFS2-specific options for expanding file systems, 30

- mkfs.gfs2 command options, 17

- mount options, 20

tracepoints,

tuning, performance, 8

U

umount command, 22

umount, system hang, 22

unmounting a file system, 22, 22

W

withdraw function, GFS2, 40