General Parallel File System

Problem Determination Guide

Version 3 Release 2.1

GA76-0415-02

General Parallel File System

Problem Determination Guide

Version 3 Release 2.1

GA76-0415-02

Note:
Before using this information and the product it supports, be sure to read the general information under “Notices” on page 219.

Third Edition (August 2008)

This edition applies to version 3, release 2, modification 1 of IBM General Parallel File System Multiplatform (product number 5724-N94), IBM General Parallel File System for POWER™ (product number 5765-G66), and to all subsequent releases and modifications until otherwise indicated in new editions. Technical changes or additions to the text and illustrations are indicated by a vertical line (|) to the left of the change.

IBM welcomes your comments. A form for your comments may be provided at the back of this publication, or you may send your comments to this address:

International Business Machines Corporation
Department 58HA, Mail Station P181
2455 South Road
Poughkeepsie, NY 12601-5400
United States of America

FAX (United States and Canada): 1+845+432-9405
FAX (Other Countries): Your International Access Code +1+845+432-9405

IBMLink™ (United States customers only): IBMUSM10(MHVRCFS)
Internet e-mail: mhvrcfs@us.ibm.com

If you want a reply, be sure to include your name, address, telephone number, or FAX number. Make sure to include the following in your comment or note:

- Title and order number of this publication
- Page number or topic related to your comment

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

Contents

About this information

This information explains how to handle problems you may encounter with the General Parallel File System™ (GPFS™) licensed program.

This edition applies to GPFS version 3.2.1 for AIX®, Linux®, and Windows®.

To find out which version of GPFS is running on a particular AIX node, enter:

```
lslpp -l gpfs\*
```

To find out which version of GPFS is running on a particular Linux node, enter:

```
rpm -qa | grep gpfs
```

To find out which version of GPFS is running on a particular Windows node, use the graphical user interface (GUI) and follow these steps:

1. Click **Control Panel**–>**Add or Remove Programs**
2. Click **IBM General Parallel File System** and choose **Click here for support information.**

Throughout this information you will see various command and component names beginning with the prefix **mm**. This is not an error. GPFS shares many components with the related products IBM® Multi-Media Server and IBM Video Charger.

This information contains explanations of GPFS error messages. Informational messages are not documented.

Who should read this information

This information is intended for system administrators of GPFS systems. This information assumes that you are, and it is particularly important that you be, experienced with and understand your operating system and the subsystems used to manage your disks. This information also assumes that you are familiar with the GPFS concepts presented in the *General Parallel File System: Concepts, Planning, and Installation Guide*.

Use this information when diagnosing problems encountered with the usage of GPFS.

Conventions used in this information

Table 1 describes the typographic conventions used in this information.

Note: Throughout this information, UNIX® file name conventions are used. For example, the name of the main GPFS log file is specified as **/var/adm/ras/mmfs.log.latest**. On Windows, the UNIX name space starts under the **%SystemRoot%\SUA** directory, and UNIX-style file names need to be converted accordingly. For example, the log file name mentioned above would become **C:\windows\SUA\var\adm\ras\mmfs.log.latest**.

Table 1. Typographic conventions

Typographic convention	Usage
Bold	**Bold** words or characters represent system elements that you must use literally, such as commands, flags, path names, directories, file names, values, and selected menu options.
Bold Underlined	**Bold Underlined** keywords are defaults. These take effect if you fail to specify a different keyword.

Table 1. Typographic conventions (continued)

Typographic convention	Usage
Italic	• *Italic* words or characters represent variable values that you must supply. • *Italics* are also used for publication titles and for general emphasis in text.
`Constant width`	All of the following are displayed in `constant width` typeface: • Displayed information • Message text • Example text • Specified text typed by the user • Field names as displayed on the screen • Prompts from the system • References to example text
[ ]	Brackets enclose optional items in format and syntax descriptions.
{ }	Braces enclose a list from which you must choose an item in format and syntax descriptions.
|	A vertical bar separates items in a list of choices. (In other words, it means ″or″)
< >	Angle brackets (less-than and greater-than) enclose the name of a key on the keyboard. For example, **<Enter>** refers to the key on your terminal or workstation that is labeled with the word Enter.
...	An ellipsis indicates that you can repeat the preceding item one or more times.
<Ctrl-*x*>	The notation **<Ctrl-*x*>** indicates a control character sequence. For example, **<Ctrl-*c*>** means that you hold down the control key while pressing **<c>**.
\	The continuation character is used in programming examples in this information for formatting purposes.

Prerequisite and related information

For updates to this information, see publib.boulder.ibm.com/infocenter/clresctr/topic/com.ibm.cluster.gpfs.doc/gpfsbooks.html.

For the latest support information, see the GPFS Frequently Asked Questions at publib.boulder.ibm.com/infocenter/clresctr/topic/com.ibm.cluster.gpfs.doc/gpfs_faqs/gpfsclustersfaq.html.

ISO 9000

ISO 9000 registered quality systems were used in the development and manufacturing of this product.

Using LookAt to look up message explanations

LookAt is an online facility that lets you look up explanations for most of the IBM messages you encounter, as well as for some system abends and codes. You can use LookAt from the following locations to find IBM message explanations for Clusters software products:

- The Internet. You can access IBM message explanations directly from the LookAt Web site: http://www.ibm.com/systems/z/os/zos/bkserv/lookat/
- Your wireless handheld device. You can use the LookAt Mobile Edition with a handheld device that has wireless access and an Internet browser (for example, Internet Explorer for Pocket PCs, Blazer, or Eudora for Palm OS, or Opera for Linux handheld devices). Link to the LookAt Mobile Edition from the LookAt Web site.

How to send your comments

Your feedback is important in helping us to produce accurate, high-quality information. If you have any comments about this information or any other GPFS documentation:

- Send your comments by e-mail to: mhvrcfs@us.ibm.com.

  Include the publication title and order number, and, if applicable, the specific location of the information you have comments on (for example, a page number or a table number).
- Fill out one of the forms at the back of this information and return it by mail, by fax, or by giving it to an IBM representative.

To contact the IBM cluster development organization, send your comments by e-mail to: cluster@us.ibm.com.

Summary of changes

The following sections summarize changes to the GPFS licensed program and the GPFS library for version 3, release 2, modification 1. Within each information unit in the library, a vertical line to the left of text and illustrations indicates technical changes or additions made to the previous edition of the book.

Summary of changes
for GPFS Version 3, Release 2, Modification 1
as updated, August 2008

Changes to GPFS and to the GPFS library for version 3, release 2, modification 1 include:

- **New information**
  - GPFS for Windows Multiplatform, V3.2.1 supports the Windows Server 2003 R2 operating system running on 64-bit architectures (AMD x64 / EM64T). GPFS on Windows participates in a new or existing GPFS V3.2 cluster in conjunction with AIX and Linux (32- or 64-bit) operating systems.
  - Identity mapping between Windows and UNIX user accounts is one of the key advancements delivered in GPFS for Windows Multiplatform. System administrators can explicitly match users and groups defined on UNIX with those defined on Windows. This allows users to maintain file ownership and access rights from either platform. System administrators are not required to define an identity map. GPFS automatically creates a mapping when one is not defined. For more information about identity mapping, see the *General Parallel File System: Concepts, Planning, and Installation Guide* and the *General Parallel File System: Advanced Administration Guide*.
  - IBM has enhanced many of the details within GPFS to support the unique semantic requirements posed by Windows. These include case insensitive names, NTFS-like file attributes, and Windows file locking. GPFS provides a bridge between a Windows and POSIX view of files, while not adversely affecting the long-standing capabilities provided on AIX and Linux operating systems.
  - GPFS for Windows Multiplatform provides the same core services to parallel and serial applications as are available on AIX and Linux operating systems. GPFS allows parallel applications simultaneous access to the same files, or different files, from any node that has the GPFS file system mounted while managing a high level of control over all file system operations. System administrators and users have a consistent command interface on AIX, Linux, and Windows operating systems.

    The following commands have been updated for Windows:

    - **mmchfs** to add the **-t** *DriveLetter* option
    - **mmcrfs** to add the **-t** *DriveLetter* option
    - **mmlsfs** to add the **-t** option to display the Windows drive letter
    - **mmmount** to add the *DefaultDriveLetter* and *DriveLetter* parameters
    - **mmumount** to add the *DefaultDriveLetter* and *DriveLetter* parameters

    With few exceptions, the commands supported on the Windows operating system are identical to the commands supported on other GPFS platforms. For a list of unsupported commands, see the *General Parallel File System: Concepts, Planning, and Installation Guide*.
  - GPFS for Windows Multiplatform, V3.2.1 does not support or has restricted support for some features. For a complete list of these limitations, see the *General Parallel File System: Concepts, Planning, and Installation Guide*.
- **Changed information:**

  Minor editorial updates marked by a vertical line to the left of the text.
- **Deleted information:**

  There has been no information deleted from the GPFS library for GPFS V3.2.1.

Chapter 1. Logs, traces and dumps

The problem determination tools provided with General Parallel File System (GPFS) are intended for use by experienced system administrators who know how to collect data and run debugging routines.

GPFS has it own error log, but the operating system error log is also useful because it contains information about hardware failures and operating system or other software failures that can affect GPFS.

GPFS also provides a snap dump, trace, and other utilities that can be used to obtain detailed information about specific problems.

The information is organized as follows:

Note: GPFS error logs and messages contain the MMFS prefix. This is intentional, because GPFS shares many components with the IBM Multi-Media LAN Server, a related licensed program.

The MMFS log

GPFS writes both operational messages and error data to the **MMFS** log file. The **MMFS** log can be found in the **/var/adm/ras** directory on each node.

The **MMFS** log file is named **mmfs.log.***date***.***nodeName***,** where *date* is the time stamp when the instance of GPFS started on the node and *nodeName* is the name of the node. The latest **mmfs** log file can be found by using the symbolic file name **/var/adm/ras/mmfs.log.latest**.

The **MMFS** log from the prior startup of GPFS can be found by using the symbolic file name **/var/adm/ras/mmfs.log.previous**. All other files have a timestamp and node name appended to the file name.

At GPFS startup, files that have not been accessed during the last ten days are deleted. If you want to save old files, copy them elsewhere.

This example shows normal operational messages that appear in the **MMFS** log file:

```
Wed Mar 26 12:54:39 EDT 2008 runmmfs starting
Loading kernel extension from /usr/lpp/mmfs/bin . . .
/usr/lpp/mmfs/bin/mmfskxload: /usr/lpp/mmfs/bin/aix64/mmfs64 is already loaded at 68647936.
Wed Mar 26 12:54:40.120 2008: mmfsd64 initializing. {Version: 3.2.1.2 Built: Mar 25 2008 20:09:09} ...
Wed Mar 26 12:54:50.368 2008: OpenSSL library loaded and initialized.
Wed Mar 26 12:54:51.627 2008: Connecting to 192.168.10.95 c5n95
Wed Mar 26 12:54:51.638 2008: Connected to 192.168.10.95 c5n95
Wed Mar 26 12:54:51.706 2008: Accepted and connected to 9.114.132.110 c5n110
Wed Mar 26 12:54:51.735 2008: Connecting to 192.168.10.109 c5n109
Wed Mar 26 12:54:52.388 2008: mmfsd ready
Wed Mar 26 12:54:52 EDT 2008: mmcommon mmfsup invoked
Wed Mar 26 12:54:52 EDT 2008: mmcommon tsmigratedEventNotify invoked: thisNode=9.114.132.97 \
  nodeList="9.114.132.95"
```

```
Wed Mar 26 12:54:54.649 2008: Accepted and connected to 9.114.132.92 c5n92
Wed Mar 26 12:54:54.795 2008: Accepted and connected to 192.168.10.109 c5n109
Wed Mar 26 12:54:54.912 2008: Accepted and connected to 9.114.132.98 c5n98
Wed Mar 26 13:29:02.890 2008: Recovering nodes: 9.114.132.98
Wed Mar 26 13:29:03 EDT 2008: mmcommon tsmigratedEventNotify invoked: thisNode=9.114.132.97 \
  nodeList="9.114.132.95"
Wed Mar 26 13:29:03.127 2008: Recovered 1 nodes.
Wed Mar 26 13:58:31.253 2008: Accepted and connected to 192.168.10.98 c5n98
Wed Mar 26 17:12:43.185 2008: Command: mount fs2 692280
Wed Mar 26 17:12:43.482 2008: Node 9.114.132.97 (c5n97) appointed as manager for fs2.
Wed Mar 26 17:12:44.445 2008: Node 9.114.132.97 (c5n97) completed take over for fs2.
Wed Mar 26 17:12:44.456 2008: Command: err 0: mount fs2 692280
```

Depending on the size and complexity of your system configuration, the amount of time to start GPFS varies. Taking your system configuration into consideration, after a reasonable amount of time if you cannot access a file system that has been mounted (either automatically or by issuing a **mount** command), examine the log file for error messages.

The GPFS log is a repository of error conditions that have been detected on each node, as well as operational events such as file system mounts. The GPFS log is the first place to look when attempting to debug abnormal events. Since GPFS is a cluster file system, events that occur on one node might affect system behavior on other nodes, and all GPFS logs can have relevant data.

Creating a master MMFS.log file

The **MMFS log** frequently shows problems on one node that actually originated on another node.

GPFS is a file system that runs on multiple nodes of a cluster. This means that problems originating on one node of a cluster often have effects that are visible on other nodes. It is often valuable to merge the GPFS logs in pursuit of a problem. Having accurate time stamps aids the analysis of the sequence of events.

Before following any of the debug steps, IBM suggests that you:

1. Synchronize all clocks of all nodes in the GPFS cluster. If this is not done, and clocks on different nodes are out of sync, there is no way to establish the real time line of events occurring on multiple nodes. Therefore, a merged error log is less useful for determining the origin of a problem and tracking its effects.
2. Merge and chronologically sort all of the **MMFS log** entries from each node in the cluster.

   Begin by creating a file named **/tmp/gpfs.allnodes** containing all nodes in the cluster. Refer to *General Parallel File System: Concepts, Planning, and Installation Guide* and search for *files to ease the installation process*.

   Collect the GPFS logs from the nodes using the script named **/usr/lpp/mmfs/samples/gatherlogs.sample.sh**. This script generates a **/tmp/logs.sorted** file that is comprised of the error logs from all the nodes listed in **/tmp/gpfs.allnodes**, prefixed by *nodename*, and merged into a single sorted file. The output of this script will be used in the problem determination procedures that follow.

The operating system error log facility

GPFS records file system or disk failures using the error logging facility provided by the operating system: **syslog** facility on Linux, **errpt** facility on AIX, and Event Logger facility on Windows.

The error logging facility is referred to as ″the error log″ regardless of operating-system specific error log facility naming conventions.

These failures can be viewed by issuing this command on an AIX node:

```
errpt -a
```

and this command on a Linux node:

```
grep "mmfs:" /var/log/messages
```

On Windows, use the Event Viewer utility and look for events with source label ″GPFS″ in the ″Application″ event category.

The error log contains information about several classes of events or errors. These classes are:

MMFS_ABNORMAL_SHUTDOWN

The MMFS_ABNORMAL_SHUTDOWN error log entry means that GPFS has determined that it must shutdown all operations on this node because of a problem. Insufficient memory on the node to handle critical recovery situations can cause this error. In general there will be other error log entries from GPFS or some other component associated with this error log entry.

MMFS_DISKFAIL

The MMFS_DISKFAIL error log entry indicates that GPFS has detected the failure of a disk and forced the disk to the stopped state. This is ordinarily not a GPFS error but a failure in the disk subsystem or the path to the disk subsystem.

MMFS_ENVIRON

MMFS_ENVIRON error log entry records are associated with other records of the MMFS_GENERIC or MMFS_SYSTEM_UNMOUNT types. They indicate that the root cause of the error is external to GPFS and usually in the network that supports GPFS. Check the network and its physical connections. The data portion of this record supplies the return code provided by the communications code.

MMFS_FSSTRUCT

The **MMFS_FSSTRUCT** error log entry indicates that GPFS has detected a problem with the on-disk structure of the file system. The severity of these errors depends on the exact nature of the inconsistent data structure. If it is limited to a single file, **EIO** errors will be reported to the application and operation will continue. If the inconsistency affects vital metadata structures, operation will cease on this file system. These errors are often associated with an **MMFS_SYSTEM_UNMOUNT** error log entry and will probably occur on all nodes. If the error occurs on all nodes, some critical piece of the file system is inconsistent. This can occur as a result of a GPFS error or an error in the disk system.

Issuing the **mmfsck** command might repair the error. This command can take a long time to complete, depending on the file system size. Contact the IBM Support Center before using the **mmfsck** command.

1. Issue the **mmfsck -n** command to collect data.
2. Issue the **mmfsck -y** command offline to repair the file system.

If the file system is not repaired after issuing the **mmfsck** command, contact the IBM Support Center.

MMFS_GENERIC

The MMFS_GENERIC error log entry means that GPFS self diagnostics have detected an internal error, or that additional information is being provided with an MMFS_SYSTEM_UNMOUNT report. If the record is associated with an MMFS_SYSTEM_UNMOUNT report, the event code fields in the records will be the same. The error code and return code fields might describe the error. See Chapter 10, “Messages,” on page 109 for a listing of codes generated by GPFS.

If the error is generated by the self diagnostic routines, service personnel should interpret the return and error code fields since the use of these fields varies by the specific error. Errors caused by the self checking logic will result in the shutdown of GPFS on this node.

MMFS_GENERIC errors can result from an inability to reach a critical disk resource. These errors might look different depending on the specific disk resource that has become unavailable, like logs and allocation maps. This type of error will usually be associated with other error indications. Other errors generated by disk subsystems, high availability components, and communications components at the same time as, or immediately preceding, the GPFS error should be pursued first because they might be the cause of these errors. MMFS_GENERIC error indications without an associated error of those types represent a GPFS problem that requires the IBM Support Center. See “Information to collect before contacting the IBM Support Center” on page 37.

MMFS_LONGDISKIO

The MMFS_LONGDISKIO error log entry indicates that GPFS is experiencing very long response time for disk requests. This is a warning message and can indicate that your disk system is overloaded or that a failing disk is requiring many I/O retries. Follow your operating system’s instructions for monitoring the performance of your I/O subsystem on this node and on any disk server nodes that might be involved. The data portion of this error record specifies the disk involved. There might be related error log entries from the disk subsystems that will pinpoint the actual cause of the problem. If the disk is attached to an AIX node, refer to *AIX 5L Performance Management Guide* at: publib.boulder.ibm.com/infocenter/pseries/index.jsp

The **mmpmon** command can be used to analyze I/O performance on a per-node basis. See the unit titled *Monitoring GPFS I/O performance with the **mmpmon** command* in the *General Parallel File System: Advanced Administration Guide*.

MMFS_QUOTA

The **MMFS_QUOTA** error log entry is used when GPFS detects a problem in the handling of quota information. This entry is created when the quota manager has a problem reading or writing the quota file. If the quota manager cannot read all entries in the quota file when mounting a file system with quotas enabled, the quota manager shuts down but file system manager initialization continues. Mounts will not succeed and will return an appropriate error message (see “File system forced unmount” on page 71). For information about running the **mmcheckquota** command, see “The mmcheckquota command” on page 30.

Quota accounting depends on a consistent mapping between user names and their numeric identifiers. This means that a single user accessing a quota enabled file system from different nodes should map to the same numeric user identifier from each node. Within a local cluster this is usually achieved by ensuring that **/etc/passwd** and **/etc/group** are identical across the cluster.

When accessing quota enabled file systems from other clusters, you need to either ensure individual accessing users have equivalent entries in **/etc/passwd** and **/etc/group**, or use the user identity mapping facility as outlined in *UID Mapping for GPFS in a Multi-Cluster Environment* at www.ibm.com/servers/eserver/clusters/library/wp_aix_lit.html.

It might be necessary to run an offline quota check (**mmcheckquota**) to repair or recreate the quota file. If the quota file is corrupted, **mmcheckquota** will not restore it. The file must be restored from the backup copy. If there is no backup copy, an empty file can be set as the new quota file. This is equivalent to recreating the quota file. To set an empty file or use the backup file, issue the **mmcheckquota** command with the appropriate operand:

- **-u** *UserQuotaFilename* for the user quota file
- **-g** *GroupQuotaFilename* for the group quota file
- **-j** *FilesetQuotaFilename* for the fileset quota file

Reissue the **mmcheckquota** command to check the file system inode and space usage.

MMFS_SYSTEM_UNMOUNT

The `MMFS_SYSTEM_UNMOUNT` error log entry means that GPFS has discovered a condition that might result in data corruption if operation with this file system continues from this node. GPFS has marked the file system as disconnected and applications accessing files within the file system will receive `ESTALE` errors. This can be the result of:

- The loss of a path to all disks containing a critical data structure. If you are using an IBM Virtual Shared Disk, locate material in the appropriate IBM Reliable Scalable Cluster Technology (RSCT) document at: publib.boulder.ibm.com/clresctr/windows/public/rsctbooks.html and search on *diagnosing IBM Virtual Shared Disk problems*. Follow the problem determination and repair actions specified.

  If you are using SAN attachment of your storage, consult the problem determination guides provided by your SAN switch vendor and your storage vendor.
- An internal processing error within the file system.

See “File system forced unmount” on page 71. Follow the problem determination and repair actions specified.

MMFS_SYSTEM_WARNING

The `MMFS_SYSTEM_WARNING` error log entry means that GPFS has detected a system level value approaching its maximum limit. This might occur as a result of the number of inodes (files) reaching its limit. Issue the **mmchfs** command to increase the number of inodes for the file system so there is at least a minimum of 5% free.

Error log entry example

This is an example of an error log entry that indicates a failure in either the storage subsystem or communication subsystem:

```
LABEL: MMFS_SYSTEM_UNMOUNT
IDENTIFIER: C954F85D

Date/Time: Thu Jul 8 10:17:10 CDT
Sequence Number: 25426
Machine Id: 000024994C00
Node Id: nos6
Class: S
Type: PERM
Resource Name: mmfs

Description
STORAGE SUBSYSTEM FAILURE

Probable Causes
STORAGE SUBSYSTEM
COMMUNICATIONS SUBSYSTEM

Failure Causes
STORAGE SUBSYSTEM
```

```
COMMUNICATIONS SUBSYSTEM

Recommended Actions
CONTACT APPROPRIATE SERVICE REPRESENTATIVE

Detail Data
EVENT CODE
15558007
STATUS CODE
212
VOLUME
gpfsd
```

Gathering data to solve GPFS problems

Use the **gpfs.snap** command as one of the main tools to gather preliminary data when a GPFS problem is encountered, such as a hung file system, a hung GPFS command, or an **mmfsd** daemon assert.

The **gpfs.snap** command gathers general information that can be used in conjunction with other GPFS information (for example, **mmfsadm dump all**) to solve a GPFS problem.

The syntax of the **gpfs.snap** command is:

gpfs.snap [**-c** *"CommandString"*] [**-d** *OutputDirectory*] [**-x** {**1** | **2**}] [**-y** | **-z**] [**-a** | **-N** {*Node*[,*Node ...*] | *NodeFile* | *NodeClass*}

These options are used with **gpfs.snap**:

-a Directs **gpfs.snap** to collect data from all nodes in the cluster. This is the default.

-c *"CommandString"*
Specifies the command string to run on the specified nodes. When this option is specified, the data collected is limited to the result of the specified command string; the standard data collected by **gpfs.snap** is not collected. *CommandString* can consist of multiple commands, which are separated by semicolons (;) and enclosed in double quotation marks (*"*). The **-d**, **-p**, **-x**, **-y**, and **-z** options have no effect when **-c** is specified.

-d *OutputDirectory*
Specifies the output directory. The default is **/tmp/gpfs.snapOut**.

-N {*Node***[,***Node ...***] | *NodeFile* | *NodeClass***}**
Specifies the nodes from which to collect snapshot data. This option can be used only in conjunction with the **-y** option. This option supports all defined node classes. For general information on how to specify node names, see "Specifying nodes as input to GPFS commands" in the *General Parallel File System: Administration and Programming Reference*.

-p Skips the problem determination sequence (applies only to master snaps).

-x {1 | 2}

1 Determines whether enough space exists for **gpfs.snap** to collect data. Does not actually collect data.

2 Collects data only. Does not check if enough space exists.

-y Collects snap data only from the specified nodes. Use this option in conjunction with the **-N** option.

-z Collects snap data only from the node on which **gpfs.snap** is invoked. No master data is collected.

Use the **-z** option to specify that a non-master snapshot is to be generated. This is useful if there are many nodes on which to take a snapshot, and only one master snapshot is needed. For a GPFS problem within a large cluster (hundreds or thousands of nodes), one strategy might call for a single master

snapshot and multiple non-master snapshots, (that is, one invocation of **gpfs.snap** with no options, and multiple non-master snapshots (**gpfs.snap** with the **-z** option).

Use the **-N** option to obtain **gpfs.snap** data from multiple nodes in the cluster. When these options are used, the **gpfs.snap** command takes non-master snapshots of all nodes specified with these options and a master snapshot of the node on which it was invoked.

Using the gpfs.snap command

Running the **gpfs.snap** command with no options generates a *master snap* on the node where the command is run. This invocation creates a file that is made up of multiple **gpfs.snap** snapshots. The file included is a master snapshot of the node from which the **gpfs.snap** command was invoked and non-master snapshots of each of the GPFS cluster configuration servers, as well as non-master snapshots of each of the GPFS file system manager nodes.

If the node on which the **gpfs.snap** command is run is not a file system manager node, **gpfs.snap** creates a non-master snapshot on the file system manager nodes.

The difference between a master snapshot and a non-master snapshot is the data that is gathered. A master snapshot gathers information from nodes in the cluster. A master snapshot contains all data that a non-master snapshot has. There are two categories of data that is collected:

1. Data that is always gathered by **gpfs.snap** (for master snapshots and non-master snapshots):
   - "Data always gathered by gpfs.snap on all platforms"
   - "Data always gathered by gpfs.snap on AIX" on page 8
   - "Data always gathered by gpfs.snap on Linux" on page 8
   - "Data always gathered by gpfs.snap on Windows" on page 8
2. Data that is gathered by **gpfs.snap** only in the case of a master snapshot. See "Data always gathered by gpfs.snap for a master snap" on page 9.

Data always gathered by gpfs.snap on all platforms

These items are always obtained by the **gpfs.snap** command when gathering data for an AIX, Linux, or Windows node:

1. The output of these commands:
   - **df -k**
   - **ifconfig** *interface*
   - **ipcs -a**
   - **ls -l /dev**
   - **mmfsadm dump alloc hist**
   - **mmfsadm dump alloc stats**
   - **mmfsadm dump allocmgr**
   - **mmfsadm dump allocmgr hist**
   - **mmfsadm dump allocmgr stats**
   - **mmfsadm dump cfgmgr**
   - **mmfsadm dump config**
   - **mmfsadm dump dealloc stats**
   - **mmfsadm dump disk**
   - **mmfsadm dump mmap**
   - **mmfsadm dump mutex**
   - **mmfsadm dump nsd**

- **mmfsadm dump sgmgr**
- **mmfsadm dump stripe**
- **mmfsadm dump tscomm**
- **mmfsadm dump version**
- **mmfsadm dump waiters**
- **netstat** with the **-i**, **-r**, **-rn**, **-s**, and **-v** options
- **ps -edf**
- **vmstat**

2. The contents of these files:
   - **/etc/filesystems**
   - **/etc/fstab**
   - **/etc/syslog.conf**
   - **/etc/trcfmt**
   - **/tmp/mmfs/internal***
   - **/tmp/mmfs/trcrpt***
   - **/var/adm/ras/mmfs.log.***
   - **/var/mmfs/gen/mmfs.cfg**
   - **/var/mmfs/gen/mmsdrfs**
   - **/var/mmfs/gen/nsdmap**
   - **/var/mmfs/ssl/*** except for **complete.map** and **id_rsa** files

Data always gathered by gpfs.snap on AIX

These items are always obtained by the **gpfs.snap** command when gathering data for an AIX node:

1. The output of these commands is:
   - **errpt -a**
   - **lssrc -a**
   - **lslpp -hac**
   - **no -a**
2. The contents of this file **/var/mmfs/gen/nsdpvol**
3. If IBM virtual shared disks are present, the output of the **lsvsd -l** command

Data always gathered by gpfs.snap on Linux

These items are always obtained by the **gpfs.snap** command when gathering data for a Linux node:

1. The output of the **rpm -qa** command
2. The contents of these files:
   - **/etc/*release**
   - **/var/log/messages***

Data always gathered by gpfs.snap on Windows

These items are always obtained by the **gpfs.snap** command when gathering data for a Windows node:

1. The output from **systeminfo.exe**
2. Any raw trace files ***.tmf** and **mmfs.trc***
3. The ***.pdb** symbols from **/usr/lpp/mmfs/bin/symbols**

Data always gathered by gpfs.snap for a master snap

When the **gpfs.snap** command is specified with no options, a master snapshot is taken on the node where the command was issued. All of the information from "Data always gathered by gpfs.snap on all platforms" on page 7, "Data always gathered by gpfs.snap on AIX" on page 8, "Data always gathered by gpfs.snap on Linux" on page 8, and "Data always gathered by gpfs.snap on Windows" on page 8 is obtained, as well as this data:

1. The output of these commands:
   - **mmauth**
   - **mmdf**
   - **mmgetstate -a**
   - **mmlsdisk**
   - **mmlsfs**
   - **mmlsfileset**
   - **mmlsnode -a**
   - **mmlspolicy**
   - **mmlssnapshot**
   - **tsstatus**
2. The contents of the **/var/adm/ras/mmfs.log.*** file (on all nodes in the cluster)

The mmfsadm command

The **mmfsadm** command is intended for use by trained service personnel. IBM suggests you do not run this command except under the direction of such personnel.

Note: The contents of **mmfsadm** output might vary from release to release and make obsolete any user programs that depend on that output. Therefore, we suggest that you do not create user programs that invoke **mmfsadm**.

The **mmfsadm** command extracts data from GPFS without using locking, so that it can collect the data in the event of locking errors. In certain rare cases, this can cause GPFS or the node to fail. Several options of this command exist and might be required for use:

cleanup
: Delete shared segments left by a previously failed GPFS daemon without actually restarting the daemon.

dump *what*
: Dumps the state of a large number of internal state values that might be useful in determining the sequence of events. The *what* parameter can be set to **all**, indicating that all available data should be collected, or to another value, indicating more restricted collection of data. The output is presented to STDOUT and should be collected by redirecting STDOUT.

showtrace
: Shows the current level for each subclass of tracing available in GPFS. Trace level 14 provides the highest level of tracing for the class and trace level 0 provides no tracing. Intermediate values exist for most classes. More tracing requires more storage and results in a higher probability of overlaying the required event.

trace *class n*
: Sets the trace class to the value specified by *n*. Actual trace gathering only occurs when the AIX **trace** command or Linux **lxtrace** command has been issued.

Other options provide interactive GPFS debugging, but are not described here. Output from the **mmfsadm** command will be required in almost all cases where a GPFS problem is being reported. The **mmfsadm**

command collects data only on the node where it is issued. Depending on the nature of the problem, **mmfsadm** output might be required from several or all nodes. The **mmfsadm** output from the file system manager is often required.

To determine where the file system manager is, issue the **mmlsmgr** command:

```
mmlsmgr
```

Output similar to this example is displayed:

```
file system      manager node
---------------- ------------------
fs3              9.114.94.65 (c154n01)
fs2              9.114.94.73 (c154n09)
fs1              9.114.94.81 (c155n01)

Cluster manager node: 9.114.94.65 (c154n01)
```

The GPFS trace facility

GPFS includes many different trace points to facilitate rapid problem determination of failures.

GPFS tracing is based on the AIX kernel trace facility on AIX, embedded GPFS trace subsystem on Linux, and the Windows ETL subsystem on Windows. The level of detail that is gathered by the trace facility is controlled by setting the trace levels using the **mmtracectl** command.

The **mmtracectl** command sets up and enables tracing using default settings for various common problem situations. Using this command improves the probability of gathering accurate and reliable problem determination information. For more information about the **mmtracectl** command, see the *GPFS: Administration and Programming Reference*.

Generating GPFS trace reports

Use the **mmtracectl** command to configure trace-related configuration variables and start and stop the trace facility on any range of nodes in the GPFS cluster.

To configure and use the trace properly:

1. Verify that a directory for dumps was created when the cluster was configured. Issue the **mmlsconfig** command and look for the **dataStructureDump** entry. The default location for trace and problem determination data is **/tmp/mmfs**. Use **mmtracectl** as instructed by service personnel to set trace configuration parameters as required if the default parameters are insufficient. For example, if the problem results in GPFS shutting down, set the **traceRecyle** variable with **--trace-recycle** as described in the **mmtracectl** command in order to ensure that GPFS traces are performed at the time the error occurs.

   If the **dataStructureDump** entry does not exist, issue these commands:

   a. If desired, specify another location for trace and problem determination data by issuing this command:

      ```
      mmchconfig dataStructureDump=path for storage of dumps
      ```

2. To start the tracing facility on all nodes, issue this command:

   ```
   mmtracectl --start
   ```

3. Re-create the problem.
4. When the event to be captured occurs, stop the trace as soon as possible by issuing this command:

   ```
   mmtracectl --stop
   ```

5. The output of the GPFS trace facility is stored in **/tmp/mmfs**, unless the location was changed using the command in Step 1a. Save this output.

6. If the problem results in a shutdown and restart of the GPFS daemon, set the **traceRecycle** variable as necessary to start tracing automatically on daemon startup and stop the trace automatically on daemon shutdown.

If the problem requires more detailed tracing, the IBM Support Center personnel might ask you to modify the GPFS trace levels. Use the **mmtracectl** command to establish the required trace classes and levels of tracing. For example:

```
mmtracectl --set --trace=def
```

Once the trace levels are established, start the tracing by issuing:

```
mmtracectl --start
```

After the trace data has been gathered, stop the tracing by issuing:

```
mmtracectl --stop
```

To clear the trace settings and make sure tracing is turned off, issue:

```
mmtracectl --off
```

Possible trace values include:

trace_class
: Can be any of these terms:

  alloc disk space allocation

  allocmgr
  : allocation manager

  basic 'basic' classes

  block block operations

  brl byte range locks

  cleanup
  : cleanup routines

  cmd ts commands

  defrag
  : defragmentation

  dentry
  : dentry operations

  dfs DFS™ export

  disk physical disk I/O

  disklease
  : disk lease

  dmapi Data Management API

  ds data shipping

  errlog error logging

  eventsExporter
  : events exporter

  file file operations

  fs file system

fsck online multinode fsck

io physical I/O

kernel kernel operations

klockl low-level vfs locking

ksvfs generic kernel vfs information

lock interprocess locking

log recovery log

malloc
malloc and free in shared segment

mb mailbox message handling

mmpmon
mmpmon command

mnode
mnode operations

msg call to routines in SharkMsg.h

mutex mutexes and condition variables

nsd network shared disk

perfmon
performance monitors

pgalloc
page allocator tracing

pin pinning to real memory

pit parallel inode tracing

quota quota management

shared
shared segments

smb SMB locks

sp SP™ message handling

tasking
tasking system but not Thread operations

thread
operations in Thread class

tm token manager

ts daemon specific code

vnode vnode layer of VFS kernel support

vnop one line per VNOP with all important information

trace_level
A value from 0 through 14, which represents an increasing level of detail. A value of 0 turns tracing off.

To display the trace level in use, issue the **mmfsadm showtrace** command.

Using ActivePython for Windows tracing

Windows GPFS diagnostic tracing tools use ActivePython, which must be the Windows (x86) build (not the x64 build).

Tracing on Windows is controlled as it is on the UNIX platforms, by using the **mmtracectl** command. Unless you change the location of GPFS diagnostic output, the trace files are located in the **/tmp/mmfs** directory. The naming convention for the trace files follows the same pattern used on the UNIX platforms. The current implementation of tracing support on Windows for **mmtracectl** makes use of the **gw.py** script to format the trace files. Similarly, the **gw.py** script is used to format any Windows minidump files (core files) produced by GPFS. Since **gw.py** is a Python script this requires that Python be installed on your system for the **gw.py** script to properly format trace and minidump files. Like the trace files the minidump files are placed in the **/tmp/mmfs** directory by default.

Note: For trace formatting to work correctly, you must have Microsoft Windows Driver Kit (WDK) installed, and have the **DDKPATH** environment variable set (for example, **C:\WINDDK\6001.18000**). If this is not set then the **gw.py** script will not be able to find the **tracefmt.exe** command.

Chapter 2. GPFS cluster state information

There are a number of GPFS commands used to obtain cluster state information.

The **mmgetstate** command obtains the state of the GPFS daemon on one or more nodes. The **mmlscluster** and **mmlsconfig** commands display detailed information about the GPFS cluster configuration. The **mmrefresh** command places the most recent GPFS cluster configuration information on the specified nodes. The **mmsdrrestore** command restores GPFS cluster configuration information to the node on which it is invoked. The **mmexpelnode** command instructs the cluster manager to expel the target node.

Note: Both **mmrefresh** and **mmsdrrestore** should be used only by experienced system administrators who are familiar with collecting data and running debugging routines.

For more information, see these commands:

The mmgetstate command

Use the **mmgetstate** command to display the state of the GPFS daemon on one or more nodes.

The **mmgetstate** command is fully described in the *Commands* topic in the *General Parallel File System: Administration and Programming Reference*.

These flags are of interest for problem determination:

-a List all nodes in the GPFS cluster. The option does not display information for nodes that cannot be reached. You may obtain more information if you specify the **-v** option.

-L Additionally display quorum, number of nodes up, and total number of nodes.

The total number of nodes may sometimes be larger than the actual number of nodes in the cluster. This is the case when nodes from other clusters have established connections for the purposes of mounting a file system that belongs to your cluster.

-s Display summary information: number of local and remote nodes that have joined in the cluster, number of quorum nodes, and so forth.

-v Display intermediate error messages.

The remaining flags have the same meaning as in the **mmshutdown** command. They can be used to specify the nodes on which to get the state of the GPFS daemon.

The GPFS *states* recognized and displayed by this command are:

active GPFS is ready for operations.

arbitrating
A node is trying to form quorum with the other available nodes.

down GPFS daemon is not running on the node.

unknown
Unknown value. Node cannot be reached or some other error occurred.

For example, to display the quorum, the number of nodes up, and the total number of nodes, issue:

```
mmgetstate -L -a
```

The system displays output similar to:

```
Node number Node name Quorum Nodes up Total nodes GPFS state Remarks
------------------------------------------------------------------
  2         k154n06     1*     3        7         active     quorum node
  3         k155n05     1*     3        7         active     quorum node
  4         k155n06     1*     3        7         active     quorum node
  5         k155n07     1*     3        7         active
  6         k155n08     1*     3        7         active
  9         k156lnx02   1*     3        7         active
 11         k155n09     1*     3        7         active
```

where *, if present, indicates that tiebreaker disks are being used.

The mmlscluster command

Use the **mmlscluster** command to display GPFS cluster configuration information.

The **mmlscluster** command is fully described in the *Commands* topic in the: *General Parallel File System: Administration and Programming Reference*. The syntax of the command is:

```
mmlscluster
```

The system displays output similar to:

```
GPFS cluster information
========================
  GPFS cluster name:         cluster1.kgn.ibm.com
  GPFS cluster id:           680681562214606028
  GPFS UID domain:           cluster1.kgn.ibm.com
  Remote shell command:      /usr/bin/rsh
  Remote file copy command:  /usr/bin/rcp

GPFS cluster configuration servers:
-----------------------------------
  Primary server:    k164n06.kgn.ibm.com
  Secondary server:  k164n05.kgn.ibm.com

 Node  Daemon node name        IP address      Admin node name       Designation
-----------------------------------------------------------------------------------
   1   k164n04.kgn.ibm.com     198.117.68.68   k164n04.kgn.ibm.com   quorum
   2   k164n05.kgn.ibm.com     198.117.68.71   k164n05.kgn.ibm.com   quorum
   3   k164n06.kgn.ibm.com     198.117.68.70   k164n06.kgn.ibm.com   quorum-manager
```

The mmlsconfig command

Use the **mmlsconfig** command to display current configuration data for a GPFS cluster.

Depending on your configuration, additional information not documented in either the **mmcrcluster** command or the **mmchconfig** command may be displayed to assist in problem determination. The **mmlsconfig** command is fully described in the *Commands* topic in the *General Parallel File System: Administration and Programming Reference*.

If a configuration parameter is not shown in the output of this command, the default value for that parameter, as documented in the **mmchconfig** command, is in effect.

The syntax of the **mmlsconfig** command is:

```
mmlsconfig
```

The system displays information similar to:

```
Configuration data for cluster dream.cluster:
---------------------------------------------
clusterName dream.cluster
clusterId 680752107138921233
clusterType lc
autoload no
minReleaseLevel 3.2.0.0
pagepool 300m
maxblocksize 4m
[c5n97g]
pagepool 3500m
[common]
cipherList EXP-RC4-MD5

File systems in cluster dream.cluster:
--------------------------------------
/dev/fs2
```

The mmrefresh command

The **mmrefresh** command is intended for use by experienced system administrators who know how to collect data and run debugging routines.

Use the **mmrefresh** command only when you suspect that something is not working as expected and the reason for the malfunction is a problem with the GPFS configuration data. For example, a **mount** command fails with a 'device not found' error, and you know that the file system exists. Another example is if any of the files in the **/var/mmfs/gen** directory were accidentally erased. Under normal circumstances, the GPFS command infrastructure maintains the cluster data files automatically and there is no need for user intervention. The **mmrefresh** command is fully described in the *Commands* topic in the *General Parallel File System: Administration and Programming Reference*.

The **mmrefresh** command places the most recent GPFS cluster configuration data files on the specified nodes. The syntax of this command is:

```
mmrefresh [-f] [ -a | -N {Node[,Node...] | NodeFile | NodeClass}]
```

The **-f** flag can be used to force the GPFS cluster configuration data files to be rebuilt whether they appear to be at the most current level or not. If no other option is specified, the command affects only the node on which it is run. The remaining flags have the same meaning as in the **mmshutdown** command, and are used to specify the nodes on which the refresh is to be performed.

For example, to place the GPFS cluster configuration data files at the latest level, on all nodes in the cluster, issue:

```
mmrefresh -a
```

The mmsdrrestore command

The **mmsdrrestore** command is intended for use by experienced system administrators.

The **mmsdrrestore** command restores the latest GPFS system files on the specified nodes. If no nodes are specified, the command restores the configuration information only on the node where it is invoked. If the local GPFS configuration file is missing, the file specified with the **-F** option from the node specified with the **-p** option is used instead. This command works best when used in conjunction with the **mmsdrbackup** user exit, which is described in the *GPFS user exits* topic in the *General Parallel File System: Administration and Programming Reference*.

The **mmsdrrestore** command has this syntax:

mmsdrrestore [**-p** *NodeName*] [**-F** *mmsdrfsFile*] [**-R** *remoteFileCopyCommand*] [**-a** | **-N** {*Node*[,*Node...*] | *NodeFile* | *NodeClass*}]

The flags used by this command are:

-a Restores the GPFS configuration files on all nodes in the cluster.

-F *mmsdrfsFile*
Specifies the path name of the GPFS configuration file for the **mmsdrrestore** command to use. This configuration file could be the current one on the primary server, or it could be a configuration file obtained by means of the **mmsdrbackup** user exit. If not specified, **/var/mmfs/gen/mmsdrfs** is used.

-N {*Node***[***,Node...***] | *NodeFile* | *NodeClass***}**
Restores the GPFS configuration files on a set of nodes. For general information on how to specify node names, see the *Specifying nodes as input to GPFS commands* topic in the *General Parallel File System: Administration and Programming Reference*. This command does not support a *NodeClass* of mount.

-p *NodeName*
Specifies the node from which to obtain a valid GPFS configuration file. This should be either the primary configuration server or a node that has a valid backup copy of the **mmsdrfs** file. If not specified, the local node is used.

-R *remoteFileCopyCommand*
Specifies the fully-qualified path name for the remote file copy program to be used for obtaining the GPFS configuration file. The default is **/usr/bin/rcp**.

Examples of the mmsdrrestore command

1. To restore the latest GPFS system files on the local node using the GPFS configuration file **/var/mmfs/gen/mmsdrfs** from the node named **primaryServer**, issue the following command:

   ```
   mmsdrrestore -p primaryServer
   ```

   The system displays output similar to:

   ```
   Tue Jul  8 18:19:53 CDT 2008: mmsdrrestore: Processing node  k164n04.kgn.ibm.com
   mmsdrrestore:  Node k164n04.kgn.ibm.com successfully restored.
   ```

2. To restore the GPFS system files on all nodes in the cluster using GPFS configuration file **/GPFSconfigFiles/mmsdrfs.120605** on the node named **GPFSarchive**, issue the following command from the node named **localNode**:

   ```
   mmsdrrestore -p GPFSarchive -F /GPFSconfigFiles/mmsdrfs.120605 -a
   ```

   The system displays output similar to:

   ```
   Tue Jul  8 18:29:28 CDT 2008: mmsdrrestore: Processing node k164n04.kgn.ibm.com
   Tue Jul  8 18:29:30 CDT 2008: mmsdrrestore: Processing node k164n05.kgn.ibm.com
   Tue Jul  8 18:29:31 CDT 2008: mmsdrrestore: Processing node k164n06.kgn.ibm.com
   mmsdrrestore: Command successfully completed
   ```

The mmexpelnode command

Use the **mmexpelnode** command to instruct the cluster manager to expel the target nodes and to run the normal recovery protocol.

The cluster manager keeps a list of the expelled nodes. Expelled nodes will not be allowed to rejoin the cluster until they are removed from the list using the **-r** or **--reset** option on the **mmexpelnode** command. The expelled nodes information will also be reset if the cluster manager node goes down or is changed with **mmchmgr -c**.

The **mmexpelnode** command has this syntax:

mmexpelnode [**-o** | **--once**] [**-f** | **--is-fenced**] **-N** *Node*[,*Node*...]

Or,

mmexpelnode {**-r** | **--reset**} **-N** *Node*[,*Node*...]

Or,

mmexpelnode {**-l** | **--list**}

The flags used by this command are:

-f | --is-fenced
: Specifies that the nodes are fenced out and precluded from accessing any GPFS disks without first rejoining the cluster (for example, the nodes were forced to reboot by turning off power). This allows GPFS to run faster recovery protocol.

-l | --list
: Lists all currently expelled nodes.

-N *Node***[,***Node***...]**
: Specifies a list of host names or IP addresses representing the nodes to be expelled. Specify the daemon interface host names or IP addresses as shown by the **mmlscluster** command. The **mmexpelnode** command does not support administration node names or node classes.

-o | --once
: Specifies that the nodes should not be prevented from rejoining the cluster after the recovery protocol completes.

-r | --reset
: Allows the specified nodes to rejoin the cluster (reset node status).

Examples of the mmexpelnode command

1. To expel node c100c1rp3, issue the command:

   ```
   mmexpelnode -N c100c1rp3
   ```

2. To show a list of expelled nodes, issue the command:

   ```
   mmexpelnode --list
   ```

   The system displays information similar to:

   ```
   Node List
   ---------------------
       192.168.100.35 (c100c1rp3.ppd.pok.ibm.com)
   ```

3. To allow node c100c1rp3 to rejoin the cluster, issue the command:

   ```
   mmexpelnode -r -N c100c1rp3
   ```

Chapter 3. GPFS file system and disk information

The problem determination tools provided with GPFS for file system, disk and NSD problem determination are intended for use by experienced system administrators who know how to collect data and run debugging routines.

The information is organized as follows:

Restricted mode mount

GPFS provides a capability to mount a file system in a restricted mode when significant data structures have been destroyed by disk failures or other error conditions.

Restricted mode mount is not intended for normal operation, but may allow the recovery of some user data. Only data which is referenced by intact directories and metadata structures would be available.

Attention:

1. Consult the IBM Support Center before using this capability.
2. Attempt this only after you have tried to repair the file system with the **mmfsck** command.
3. Use this procedure only if the failing disk is attached to an AIX or Linux node.

Some disk failures can result in the loss of enough metadata to render the entire file system unable to mount. In that event it might be possible to preserve some user data through a *restricted mode mount*. This facility should only be used if a normal mount does not succeed, and should be considered a last resort to save some data after a fatal disk failure.

| Restricted mode mount is invoked by using the **mmmount** command with the **-o rs** flags. After a restricted mode mount is done, some data *may* be sufficiently accessible to allow copying to another file system. The success of this technique depends on the actual disk structures damaged.

Read-only mode mount

Some disk failures can result in the loss of enough metadata to make the entire file system unable to mount. In that event, it might be possible to preserve some user data through a *read-only mode mount*.

Attention: Attempt this only after you have tried to repair the file system with the **mmfsck** command.

| This facility should be used only if a normal mount does not succeed, and should be considered a last resort to save some data after a fatal disk failure.

Read-only mode mount is invoked by using the **mmmount** command with the **-o ro** flags. After a read-only mode mount is done, some data *may* be sufficiently accessible to allow copying to another file system. The success of this technique depends on the actual disk structures damaged.

The lsof command

Use the **lsof** (list open files) command to return the user processes that are actively using a file system. It is sometimes helpful in determining why a file system remains in use and cannot be unmounted.

The **lsof** command is available using anonymous ftp from **lsof.itap.purdue.edu** (**cd** to **/pub/tools/unix/lsof**). The inventor of the **lsof** command is Victor A. Abell (abe@purdue.edu), Purdue University Computing Center.

The mmlsmount command

Use the **mmlsmount** command to see which nodes have the file system in use.

Use the **-L** option to see the node name and IP address of each node that has the file system in use. This command can be used for all file systems, all remotely mounted file systems, or file systems mounted on nodes of certain clusters.

The **mmlsmount** command is fully described in the *Commands* unit of: *General Parallel File System: Administration and Programming Reference*.

While not specifically intended as a service aid, the **mmlsmount** command is useful in these situations:

1. When writing and debugging new file system administrative procedures, to determine which nodes have a file system mounted and which do not.
2. When mounting a file system on multiple nodes, to determine which nodes have successfully completed the mount and which have not.
3. When a file system is mounted, but appears to be inaccessible to some nodes but accessible to others, to determine the extent of the problem.
4. When a normal (not force) unmount has not completed, to determine the affected nodes.
5. When a file system has force unmounted on some nodes but not others, to determine the affected nodes.

For example, to list the nodes having all file systems mounted:

```
mmlsmount all -L
```

Output is similar to this:

```
File system fs1 is mounted on 7 nodes:
119.124.94.69 k154n05 dq.cluster
119.124.94.87 k155n07 dq.cluster
119.124.94.102 k156n06 remote.cluster
119.124.94.82 k155n02 remote.cluster
119.124.94.86 k155n06 dq.cluster
119.124.94.89 k155n09 dq.cluster
119.124.94.81 k155n01 remote.cluster

File system remotefs1 (remote.cluster:fs1)is mounted on 4 nodes:
119.124.94.115 kolt1_r1b42 remote.cluster
119.124.94.65 k154n01 remote.cluster
119.124.94.69 k154n05 dq.cluster
119.124.94.73 k154n09 remote.cluster
```

The mmapplypolicy -L command

Use the **-L** flag of the **mmapplypolicy** command when you are using policy files to manage storage resources and the data stored on those resources. This command has different levels of diagnostics to help debug and interpret the actions of a policy file.

The **-L** flag, used in conjunction with the **-I test** flag, allows you to display the actions that would be performed by a policy file without actually applying it. This way, potential errors and misunderstandings can be detected and corrected without actually making these mistakes.

These are the trace levels for the **mmapplypolicy -L** flag:

Value	Description
0	Displays only serious errors.
1	Displays some information as the command runs, but not for each file.
2	Displays each chosen file and the scheduled migration or deletion action.
3	All of the above, plus displays each candidate file and the applicable rule.
4	All of the above, plus displays each explicitly excluded file, and the applicable rule.
5	All of the above, plus displays the attributes of candidate and excluded files.
6	All of the above, plus displays files that are not candidate files, and their attributes.

These terms are used:

candidate file
A file that matches a **MIGRATE** or **DELETE** policy rule.

chosen file
A candidate file that has been scheduled for migration or deletion.

This policy file is used in the examples that follow:

```
/* Exclusion rule */
RULE 'exclude *.save files' EXCLUDE WHERE NAME LIKE '%.save'

/* Deletion rule */
RULE 'delete' DELETE FROM POOL 'sp1' WHERE NAME LIKE '%tmp%'

/* Migration rule */
RULE 'migration to system pool' MIGRATE FROM POOL 'sp1' TO POOL 'system' WHERE NAME LIKE '%file%'

/*  Typo in rule : removed later */
RULE 'exclude 2' EXCULDE
```

These are some of the files in file system **/fs1**:

```
.  ..  data1  file.tmp0  file.tmp1  file0  file1  file1.save  file2.save
```

mmapplypolicy -L 0

Use this option to display only serious errors.

In this example, there is an error in the policy file. This command:

```
mmapplypolicy fs1 -P policyfile -I test -L 0
```

produces output similar to this:

```
Error while loading policy rules from file policyfile: code -1
PCSQLERR: Unexpected SQL identifier token - 'EXCULDE'.
PCSQLCTX: at line 16 of 16: RULE 'exclude 2 ' {{{EXCULDE}}}
```

The error in the policy file is corrected by removing these lines:

```
/* Typo in rule */
RULE 'exclude 2' EXCULDE
```

Now rerun the command:

```
mmapplypolicy fs1 -P policyfile -I test -L 0
```

No messages are produced because no serious errors were detected.

mmapplypolicy -L 1

Use this option to display all of the information (if any) from the previous level, plus some information as the command runs, but not for each file.

Total numbers for file migration and deletion are displayed, for example.

This command:

```
mmapplypolicy fs1 -P policyfile -I test -L 1
```

produces output similar to this:

```
GPFS Current Data Pool Utilization in KB and %
sp1     2560    9765632 0.026214%
system  154880  9765632 1.585970%
Loaded policy rules from policyfile.
Evaluating MIGRATE/DELETE/EXCLUDE rules with CURRENT_TIMESTAMP = 2007-06-27@14:38:40 UTC
parsed 0 Placement Rules, 0 Restore Rules, 3 Migrate/Delete/Exclude Rules
/* Exclusion rule */
RULE 'exclude *.save files' EXCLUDE WHERE NAME LIKE '%.save'

/* Deletion rule */
RULE 'delete' DELETE FROM POOL 'sp1' WHERE NAME LIKE '%tmp%'

/* Migration rule */
RULE 'migration to system pool' MIGRATE FROM POOL 'sp1' TO POOL 'system' WHERE NAME LIKE '%file%'

Directories scan: 7 files, 1 directories, 0 other objects, 0 'skipped' files and/or errors.
Inodes scan: 7 files, 1 directories, 0 other objects, 0 'skipped' files and/or errors.
Summary of Rule Applicability and File Choices:
 Rule#  Hit_Cnt Chosen  KB_Chosen       KB_Ill  Rule
  0     2       0       0       0       RULE 'exclude *.save files' EXCLUDE WHERE(.)
  1     2       2       16      0       RULE 'delete' DELETE FROM POOL 'sp1' WHERE(.)
  2     2       2       16      0       RULE 'migration to system pool' MIGRATE FROM POOL
                                        'sp1' TO POOL 'system' WHERE(.)

Files with no applicable rules: 2.

GPFS Policy Decisions and File Choice Totals:
 Chose to migrate 16KB: 2 of 2 candidates;
 Chose to premigrate 0KB: 0 candidates;
 Already co-managed 0KB: 0 candidates;
 Chose to delete 16KB: 2 of 2 candidates;
 0KB of chosen data is illplaced or illreplicated;
Predicted Data Pool Utilization in KB and %:
sp1     2528    9765632 0.025887%
system  154896  9765632 1.586134%
```

mmapplypolicy -L 2

Use this option to display all of the information from the previous levels, plus each chosen file and the scheduled migration or deletion action.

This command:

```
mmapplypolicy fs1 -P policyfile -I test -L 2
```

produces output similar to this:

```
GPFS Current Data Pool Utilization in KB and %
sp1     2560    9765632 0.026214%
system  154624  9765632 1.583349%
Loaded policy rules from policyfile.
Evaluating MIGRATE/DELETE/EXCLUDE rules with CURRENT_TIMESTAMP = 2007-06-27@14:50:22 UTC
parsed 0 Placement Rules, 0 Restore Rules, 3 Migrate/Delete/Exclude Rules
/* Exclusion rule */
RULE 'exclude *.save files' EXCLUDE WHERE NAME LIKE '%.save'

/* Deletion rule */
RULE 'delete' DELETE FROM POOL 'sp1' WHERE NAME LIKE '%tmp%'

/* Migration rule */
RULE 'migration to system pool' MIGRATE FROM POOL 'sp1' TO POOL 'system' WHERE NAME LIKE '%file%'

Directories scan: 7 files, 1 directories, 0 other objects, 0 'skipped' files and/or errors.
Inodes scan: 7 files, 1 directories, 0 other objects, 0 'skipped' files and/or errors.
WEIGHT(inf) MIGRATE /fs1/file1 TO POOL system
WEIGHT(inf) MIGRATE /fs1/file0 TO POOL system
WEIGHT(inf) DELETE /fs1/file.tmp0
WEIGHT(inf) DELETE /fs1/file.tmp1
Summary of Rule Applicability and File Choices:
 Rule#  Hit_Cnt Chosen  KB_Chosen       KB_Ill  Rule
  0     2       0       0       0       RULE 'exclude *.save files' EXCLUDE WHERE(.)
  1     2       2       16      0       RULE 'delete' DELETE FROM POOL 'sp1' WHERE(.)
  2     2       2       16      0       RULE 'migration to system pool' MIGRATE FROM POOL
                                        'sp1' TO POOL 'system' WHERE(.)

Files with no applicable rules: 2.

GPFS Policy Decisions and File Choice Totals:
 Chose to migrate 16KB: 2 of 2 candidates;
 Chose to premigrate 0KB: 0 candidates;
 Already co-managed 0KB: 0 candidates;
 Chose to delete 16KB: 2 of 2 candidates;
 0KB of chosen data is illplaced or illreplicated;
Predicted Data Pool Utilization in KB and %:
sp1     2528    9765632 0.025887%
system  154640  9765632 1.583512%
```

The lines:

```
WEIGHT(inf) MIGRATE /fs1/file0 TO POOL system
WEIGHT(inf) MIGRATE /fs1/file1 TO POOL system
WEIGHT(inf) DELETE /fs1/file.tmp0
WEIGHT(inf) DELETE /fs1/file.tmp1
```

show the four chosen files and the scheduled migrate or delete action.

mmapplypolicy -L 3

Use this option to display all of the information from the previous levels, plus each candidate file and the applicable rule.

This command:

```
mmapplypolicy fs1-P policyfile -I test -L 3
```

produces output similar to this:

```
GPFS Current Data Pool Utilization in KB and %
sp1     2560    9765632 0.026214%
system  154624  9765632 1.583349%
```

```
Loaded policy rules from policyfile.
Evaluating MIGRATE/DELETE/EXCLUDE rules with CURRENT_TIMESTAMP = 2007-06-27@14:50:56 UTC
parsed 0 Placement Rules, 0 Restore Rules, 3 Migrate/Delete/Exclude Rules
/* Exclusion rule */
RULE 'exclude *.save files' EXCLUDE WHERE NAME LIKE '%.save'

/* Deletion rule */
RULE 'delete' DELETE FROM POOL 'sp1' WHERE NAME LIKE '%tmp%'

/* Migration rule */
RULE 'migration to system pool' MIGRATE FROM POOL 'sp1' TO POOL 'system' WHERE NAME LIKE '%file%'

Directories scan: 7 files, 1 directories, 0 other objects, 0 'skipped' files and/or errors.
/fs1/file1       RULE 'migration to system pool' MIGRATE FROM POOL 'sp1' TO POOL 'system' WEIGHT(inf)
/fs1/file0       RULE 'migration to system pool' MIGRATE FROM POOL 'sp1' TO POOL 'system' WEIGHT(inf)
/fs1/file.tmp0   RULE 'delete' DELETE FROM POOL 'sp1' WEIGHT(inf)
/fs1/file.tmp1   RULE 'delete' DELETE FROM POOL 'sp1' WEIGHT(inf)
Inodes scan: 7 files, 1 directories, 0 other objects, 0 'skipped' files and/or errors.
WEIGHT(inf) MIGRATE /fs1/file1 TO POOL system
WEIGHT(inf) MIGRATE /fs1/file0 TO POOL system
WEIGHT(inf) DELETE /fs1/file.tmp0
WEIGHT(inf) DELETE /fs1/file.tmp1
Summary of Rule Applicability and File Choices:
 Rule#  Hit_Cnt Chosen  KB_Chosen       KB_Ill  Rule
  0     2       0       0       0       RULE 'exclude *.save files' EXCLUDE WHERE(.)
  1     2       2       16      0       RULE 'delete' DELETE FROM POOL 'sp1' WHERE(.)
  2     2       2       16      0       RULE 'migration to system pool' MIGRATE FROM POOL
                                        'sp1' TO POOL 'system' WHERE(.)

Files with no applicable rules: 2.

GPFS Policy Decisions and File Choice Totals:
 Chose to migrate 16KB: 2 of 2 candidates;
 Chose to premigrate 0KB: 0 candidates;
 Already co-managed 0KB: 0 candidates;
 Chose to delete 16KB: 2 of 2 candidates;
 0KB of chosen data is illplaced or illreplicated;
Predicted Data Pool Utilization in KB and %:
sp1     2528    9765632 0.025887%
system  154640  9765632 1.583512%
```

The lines:

```
/fs1/file0       RULE 'migration to system pool' MIGRATE FROM POOL 'sp1' TO POOL 'system' WEIGHT(inf)
/fs1/file1       RULE 'migration to system pool' MIGRATE FROM POOL 'sp1' TO POOL 'system' WEIGHT(inf)
/fs1/file.tmp0   RULE 'delete' DELETE FROM POOL 'sp1' WEIGHT(inf)
/fs1/file.tmp1   RULE 'delete' DELETE FROM POOL 'sp1' WEIGHT(inf)
```

show the four candidate files, and the applicable rules.

mmapplypolicy -L 4

Use this option to display all of the information from the previous levels, plus the name of each explicitly excluded file, and the applicable rule.

This command:

```
mmapplypolicy fs1 -P policyfile -I test -L 4
```

produces output similar to this:

```
GPFS Current Data Pool Utilization in KB and %
sp1     2560    9765632 0.026214%
system  154624  9765632 1.583349%
Loaded policy rules from policyfile.
Evaluating MIGRATE/DELETE/EXCLUDE rules with CURRENT_TIMESTAMP = 2007-06-27@14:53:17 UTC
parsed 0 Placement Rules, 0 Restore Rules, 3 Migrate/Delete/Exclude Rules
```

```
/* Exclusion rule */
RULE 'exclude *.save files' EXCLUDE WHERE NAME LIKE '%.save'

/* Deletion rule */
RULE 'delete' DELETE FROM POOL 'sp1' WHERE NAME LIKE '%tmp%'

/* Migration rule */
RULE 'migration to system pool' MIGRATE FROM POOL 'sp1' TO POOL 'system' WHERE NAME LIKE '%file%'

Directories scan: 7 files, 1 directories, 0 other objects, 0 'skipped' files and/or errors.
/fs1/file1       RULE 'migration to system pool' MIGRATE FROM POOL 'sp1' TO POOL 'system' WEIGHT(inf)
/fs1/file1.save  RULE 'exclude *.save files' EXCLUDE
/fs1/file0       RULE 'migration to system pool' MIGRATE FROM POOL 'sp1' TO POOL 'system' WEIGHT(inf)
/fs1/file2.save  RULE 'exclude *.save files' EXCLUDE
/fs1/file.tmp0   RULE 'delete' DELETE FROM POOL 'sp1' WEIGHT(inf)
/fs1/file.tmp1   RULE 'delete' DELETE FROM POOL 'sp1' WEIGHT(inf)
Inodes scan: 7 files, 1 directories, 0 other objects, 0 'skipped' files and/or errors.
WEIGHT(inf) MIGRATE /fs1/file1 TO POOL system
WEIGHT(inf) MIGRATE /fs1/file0 TO POOL system
WEIGHT(inf) DELETE /fs1/file.tmp0
WEIGHT(inf) DELETE /fs1/file.tmp1
Summary of Rule Applicability and File Choices:
 Rule#  Hit_Cnt Chosen  KB_Chosen       KB_Ill  Rule
  0     2       0       0       0       RULE 'exclude *.save files' EXCLUDE WHERE(.)
  1     2       2       16      0       RULE 'delete' DELETE FROM POOL 'sp1' WHERE(.)
  2     2       2       16      0       RULE 'migration to system pool' MIGRATE FROM POOL
                                        'sp1' TO POOL 'system' WHERE(.)

Files with no applicable rules: 2.

GPFS Policy Decisions and File Choice Totals:
 Chose to migrate 16KB: 2 of 2 candidates;
 Chose to premigrate 0KB: 0 candidates;
 Already co-managed 0KB: 0 candidates;
 Chose to delete 16KB: 2 of 2 candidates;
 0KB of chosen data is illplaced or illreplicated;
Predicted Data Pool Utilization in KB and %:
sp1     2528    9765632 0.025887%
system  154640  9765632 1.583512%
```

The lines:

```
/fs1/file1.save  RULE 'exclude *.save files' EXCLUDE
/fs1/file2.save  RULE 'exclude *.save files' EXCLUDE
```

indicate that there are two excluded files, **/fs1/file1.save** and **/fs1/file2.save**.

mmapplypolicy -L 5

Use this option to display all of the information from the previous levels, plus the attributes of candidate and excluded files.

These attributes include:

- **MODIFICATION_TIME**
- **USER_ID**
- **GROUP_ID**
- **FILE_SIZE**
- **POOL_NAME**
- **ACCESS_TIME**
- **KB_ALLOCATED**
- **FILESET_NAME**

This command:

```
mmapplypolicy fs1 -P policyfile -I test -L 5
```

produces output similar to this:

```
GPFS Current Data Pool Utilization in KB and %
sp1     2560    9765632 0.026214%
system  154624  9765632 1.583349%
Loaded policy rules from policyfile.
Evaluating MIGRATE/DELETE/EXCLUDE rules with CURRENT_TIMESTAMP = 2007-06-27@14:56:21 UTC
parsed 0 Placement Rules, 0 Restore Rules, 3 Migrate/Delete/Exclude Rules
/* Exclusion rule */
RULE 'exclude *.save files' EXCLUDE WHERE NAME LIKE '%.save'

/* Deletion rule */
RULE 'delete' DELETE FROM POOL 'sp1' WHERE NAME LIKE '%tmp%'

/* Migration rule */
RULE 'migration to system pool' MIGRATE FROM POOL 'sp1' TO POOL 'system' WHERE NAME LIKE '%file%'

Directories scan: 7 files, 1 directories, 0 other objects, 0 'skipped' files and/or errors.
/fs1/file1      [2007-06-27@14:30:54 0 0 29 sp1 2007-06-27@14:32:19 8 root] RULE 'migration to system pool'
                                                    MIGRATE FROM POOL 'sp1' TO POOL 'system' WEIGHT(inf)
/fs1/file1.save [2007-06-27@14:30:57 0 0 29 sp1 2007-06-27@14:30:57 8 root] RULE 'exclude *.save files'
                                                    EXCLUDE
/fs1/file0      [2007-06-27@14:31:00 0 0 29 sp1 2007-06-27@14:31:00 8 root] RULE 'migration to system pool'
                                                    MIGRATE FROM POOL 'sp1' TO POOL 'system' WEIGHT(inf)
/fs1/file2.save [2007-06-27@14:31:03 0 0 29 sp1 2007-06-27@14:31:03 8 root] RULE 'exclude *.save files'
                                                    EXCLUDE
/fs1/file.tmp0  [2007-06-27@14:32:13 0 0 29 sp1 2007-06-27@14:32:13 8 root] RULE 'delete' DELETE FROM POOL
                                                                            'sp1' WEIGHT(inf)
/fs1/file.tmp1  [2007-06-27@14:32:19 0 0 29 sp1 2007-06-27@14:32:19 8 root] RULE 'delete' DELETE FROM POOL
                                                                            'sp1' WEIGHT(inf)
Inodes scan: 7 files, 1 directories, 0 other objects, 0 'skipped' files and/or errors.
WEIGHT(inf) MIGRATE /fs1/file1 TO POOL system
WEIGHT(inf) MIGRATE /fs1/file0 TO POOL system
WEIGHT(inf) DELETE /fs1/file.tmp0
WEIGHT(inf) DELETE /fs1/file.tmp1
Summary of Rule Applicability and File Choices:
 Rule#  Hit_Cnt Chosen  KB_Chosen       KB_Ill  Rule
  0     2       0       0       0       RULE 'exclude *.save files' EXCLUDE WHERE(.)
  1     2       2       16      0       RULE 'delete' DELETE FROM POOL 'sp1' WHERE(.)
  2     2       2       16      0       RULE 'migration to system pool' MIGRATE FROM POOL
                                        'sp1' TO POOL 'system' WHERE(.)

Files with no applicable rules: 2.

GPFS Policy Decisions and File Choice Totals:
 Chose to migrate 16KB: 2 of 2 candidates;
 Chose to premigrate 0KB: 0 candidates;
 Already co-managed 0KB: 0 candidates;
 Chose to delete 16KB: 2 of 2 candidates;
 0KB of chosen data is illplaced or illreplicated;
Predicted Data Pool Utilization in KB and %:
sp1     2528    9765632 0.025887%
system  154640  9765632 1.583512%
(10:56:21) c109c2rp1:/ #
```

These lines:

```
/fs1/file1.save [2006-03-22@14:46:23 0 0 1024 sp1 2006-03-22@14:46:23 8 root]
RULE 'exclude *.save files' EXCLUDE
/fs1/file2.save [2006-03-22@14:46:23 0 0 1024 sp1 2006-03-22@14:46:23 8 root]
RULE 'exclude *.save files' EXCLUDE
```

give the attributes of the two excluded files **/fs1/file1.save** and **/fs1/file2.save**.

mmapplypolicy -L 6

Use this option to display all of the information from the previous levels, plus files that are not candidate files, and their attributes.

These attributes include:

- **MODIFICATION_TIME**
- **USER_ID**
- **GROUP_ID**
- **FILE_SIZE**
- **POOL_NAME**
- **ACCESS_TIME**
- **KB_ALLOCATED**
- **FILESET_NAME**

This command:

```
mmapplypolicy fs1 -P policyfile -I test -L 6
```

produces output similar to this:

```
GPFS Current Data Pool Utilization in KB and %
sp1     2560    9765632 0.026214%
system  154624  9765632 1.583349%
Loaded policy rules from policyfile.
Evaluating MIGRATE/DELETE/EXCLUDE rules with CURRENT_TIMESTAMP = 2007-06-27@14:57:40 UTC
parsed 0 Placement Rules, 0 Restore Rules, 3 Migrate/Delete/Exclude Rules
/* Exclusion rule */
RULE 'exclude *.save files' EXCLUDE WHERE NAME LIKE '%.save'

/* Deletion rule */
RULE 'delete' DELETE FROM POOL 'sp1' WHERE NAME LIKE '%tmp%'

/* Migration rule */
RULE 'migration to system pool' MIGRATE FROM POOL 'sp1' TO POOL 'system' WHERE NAME LIKE '%file%'

Directories scan: 7 files, 1 directories, 0 other objects, 0 'skipped' files and/or errors.
/fs1/.  [2007-06-27@14:32:36 0 0 8192 system 2007-06-27@14:56:21 8 root] NO RULE APPLIES
/fs1/data1      [2007-06-27@14:30:49 0 0 29 sp1 2007-06-27@14:30:49 8 root] NO RULE APPLIES
/fs1/file1      [2007-06-27@14:30:54 0 0 29 sp1 2007-06-27@14:32:19 8 root] RULE 'migration to system pool'
                                                MIGRATE FROM POOL 'sp1' TO POOL 'system' WEIGHT(inf)
/fs1/file1.save [2007-06-27@14:30:57 0 0 29 sp1 2007-06-27@14:30:57 8 root] RULE 'exclude *.save files'
                                                EXCLUDE
/fs1/file0      [2007-06-27@14:31:00 0 0 29 sp1 2007-06-27@14:31:00 8 root] RULE 'migration to system pool'
                                                MIGRATE FROM POOL 'sp1' TO POOL 'system' WEIGHT(inf)
/fs1/file2.save [2007-06-27@14:31:03 0 0 29 sp1 2007-06-27@14:31:03 8 root] RULE 'exclude *.save files'
                                                EXCLUDE
/fs1/file.tmp0  [2007-06-27@14:32:13 0 0 29 sp1 2007-06-27@14:32:13 8 root] RULE 'delete' DELETE FROM POOL
                                                                            'sp1' WEIGHT(inf)
/fs1/file.tmp1  [2007-06-27@14:32:19 0 0 29 sp1 2007-06-27@14:32:19 8 root] RULE 'delete' DELETE FROM POOL
                                                                            'sp1' WEIGHT(inf)
Inodes scan: 7 files, 1 directories, 0 other objects, 0 'skipped' files and/or errors.
WEIGHT(inf) MIGRATE /fs1/file1 TO POOL system
WEIGHT(inf) MIGRATE /fs1/file0 TO POOL system
WEIGHT(inf) DELETE /fs1/file.tmp0
WEIGHT(inf) DELETE /fs1/file.tmp1
Summary of Rule Applicability and File Choices:
 Rule#  Hit_Cnt Chosen  KB_Chosen       KB_Ill  Rule
  0     2       0       0       0       RULE 'exclude *.save files' EXCLUDE WHERE(.)
  1     2       2       16      0       RULE 'delete' DELETE FROM POOL 'sp1' WHERE(.)
  2     2       2       16      0       RULE 'migration to system pool' MIGRATE FROM POOL
                                        'sp1' TO POOL 'system' WHERE(.)

Files with no applicable rules: 2.

GPFS Policy Decisions and File Choice Totals:
```

```
 Chose to migrate 16KB: 2 of 2 candidates;
 Chose to premigrate 0KB: 0 candidates;
 Already co-managed 0KB: 0 candidates;
 Chose to delete 16KB: 2 of 2 candidates;
 0KB of chosen data is illplaced or illreplicated;
Predicted Data Pool Utilization in KB and %:
sp1     2528    9765632 0.025887%
system  154640  9765632 1.583512%
```

The line:

```
/fs1/data1      [2006-03-22@15:15:27 0 0 29 system 2006-03-22@15:15:15 8 root] NO RULE APPLIES
```

contains information about the one file that is not a candidate file, **data1**.

The mmcheckquota command

Use the **mmcheckquota** command to count inode and space usage for a file system and write the collected data into quota files.

The **mmcheckquota** command is fully described in the *Commands* topic of the *General Parallel File System: Administration and Programming Reference*. Indications leading you to the conclusion you should run the **mmcheckquota** command include:

- **MMFS_QUOTA** error log entries. This error log entry is created when the quota manager has a problem reading or writing the quota file.
- Quota information is lost due to node failure. Node failure could leave users unable to open files or deny them disk space that their quotas should allow.
- The *in doubt* value is approaching the quota limit. The sum of the *in doubt* value and the current usage may not exceed the hard limit. Consequently, the actual block space and number of files available to the user of the group may be constrained by the *in doubt* value. Should the *in doubt* value approach a significant percentage of the quota, use the **mmcheckquota** command to account for the lost space and files.
- User, group, or fileset quota files are corrupted.

During the normal operation of file systems with quotas enabled (not running **mmcheckquota** online), the usage data reflects the actual usage of the blocks and inodes in the sense that if you delete files you should see the usage amount decrease. The *in doubt* value does not reflect how much the user has used already, it is just the amount of quotas that the quota server has assigned to its clients. The quota server does not know whether the assigned amount has been used or not. The only situation where the *in doubt* value is important to the user is when the sum of the usage and the *in doubt* value is greater than the user's quota hard limit. In this case, the user is not allowed to allocate more blocks or inodes unless he brings the usage down.

The mmlsnsd command

Use the **mmlsnsd** command to display information about the currently defined disks in the cluster.

The **mmlsnsd** command is fully described in the *Commands* topic in the *General Parallel File System: Administration and Programming Reference*.

For example, if you issue **mmlsnsd**, your output is similar to this:

```
 File system   Disk name    NSD servers
---------------------------------------------------------------------------
 fs2           hd3n97       c5n97g.ppd.pok.ibm.com,c5n98g.ppd.pok.ibm.com,c5n99g.ppd.pok.ibm.com
 fs2           hd4n97       c5n97g.ppd.pok.ibm.com,c5n98g.ppd.pok.ibm.com,c5n99g.ppd.pok.ibm.com
 fs2           hd5n98       c5n98g.ppd.pok.ibm.com,c5n97g.ppd.pok.ibm.com,c5n99g.ppd.pok.ibm.com
 fs2           hd6n98       c5n98g.ppd.pok.ibm.com,c5n97g.ppd.pok.ibm.com,c5n99g.ppd.pok.ibm.com
 fs2           hd7vsdn97    c5n97g.ppd.pok.ibm.com,c5n98g.ppd.pok.ibm.com,c5n99g.ppd.pok.ibm.com
```

```
fs2            hd8vsdn97    c5n97g.ppd.pok.ibm.com,c5n98g.ppd.pok.ibm.com,c5n99g.ppd.pok.ibm.com
fs2            hd9vsdn97    c5n97g.ppd.pok.ibm.com,c5n98g.ppd.pok.ibm.com,c5n99g.ppd.pok.ibm.com
fs2            hd10vsdn98   c5n98g.ppd.pok.ibm.com,c5n97g.ppd.pok.ibm.com,c5n99g.ppd.pok.ibm.com
fs2            hd11vsdn98   c5n98g.ppd.pok.ibm.com,c5n97g.ppd.pok.ibm.com
fs2            hd12vsdn98   c5n98g.ppd.pok.ibm.com,c5n97g.ppd.pok.ibm.com
fs2            sdbnsd       c5n94g.ppd.pok.ibm.com,c5n96g.ppd.pok.ibm.com
fs2            sdcnsd       c5n94g.ppd.pok.ibm.com,c5n96g.ppd.pok.ibm.com
fs2            sddnsd       c5n94g.ppd.pok.ibm.com,c5n96g.ppd.pok.ibm.com
fs2            sdensd       c5n94g.ppd.pok.ibm.com,c5n96g.ppd.pok.ibm.com
fs2            sdgnsd       c5n94g.ppd.pok.ibm.com,c5n96g.ppd.pok.ibm.com
fs2            sdfnsd       c5n94g.ppd.pok.ibm.com,c5n96g.ppd.pok.ibm.com
fs2            sdhnsd       c5n94g.ppd.pok.ibm.com,c5n96g.ppd.pok.ibm.com
(free disk)    hd2n97       c5n97g.ppd.pok.ibm.com,c5n98g.ppd.pok.ibm.com
```

To find out the local device names for these disks, use the **mmlsnsd** command with the **-m** option. For example, issuing **mmlsnsd -m** produces output similar to this:

```
 Disk name    NSD volume ID      Device         Node name                Remarks
---------------------------------------------------------------------------------------
 hd10vsdn98   0972846245C8E93C   /dev/hd10vsdn98 c5n97g.ppd.pok.ibm.com   server node
 hd10vsdn98   0972846245C8E93C   /dev/hd10vsdn98 c5n98g.ppd.pok.ibm.com   server node
 hd11vsdn98   0972846245C8E93F   /dev/hd11vsdn98 c5n97g.ppd.pok.ibm.com   server node
 hd11vsdn98   0972846245C8E93F   /dev/hd11vsdn98 c5n98g.ppd.pok.ibm.com   server node
 hd12vsdn98   0972846245C8E941   /dev/hd12vsdn98 c5n97g.ppd.pok.ibm.com   server node
 hd12vsdn98   0972846245C8E941   /dev/hd12vsdn98 c5n98g.ppd.pok.ibm.com   server node
 hd2n97       0972846145C8E924   /dev/hdisk2     c5n97g.ppd.pok.ibm.com   server node
 hd2n97       0972846145C8E924   /dev/hdisk2     c5n98g.ppd.pok.ibm.com   server node
 hd3n97       0972846145C8E927   /dev/hdisk3     c5n97g.ppd.pok.ibm.com   server node
 hd3n97       0972846145C8E927   /dev/hdisk3     c5n98g.ppd.pok.ibm.com   server node
 hd4n97       0972846145C8E92A   /dev/hdisk4     c5n97g.ppd.pok.ibm.com   server node
 hd4n97       0972846145C8E92A   /dev/hdisk4     c5n98g.ppd.pok.ibm.com   server node
 hd5n98       0972846245EB501C   /dev/hdisk5     c5n97g.ppd.pok.ibm.com   server node
 hd5n98       0972846245EB501C   /dev/hdisk5     c5n98g.ppd.pok.ibm.com   server node
 hd6n98       0972846245DB3AD8   /dev/hdisk6     c5n97g.ppd.pok.ibm.com   server node
 hd6n98       0972846245DB3AD8   /dev/hdisk6     c5n98g.ppd.pok.ibm.com   server node
 hd7vsdn97    0972846145C8E934   /dev/hd7vsdn97  c5n97g.ppd.pok.ibm.com   server node
```

To obtain extended information for NSDs, use the **mmlsnsd** command with the **-X** option. For example, issuing **mmlsnsd -X** produces output similar to this:

```
 Disk name    NSD volume ID      Device         Devtype  Node name                Remarks
-------------------------------------------------------------------------------------------------
 hd3n97       0972846145C8E927   /dev/hdisk3    hdisk    c5n97g.ppd.pok.ibm.com   server node,pr=no
 hd3n97       0972846145C8E927   /dev/hdisk3    hdisk    c5n98g.ppd.pok.ibm.com   server node,pr=no
 hd5n98       0972846245EB501C   /dev/hdisk5    hdisk    c5n97g.ppd.pok.ibm.com   server node,pr=no
 hd5n98       0972846245EB501C   /dev/hdisk5    hdisk    c5n98g.ppd.pok.ibm.com   server node,pr=no
 sdfnsd       0972845E45F02E81   /dev/sdf       generic  c5n94g.ppd.pok.ibm.com   server node
 sdfnsd       0972845E45F02E81   /dev/sdm       generic  c5n96g.ppd.pok.ibm.com   server node
```

The mmfileid command

Use the **mmfileid** command to determine which files are located on areas of a disk that are damaged or considered to be suspect.

Attention: Use this command only when directed by the IBM Support Center.

Before running **mmfileid**, you must run a disk analysis utility and obtain the disk sector numbers that are suspect or known to be damaged. These sectors are input to the **mmfileid** command.

The syntax is:

mmfileid *Device* **{-d [***NodeName***]:{***DiskName***|***DiskNum***|BROKEN}:[***PhysAddr1***[-***PhysAddr2***]] | -F** *DescFile***} [-o** *OutputFile***] [-f** *NumThreads***] [-t** *Directory***]**

The input parameters are:

Device
The device name for the file system on which this utility is to be run. This must be the first parameter and is required.

-d [*NodeName***]:{***DiskName***|***DiskNum***|BROKEN}:[***PhysAddr1***-[***PhysAddr2***]]**

NodeName
Specifies a node in the GPFS cluster that has access to the disk in question.

DiskName
Specifies the physical volume name of the disk.

DiskNum
Specifies the GPFS disk ID number of the disk as currently displayed by the **mmlsdisk -L** command.

BROKEN
Specifies to find all files that have broken addresses resulting in lost data. **BROKEN** must be all upper case letters.

*PhysAddr1***-***PhysAddr2*
Specified the range of physical disk addresses, as reported by the error logging utility, for the disk that has encountered a failure. The default value for *PhysAddr1* is zero. The default value for *PhysAddr2* is the value for *PhysAddr1*. Separate these two values by a hyphen (-).

If both *PhysAddr1* and *PhysAddr2* are zero, the entire disk is searched.

-F *DescFile*
Specifies a file containing a list of disks and corresponding disk addresses, for the file system on which the **mmfileid** command is to be run. The file uses the same conventions as the **-d** option.

The format of the file is:

```
#optional comment line
nodename1:hdisk1:PhysAddr1-PhysAddr2
nodename1:hdisk1:PhysAddr3-PhysAddr4
nodename3:hdisk2:PhysAddr1-PhysAddr2
nodename1:hdiskn:PhysAddr1-PhysAddr2
```

For example:

```
:gpfs1nsd
:10:27645856
:BROKEN
```

-f *NumThreads*
Specifies the number of worker threads that are to be created by the **mmfileid** command.

The default is value 16. The minimum value is 1. If *numThreads* is less than 2, only the main process performs work. The maximum value can be as large as is allowed by the operating system **pthread_create** function for a single process. A suggested value is twice the number of disks in the file system. If an error occurs creating a worker thread, **mmfileid** exits with an error.

-o *OutputFile*
Specifies the name of a file for the **mmfileid** command to write its output. If not specified, output is sent to standard output.

-t *Directory*
Specifies the directory to use for the **mmfileid** temporary files and the sort program intermediate files. The default directory is **/tmp**.

The output can be redirected to a file (using the **-o** flag) and sorted on the inode number, using the **sort** command.

The **mmfileid** command output contains one line for each inode found to be located on the corrupt disk sector. The file name is relative to the root of the file system in which it resides. Each line of the command output has this format:

```
inode_number   Logical_Disk_Address    snapshotid    filename
```

inode_number
Indicates the inode number of the file identified by **mmfileid**.

Logical_Disk_Addres
Indicates the disk block (disk sector) number of the file identified by **mmfileid**.

snapshotid
Indicates the snapshotid for the file. A **snapshotid** of 0 means that the file is not included in any snapshot.

filename
Indicates the name of the file identified by **mmfileid**.

This example assumes that a disk analysis tool reported that **hdisk6**, **hdisk7**, **hdisk8**, and **hdisk9** contained bad sectors. This command:

```
mmfileid /dev/gpfsB -F addr.in
```

where **addr.in** contains this:

```
k148n07:hdisk9:2206310-2206810
k148n07:hdisk8:2211038-2211042
k148n07:hdisk8:2201800-2202800
k148n01:hdisk6:2921879-2926880
k148n09:hdisk7:1076208-1076610
```

produces output similar to this:

```
Address 2201958 is contained in the Block allocation map (inode 1)
Address 2206688 is contained in the ACL Data file (inode 4, snapId 0)
Address 2211038 is contained in the Log File (inode 7, snapId 0)
14336 1076256 0 /gpfsB/tesDir/testFile.out
14344 2922528 1 /gpfsB/x.img
```

The lines starting with `Address` are special GPFS files. If your output contains any of these lines, *do not* attempt to replace or repair the indicated files. Call the IBM Support Center for assistance.

The line:

```
14336 1072256 0 /gpfsB/tesDir/testFile.out
```

indicates that inode number 14336, disk address 1072256 contains file **/gpfsB/tesDir/testFile.out**, which does not belong to a snapshot (0 to the left of the name). This file is located on a bad disk sector area.

The line

```
14344 2922528 1 /gpfsB/x.img
```

indicates that inode number 14344, disk address 2922528 contains file **/gpfsB/x.img**, which belongs to snapshot number 1 (1 to the left of the name). This file is located on a bad disk sector area.

The SHA digest

The Secure Hash Algorithm (SHA) digest is relevant only when using GPFS in a multi-cluster environment.

The SHA digest is a short and convenient way to identify a key registered with either the **mmauth show** or **mmremotecluster** command. In theory, two keys may have the same SHA digest. In practice, this is extremely unlikely. The SHA digest can be used by the administrators of two GPFS clusters to determine if they each have received (and registered) the right key file from the other administrator.

An example is the situation of two administrators named **Admin1** and **Admin2** who have registered the others' respective key file, but find that mount attempts by **Admin1** for file systems owned by **Admin2** fail with the error message: 'Authorization failed'. To determine which administrator has registered the wrong key, they each run **mmauth show** and send the local clusters SHA digest to the other administrator. **Admin1** then runs the **mmremotecluster** command and verifies that the SHA digest for **Admin2**'s cluster matches the SHA digest for the key that **Admin1** has registered. **Admin2** then runs the **mmauth show** command and verifies that the SHA digest for **Admin1**'s cluster matches the key that **Admin2** has authorized.

If **Admin1** finds that the SHA digests do not match, **Admin1** runs the **mmremotecluster update** command, passing the correct key file as input.

If **Admin2** finds that the SHA digests do not match, **Admin2** runs the **mmauth update** command, passing the correct key file as input.

This is an example of a **mmauth show** command:

```
Cluster name: fksdcm.pok.ibm.com
Cipher list: EXP1024-RC2-CBC-MD5
SHA digest: d5eb5241eda7d3ec345ece906bfcef0b6cd343bd
File system access: fs1 (rw, root allowed)


Cluster name: kremote.cluster
Cipher list: EXP1024-RC4-SHA
SHA digest: eb71a3aaa89c3979841b363fd6d0a36a2a460a8b
File system access: fs1 (rw, root allowed)

Cluster name: dkq.cluster (this cluster)
Cipher list: AUTHONLY
SHA digest: 090cd57a2e3b18ac163e5e9bd5f26ffabaa6aa25
File system access: (all rw)
```

Chapter 4. Other problem determination tools

Other problem determination tools include the kernel debugging facilities and the **mmpmon** command.

If your problem occurs on the AIX operating system, see the appropriate kernel debugging documentation at: publib.boulder.ibm.com/infocenter/pseries/index.jsp for information about the AIX **kdb** command.

If your problem occurs on the Linux operating system, see the documentation for your distribution vendor.

If your problem occurs on the Windows operating system, the following tools that are available from Microsoft® at: http://www.microsoft.com/en/us/default.aspx, might be useful in troubleshooting:

- Debugging Tools for Windows
- Process Monitor
- Process Explorer
- Microsoft Windows Driver Kit

The **mmpmon** command is intended for system administrators to analyze their I/O on the node on which it is run. It is not primarily a diagnostic tool, but may be used as one for certain problems. For example, running **mmpmon** on several nodes may be used to detect nodes that are experiencing poor performance or connectivity problems.

The syntax of the **mmpmon** command is described in the *Commands* unit of: *General Parallel File System: Administration and Programming Reference*. For details on the **mmpmon** command, see the unit titled *Monitoring GPFS I/O performance with the **mmpmon** command* in *General Parallel File System: Advanced Administration Guide*

Chapter 5. Contacting IBM

Specific information about a problem such as: symptoms, traces, error logs, MMFS logs, and file system status is vital to IBM in order to resolve a GPFS problem.

Obtain this information as quickly as you can after a problem is detected, so that error logs will not wrap and system parameters that are always changing, will be captured as close to the point of failure as possible. When a serious problem is detected, collect this information and then call IBM. For more information, see:

- "Information to collect before contacting the IBM Support Center"
- "How to contact the IBM Support Center" on page 38

Information to collect before contacting the IBM Support Center

For effective communication with the IBM Support Center to help with problem diagnosis, you need to collect certain information.

This data should be available when you contact the IBM Support Center:

1. A description of the problem.
2. Output of the failing application, command, and so forth.
3. Tar file generated by the **gpfs.snap** command containing data from the affected machines. The **gpfs.snap** command can collect data from multiple nodes simultaneously by using the **-w** option. For example,

   ```
   gpfs.snap -w hostname2,hostname3
   ```

   The **gpfs.snap** command is fully described in: "Gathering data to solve GPFS problems" on page 6.
4. Perform this step only on nodes for which you are certain that there is some sort of deadlock or the GPFS daemon is hung. Issue the **mmfsadm dump all** command on at least the failing node and the file system manager node. The analysis of some errors requires this output from all nodes.

   If the node is running Hierarchical Storage Manager (HSM), or the node performs exports to NFS or Samba, issue the **mmfsadm dump kthread** command.

If the **gpfs.snap** command cannot be run, collect these items:

1. Any error log entries relating to the event:

   On an AIX node, issue this command:

   ```
   errpt -a
   ```

   On a Linux node, issue this command:

   ```
   grep "mmfs:" /var/log/messages
   ```

   On a Windows node, use the **Export List...** dialog in the Event Viewer application to save the event log to a file.
2. Perform this step only on nodes for which you are certain that there is some sort of deadlock or the GPFS daemon is hung. Issue the **mmfsadm dump all** command on at least the failing node and the file system manager node. The analysis of some errors requires this output from all nodes.

   If the node is running Hierarchical Storage Manager (HSM), or the node performs exports to NFS or Samba, issue the **mmfsadm dump kthread** command.
3. Create a master **MMFS** log file that is merged and chronologically sorted for the date of the failure (see "Creating a master MMFS.log file" on page 2).
4. If the cluster was configured to store dumps, collect any internal GPFS dumps written to that directory relating to the time of the failure. The default directory is **/tmp/mmfs**.
5. On a failing Linux node, issue this command:

   ```
   rpm -qa
   ```

6. On an failing AIX node, issue this command:
   `lslpp -l`
7. For all of the failing file systems, issue:
   `mmlsfs Device`
8. For all of the failing file systems, issue:
   `mmlsdisk Device`
9. A copy of file **/var/mmfs/gen/mmsdrfs** from the primary cluster configuration server.

How to contact the IBM Support Center

IBM support is available for both customers with an IBM software maintenance contract, and those who do not have one.

- If you have an IBM Software Maintenance service contract, you may phone IBM at:
  1. In the United States the number for support is **1-800-IBM-SERV**.
  2. Outside the United States, contact your local IBM Support Center or see the directory of worldwide contacts at www.ibm.com/planetwide

  Contact the IBM Support Center, for these problems:
  – Node halt or crash not related to a hardware failure
  – Node hang or response problems
  – Failure in other software supplied by IBM.

  You will be asked for the information you collected from “Information to collect before contacting the IBM Support Center” on page 37. You will be given a time period during which an IBM representative will return your call.

  For failures in non-IBM software, follow the problem reporting procedures documented for that product.

  For IBM hardware failures, contact IBM Hardware Support at the number above.

  For any problems reported to the IBM Support Center, a Problem Management Record (PMR) is created. A PMR is an online software record used to keep track of software problems reported by customers.
  – The IBM Support Center representative will create the PMR and give you its number.
  – Have the information you collected available for inclusion in the PMR.
  – Record the PMR number. You will need it to send data to the IBM Support Center. You will also need it on subsequent phone calls to the IBM Support Center to discuss this problem.

  Be sure that the person you identified as your contact can be reached at the phone number you provided in the PMR.
- If you do not have an IBM Software Maintenance service contract, contact your IBM sales representative. You should have the information you collected from “Information to collect before contacting the IBM Support Center” on page 37 available.

Chapter 6. GPFS installation, configuration, and operation problems

You might encounter errors with GPFS installation, configuration, and operation. Use the information in this topic to help you identify and correct errors.

A GPFS installation problem should be suspected when GPFS modules are not loaded successfully, GPFS commands do not work, either on the node that you are working on or on other nodes, new command operands added with a new release of GPFS are not recognized, or there are problems with the kernel extension.

A GPFS configuration problem should be suspected when the GPFS daemon will not activate, it will not remain active, or it fails on some nodes but not on others. Suspect a configuration problem also if quorum is lost, certain nodes appear to hang or do not communicate properly with GPFS, nodes cannot be added to the cluster or are expelled, or GPFS performance is very noticeably degraded once a new release of GPFS is installed or configuration parameters have been changed.

These are some of the errors encountered with GPFS installation, configuration and operation:

Installation and configuration problems

The *General Parallel File System: Concepts, Planning, and Installation Guide* provides the step by step procedure for installing and migrating GPFS, however, some problems might occur if the procedures were not properly followed.

Some of those problems might include:

- Not being able to start GPFS after installation of the latest level. Did you reboot your GPFS nodes subsequent to the last invocation of GPFS at the old level and the first one at the new level? If you did, see "GPFS daemon will not come up" on page 45. If not, reboot. Refer to *General Parallel File System: Concepts, Planning, and Installation Guide* and search for *initialization*.
- Not being able to access a file system. See "File system will not mount" on page 61.
- New GPFS functions do not operate. See "GPFS commands are unsuccessful" on page 56.

What to do after a node of a GPFS cluster crashes and has been reinstalled

After reinstalling GPFS code, check whether the **/var/mmfs/gen/mmsdrfs** file was lost. If it was lost, and an up-to-date version of the file is present on the primary GPFS cluster configuration server, restore the file by issuing this command from the node on which it is missing.

```
mmsdrrestore
```

If the **/var/mmfs/gen/mmsdrfs** file is not present on the primary GPFS cluster configuration server, but it is present on some other node in the cluster, restore the file by issuing these commands:

```
mmsdrrestore -p remoteNode -F remoteFile
mmchcluster -p LATEST
```

where *remoteNode* is the node that has an up-to-date version of the **/var/mmfs/gen/mmsdrfs** file, and *remoteFile* is the full path name of that file on that node.

One way to ensure that the latest version of the **/var/mmfs/gen/mmsdrfs** file is always available is to use the **mmsdrbackup** user exit.

If you have made modifications to any of the user exits in **/var/mmfs/etc**, you will have to restore them before starting GPFS.

For additional information, see “Recovery from loss of GPFS cluster configuration data file” on page 43.

Problems with the /etc/hosts file

The **/ect/hosts** file must have a unique node name for each node interface to be used by GPFS. Violation of this requirement result in the message:

6027-1941
Cannot handle multiple interfaces for host *hostName*.

If you receive this message, correct the **/ect/hosts** file so that each node interface to be used by GPFS appears only once in the file.

Linux configuration considerations

Note: This information applies only to Linux nodes.

Depending on your system configuration, you may need to consider:

1. **Why can only one host successfully attach to the fibre channel loop and see the fibre channel disks?**

   Your host bus adapter may be configured with an enabled *Hard Loop ID* that conflicts with other host bus adapters on the same fibre channel loop.

   To see if that is the case, reboot your machine and enter the adapter bios with **<Alt-Q>** when the fibre channel adapter bios prompt appears. Under the Configuration Settings menu, select `Host Adapter Settings` and either ensure that the Adapter Hard Loop ID option is disabled or assign a unique Hard Loop ID per machine on the fibre channel loop.
2. **Why does GPFS not start automatically after rebooting nodes when using Myrinet?**

   If you have enabled GPFS, with the **mmcrcluster** or **mmchconfig** commands, to autostart on reboot ensure that the Myrinet driver is loaded and that the adapter is configured to run over IP when the machine boots up.
3. **Could the GPFS daemon be terminated due to a memory shortage?**

   The Linux virtual memory manager (VMM) exhibits undesirable behavior for low memory situations on nodes, where the processes with the largest memory usage are killed by the kernel (using OOM killer),

yet no mechanism is available for prioritizing important processes that should not be initial candidates for the OOM killer. The GPFS **mmfsd** daemon uses a large amount of pinned memory in the pagepool for caching data and metadata, and so the **mmfsd** process is a likely candidate for termination if memory must be freed up.

4. **What are the performance tuning suggestions?**

   For an up-to-date list of tuning suggestions, see the Frequently Asked Questions at publib.boulder.ibm.com/infocenter/clresctr/topic/com.ibm.cluster.gpfs.doc/gpfs_faqs/gpfsclustersfaq.html.

Problems with running commands on other nodes

Many of the GPFS administration commands perform operations on nodes other than the node on which the command was issued. This is achieved by utilizing a remote invocation shell and a remote file copy command. By default these items are **/usr/bin/rsh** and **/usr/bin/rcp**. You also have the option of specifying your own remote shell and remote file copy commands to be used instead of the default **rsh** and **rcp**. The remote shell and copy commands must adhere to the same syntax forms as **rsh** and **rcp** but may implement an alternate authentication mechanism. For details, see the **mmcrcluster** and **mmchcluster** commands. These are problems you may encounter with the use of remote commands.

Authorization problems

The **rsh** and **rcp** commands are used by GPFS administration commands to perform operations on other nodes. The **rsh** daemon (**rshd**) on the remote node must recognize the command being run and must obtain authorization to invoke it.

Note: The **rsh** and **rcp** commands that are shipped with the SUA subsystem are not supported on Windows. Use the **ssh** and **scp** commands that are shipped with the Interix openssh package instead.

For the **rsh** and **rcp** commands issued by GPFS administration commands to succeed, each node in the cluster must have an **.rhosts** file in the home directory for the root user, with file permission set to 600. This **.rhosts** file must list each of the nodes and the root user. If such an **.rhosts** file does not exist on each node in the cluster, the **rsh** and **rcp** commands issued by GPFS commands will fail with permission errors, causing the GPFS commands to fail in turn.

If you elected to use installation specific remote invocation shell and remote file copy commands, you must ensure:

1. Proper authorization is granted to all nodes in the GPFS cluster.
2. The nodes in the GPFS cluster can communicate without the use of a password, and without any extraneous messages.

Connectivity problems

Another reason why **rsh** may fail is that connectivity to a needed node has been lost. Error messages from **mmdsh** may indicate that connectivity to such a node has been lost. Here is an example:

```
mmdelnode -N k145n04
Verifying GPFS is stopped on all affected nodes ...
mmdsh: 6027-1617 There are no available nodes on which to run the command.
mmdelnode: 6027-1271 Unexpected error from verifyDaemonInactive: mmcommon onall.
  Return code: 1
```

If error messages indicate that connectivity to a node has been lost, use the **ping** command to verify whether the node can still be reached:

```
ping k145n04
PING k145n04: (119.114.68.69): 56 data bytes
<Ctrl- C>
----k145n04 PING Statistics----
3 packets transmitted, 0 packets received, 100% packet loss
```

If connectivity has been lost, restore it, then reissue the GPFS command.

GPFS error messages for rsh problems

When **rsh** problems arise, the system may display information similar to these error messages:

6027-1615
nodeName remote shell process had return code *value*.

6027-1617
There are no available nodes on which to run the command.

GPFS cluster configuration data files are locked

GPFS uses a file to serialize access of administration commands to the GPFS cluster configuration data files. This lock file is kept on the primary GPFS cluster configuration server in the **/var/mmfs/etc/mmlock** directory. If a system failure occurs before the cleanup of this lock file, the file will remain and subsequent administration commands may report that the GPFS cluster configuration data files are locked. Besides a serialization lock, certain GPFS commands may obtain an additional lock. This lock is designed to prevent GPFS from coming up, or file systems from being mounted, during critical sections of the command processing. If this happens you will see a message that shows the name of the blocking command, similar to message:

6027-1242
GPFS is waiting for ***name***.

To release the lock:

1. Look for **/var/mmfs/etc/mmlock/haslock**. If the **haslock** file exists, it will contain a process identification (PID) number for a process. If this PID number exists in the system, the **haslock** file will prevent running of GPFS commands that need the lock. Examine the contents of the **haslock** file to obtain the PID number and name of the node to which the PID number pertains.
2. Determine who owns the lock. Go to the node shown and issue a **ps** command for the PID. If it is a GPFS administration command that is not responding, stopping the command (be sure you want to do this) will free the lock. If anything else has the PID, another error occurred to the original GPFS command causing it to die without freeing the lock, and this new process has the same PID. If this is the case, do not kill the process.
3. If any locks are held and you want to release them manually, from any node in the GPFS cluster issue the command:

   ```
   mmcommon freeLocks
   ```

GPFS error messages for cluster configuration data file problems

When GPFS commands are unable to retrieve or update the GPFS cluster configuration data files, the system may display information similar to these error messages:

6027-1628
Cannot determine basic environment information. Not enough nodes are available.

6027-1630
The GPFS cluster data on *nodeName* is back level.

6027-1631
The commit process failed.

6027-1632
The GPFS cluster data on *nodeName* is different than the data on *nodeName*.

6027-1633
Failed to create a backup copy of the GPFS cluster data on *nodeName*.

Recovery from loss of GPFS cluster configuration data file

A copy of the GPFS cluster configuration data files is stored in the **/var/mmfs/gen/mmsdrfs** file on each node. For proper operation, this file must exist on each node in the GPFS cluster. The latest level of this file is guaranteed to be on the primary, and secondary if specified, GPFS cluster configuration server nodes that were defined when the GPFS cluster was first created with the **mmcrcluster** command.

If the **/var/mmfs/gen/mmsdrfs** file is removed by accident from any of the nodes, and an up-to-date version of the file is present on the primary GPFS cluster configuration server, restore the file by issuing this command from the node on which it is missing:

```
mmsdrrestore
```

If the **/var/mmfs/gen/mmsdrfs** file is not present on the primary GPFS cluster configuration server, but is present on some other node in the cluster, restore the file by issuing these commands:

```
mmsdrrestore -p remoteNode -F remoteFile
mmchcluster -p LATEST
```

where *remoteNode* is the node that has an up-to-date version of the **/var/mmfs/gen/mmsdrfs** file and *remoteFile* is the full path name of that file on that node.

One way to ensure that the latest version of the **/var/mmfs/gen/mmsdrfs** file is always available is to use the **mmsdrbackup** user exit.

Automatic backup of the GPFS cluster data

GPFS provides an exit, **mmsdrbackup**, that can be used to automatically back up the GPFS configuration data every time it changes. To activate this facility, follow these steps:

1. Modify the GPFS-provided version of **mmsdrbackup** as described in its prologue, to accomplish the backup of the **mmsdrfs** file however the user desires. This file is **/usr/lpp/mmfs/samples/mmsdrbackup.sample**.
2. Copy this modified **mmsdrbackup.sample** file to **/var/mmfs/etc/mmsdrbackup** on all of the nodes in the GPFS cluster. Make sure that the permission bits for **/var/mmfs/etc/mmsdrbackup** are set to permit execution by root.

GPFS will invoke the user-modified version of **mmsdrbackup** in **/var/mmfs/etc** every time a change is made to the **mmsdrfs** file. This will perform the backup of the **mmsdrfs** file according to the user's specifications. See the unit titled *GPFS user exits* in *General Parallel File System: Administration and Programming Reference*.

Error numbers specific to GPFS applications calls

When experiencing installation and configuration problems, GPFS may report these error numbers in the operating system error log facility, or return them to an application:

ECONFIG = 215, Configuration invalid or inconsistent between different nodes.

This error is returned when the levels of software on different nodes cannot coexist. For information about which levels may coexist, see the Frequently Asked Questions at publib.boulder.ibm.com/infocenter/clresctr/topic/com.ibm.cluster.gpfs.doc/gpfs_faqs/gpfsclustersfaq.html.

ENO_QUOTA_INST = 237, No Quota management enabled.

To enable quotas for the file system issue the **mmchfs -Q yes** command. To disable quotas for the file system issue the **mmchfs -Q no** command.

EOFFLINE = 208, Operation failed because a disk is offline

This is most commonly returned when an open of a disk fails. Since GPFS will attempt to continue operation with failed disks, this will be returned when the disk is first needed to complete a

command or application request. If this return code occurs, check your disk subsystem for stopped states and check to determine if the network path exists. In rare situations, this will be reported if disk definitions are incorrect.

EALL_UNAVAIL = 218, A replicated read or write failed because none of the replicas were available.
Multiple disks in multiple failure groups are unavailable. Follow the procedures in Chapter 8, "GPFS disk problems," on page 89 for unavailable disks.

6027-341
Node *nodeName* is incompatible because its maximum compatible version (*number*) is less than the version of this node (*number*).

6027-342
Node *nodeName* is incompatible because its minimum compatible version (*number*) is greater than the version of this node (*number*).

6027-343
Node *nodeName* is incompatible because its version (*number*) is less than the minimum compatible version of this node (*number*).

6027-344
Node *nodeName* is incompatible because its version (*number*) is greater than the maximum compatible version of this node (*number*).

GPFS modules cannot be loaded

This topic applies only to Linux nodes.

You must build the GPFS portability layer binaries based on the kernel configuration of your system. See the *General Parallel File System: Concepts, Planning, and Installation Guide* and search for *GPFS open source portability layer*. During **mmstartup** processing, GPFS loads the **mmfslinux** module.

Some of the more common problems that you may encounter are:

- If the **mmfslinux** module was not built, you may see messages similar to:

```
Loading modules from /usr/lpp/mmfs/bin
Error: /usr/lpp/mmfs/bin/mmfslinux kernel extension does not exist.
       Please compile a custom mmfslinux module for your kernel.
       See /usr/lpp/mmfs/src/README for directions.
Error: unable to verify kernel/module configuration
```

  Follow the directions in the *General Parallel File System: Concepts, Planning, and Installation Guide* to build an appropriate **mmfslinux** module.
- If the **mmfslinux** module is incompatible with your system, you may experience a kernel panic on GPFS startup. Ensure that the **site.mcr** has been configured properly from the **site.mcr.proto**, and GPFS has been built and installed properly. Follow the directions in the *General Parallel File System: Concepts, Planning, and Installation Guide* to build an appropriate **mmfslinux** module.

Unable to load a GPFS kernel extension on Linux

There are a few reasons why you might not be able to load a GPFS kernel extension on Linux.

This problems occurs if:

1. The portability layer is not built.
2. The GPFS kernel modules, **mmfslinux** and **tracedev**, are built with a kernel version that differs from that of the currently running Linux kernel. This situation can occur if the modules are built on another node with a different kernel version and copied to this node, or if the node is rebooted using a kernel with a different version.

Here is an example log:

```
Wed Oct 6 16:00:49 EDT 2004 runmmfs starting
Removing old /var/adm/ras/mmfs.log.* files:
Unloading modules from /usr/lpp/mmfs/bin
Error: /usr/lpp/mmfs/bin/mmfslinux kernel extension does not exist.
Please compile a custom mmfslinux module for your kernel.
See /usr/lpp/mmfs/src/README for directions.
Error: unable to verify kernel/module configuration
Loading modules from /usr/lpp/mmfs/bin
Error: /usr/lpp/mmfs/bin/mmfslinux kernel extension does not exist.
Please compile a custom mmfslinux module for your kernel.
See /usr/lpp/mmfs/src/README for directions.
Error: unable to verify kernel/module configuration
Wed Oct 6 16:00:49 EDT 2004 runmmfs: error in loading or unloading the mmfs kernel extension
Wed Oct 6 16:00:49 EDT 2004 runmmfs: stopping GPFS
```

GPFS daemon will not come up

There are several indications that could lead you to the conclusion that the GPFS daemon (**mmfsd**) will not come up and there are some steps to follow to correct the problem.

Those indications include:

- The file system has been enabled to mount automatically, but the mount has not completed.
- You issue a GPFS command and receive the message:

  6027-665
  Failed to connect to file system daemon: Connection refused.
- The **MMFS log** does not contain the message:

  6027-300
  mmfsd ready.
- The **MMFS log** file contains this error message: 'Error: daemon and kernel extension do not match.' This error indicates that the kernel extension currently loaded in memory and the daemon currently starting have mismatching versions. This situation may arise if a GPFS code update has been applied, and the node has not been rebooted prior to starting GPFS.

  While GPFS scripts attempt to unload the old kernel extension during update and install operations, such attempts may fail if the operating system is still referencing GPFS code and data structures. To recover from this error, ensure that all GPFS file systems are successfully unmounted, and reboot the node. The **mmlsmount** command can be used to ensure that all file systems are unmounted.

Steps to follow if the GPFS daemon does not come up

1. See "GPFS modules cannot be loaded" on page 44 if your node is running Linux, to verify that you have built the portability layer.
2. Verify that the GPFS daemon is active. Issue:

   ```
   ps -e | grep mmfs
   ```

   The output of this command should list **mmfsd** as operational.

   For example,

   ```
   12230 pts/8 00:00:00 mmfsd
   ```

   If the output does not show this, the GPFS daemon needs to be started with the **mmstartup** command.
3. If you did not specify the **autoload** option on the **mmcrcluster** or the **mmchconfig** command, you need to manually start the daemon by issuing the **mmstartup** command.

If you specified the **autoload** option, someone may have issued the **mmshutdown** command. In this case, issue the **mmstartup** command. When using **autoload** for the first time, **mmstartup** must be run manually. The **autoload** takes effect on the next reboot.

4. Verify that the network upon which your GPFS cluster depends is up. Issue:

   ```
   ping nodename
   ```

   to each node in the cluster. A properly working network and node will correctly reply to the ping with no lost packets.

   Query the network interface that GPFS is using with:

   ```
   netstat -i
   ```

   A properly working network will report no transmission errors.

5. Verify that the GPFS cluster configuration data is available by looking in the **MMFS log**. If you see the message:

   ```
   027-1592 Unable to retrieve GPFS cluster files from nodename.
    Determine the problem with accessing node nodename and correct it.
   ```

6. Verify that the GPFS environment is properly initialized by issuing these commands and ensuring that the output is as expected. If any of the commands produce unexpected results, this may be an indication of corrupted GPFS cluster configuration data file information. Contact the IBM Support Center.

   - List the cluster configuration. This will also update the GPFS configuration data on the node. Correct any reported errors before continuing.

     ```
     mmlscluster
     ```

   - List all file systems that were created in this cluster. For an AIX node, issue:

     ```
      lsfs -v mmfs
     ```

     For a Linux node, issue:

     ```
     cat /etc/fstab | grep gpfs
     ```

7. GPFS requires a quorum of nodes to be active before any file system operations can be honored. This requirement guarantees that a valid single token management domain exists for each GPFS file system. Prior to the existence of a quorum, most requests are rejected with a message indicating that quorum does not exist.

   To identify which nodes in the cluster have daemons **up** or **down**, issue:

   ```
   mmgetstate -L -a
   ```

   If insufficient nodes are active to achieve quorum, go to any nodes not listed as **active** and perform problem determination steps on these nodes. A node indicates that it is part of a quorum by writing a **mmfsd ready** message to the **MMFS log**. Remember that your system may have quorum nodes and non-quorum nodes, and only quorum nodes are counted to achieve the quorum.

8. This step applies only to AIX nodes. Verify that GPFS kernel extension is not having problems with its shared segment by invoking:

   ```
   cat /var/adm/ras/mmfs.log.latest
   ```

   Messages such as:

   6027-319
   Could not create shared segment.

   must be corrected by the following procedure:

   a. Issue **mmshutdown**.
   b. Wait for the command to finish. If messages are displayed indicating that it is waiting for a virtual shared disk to complete a pending I/O, do not remove the shared segment as outlined in the following procedure. Otherwise the machine will panic. If you are using IBM Virtual Shared Disks you may stop pending I/O by issuing the **suspendvsd** and **stopvsd** commands for the affected

disks. Any pending I/O will complete with an error return code. Issue the **startvsd** command to bring the IBM Virtual Shared Disk back to its previous state.

c. **Removing the shared segment in an AIX 32-bit environment**

1) Determine if there are any programs with a hold on a shared segment by issuing:

```
ipcs -mb
```

A `D` will appear in the attributes for any program holding a shared segment and the segment will stay until the program exits.

2) When you are ready to delete the shared segment, issue:

```
ipcs -m | grep 0x9283a | awk '{print $2}' | xargs -n1 ipcrm -m
```

d. **Removing the shared segment in an AIX 64-bit environment:**

1) Issue the **mmshutdown** command.

2) Issue the **mmfsadm cleanup** command.

e. If you are still unable to resolve the problem, reboot the node.

9. If the previous GPFS daemon was brought down and you are trying to start a new daemon but are unable to, this is an indication that the original daemon did not completely go away. Go to that node and check the state of GPFS. Stopping and restarting GPFS or rebooting this node will often return GPFS to normal operation. If this fails, see "Information to collect before contacting the IBM Support Center" on page 37 and call the IBM Support Center.

Unable to start GPFS after the installation of a new release of GPFS

If one or more nodes in the cluster will not start GPFS, these are the possible causes:

- If message:

  6027-2700
  A node join was rejected. This could be due to incompatible daemon versions, failure to find the node in the configuration database, or no configuration manager found.

  is written to the **MMFS log**, incompatible versions of GPFS code exist on nodes within the same cluster.
- If messages stating that functions are not supported are written to the **MMFS log**, you may not have the correct kernel extensions loaded.
  1. Ensure that the latest GPFS install packages are loaded on your system.
  2. If running on Linux, ensure that the latest kernel extensions have been installed and built. See *General Parallel File System: Concepts, Planning, and Installation Guide* and search for *Building your GPFS portability module*.
  3. Reboot the GPFS node after an installation to ensure that the latest kernel extension is loaded.
- The daemon will not start because the configuration data was not migrated. See "Installation and configuration problems" on page 39.

GPFS error messages for shared segments and network problems

Shared segment problems. Follow the problem determination and repair actions specified in the accompanying messages.

6027-319
Could not create shared segment.

6027-320
Could not map shared segment.

6027-321
Shared segment mapped at wrong address (is *value*, should be *value*).

6027-322
Could not map shared segment in kernel extension.

Network problems:

6027-306
Could not initialize internode communication.

Error numbers specific to GPFS application calls when the daemon is unable to come up

When the daemon is unable to come up, GPFS may report these error numbers in the operating system error log, or return them to an application:

ECONFIG = 215, Configuration invalid or inconsistent between different nodes.
This error is returned when the levels of software on different nodes cannot coexist. For information about which levels may coexist, see the Frequently Asked Questions at publib.boulder.ibm.com/infocenter/clresctr/topic/com.ibm.cluster.gpfs.doc/gpfs_faqs/gpfsclustersfaq.html.

6027-341
Node *nodeName* is incompatible because its maximum compatible version (*number*) is less than the version of this node (*number*).

6027-342
Node *nodeName* is incompatible because its minimum compatible version (*number*) is greater than the version of this node (*number*).

6027-343
Node *nodeName* is incompatible because its version (*number*) is less than the minimum compatible version of this node (*number*).

6027-344
Node *nodeName* is incompatible because its version (*number*) is greater than the maximum compatible version of this node (*number*).

GPFS daemon went down

There are a number of conditions that can cause the GPFS daemon to exit.

These are all conditions where the GPFS internal checking has determined that continued operation would be dangerous to the consistency of your data. Some of these conditions are errors within GPFS processing but most represent a failure of the surrounding environment.

In most cases, the daemon will exit and restart after recovery. If it is not safe to simply force the unmounted file systems to recover, the GPFS daemon will exit.

Indications leading you to the conclusion that the daemon went down:

- Applications running at the time of the failure will see either `ENODEV` or `ESTALE` errors. The `ENODEV` errors are generated by the operating system until the daemon has restarted. The `ESTALE` error is generated by GPFS as soon as it restarts.

  When quorum is lost, applications with open files receive an `ESTALE` error return code until the files are closed and reopened. New file open operations will fail until quorum is restored and the file system is remounted. Applications accessing these files prior to GPFS return may receive a `ENODEV` return code from the operating system.
- The **MMFS log** contains the message:

  6027-650
  The mmfs daemon is shutting down abnormally.

Most GPFS daemon down error messages are in the **mmfs.log.previous** log for the instance that failed. If the daemon restarted, it generates a new **mmfs.log.latest**. Begin problem determination for these errors by examining the operating system error log.

If an existing quorum is lost, GPFS stops all processing within the cluster to protect the integrity of your data. GPFS will attempt to rebuild a quorum of nodes and will remount the file system if automatic mounts are specified.

- Open requests are rejected with `no such file` or `no such directory` errors.

  When quorum has been lost, requests are rejected until the node has rejoined a valid quorum and mounted its file systems. If messages indicate lack of quorum, follow the procedures in "GPFS daemon will not come up" on page 45.
- For all other errors follow the procedures in "Information to collect before contacting the IBM Support Center" on page 37.
- Removing the setuid bit from the permissions of these commands may produce errors for non-root users:
  mmdf
  mmgetacl
  mmlsdisk
  mmlsfs
  mmlsmgr
  mmlspolicy
  mmlsquota
  mmlssnapshot
  mmputacl
  mmsnapdir
  mmsnaplatest

  The GPFS system-level versions of these commands (prefixed by **ts**) may need to be checked for how permissions are set if non-root users see the following message:

  ```
  6027-1209 GPFS is down on this node
  ```

  If the setuid bit is removed from the permissions on the system-level commands, the command cannot be executed and the node is perceived as being down. The system-level versons of the commands are:
  tsdf64
  tslsdisk64
  tslsfs64
  tslsmgr64
  tslspolicy64
  tslsquota64
  tslssnapshot64
  tssnapdir64
  tssnaplatest64

  These are found in the following directories:
  - On 32-bit AIX machines, in directory **/usr/lpp/mmfs/bin/aix32**
  - On 64-bit AIX machines, in the directory **/usr/lpp/mmfs/bin/aix64**
  - On Linux and Windows machines, in the directory **/usr/lpp/mmfs/bin**.

  Note: The mode bits for all listed commands are 4555 or **-r-sr-xr-x**. To restore the default (shipped) permission, enter:

  ```
  chmod 4555 tscommand
  ```

Attention: Only administration-level versions of GPFS commands (prefixed by **mm**) should be executed. Executing system-level commands (prefixed by **ts**) directly will produce unexpected results.

GPFS failures due to a network failure

For proper functioning, GPFS depends both directly and indirectly on correct network operation.

This dependence is direct because various GPFS internal messages flow on the network, and may be indirect if the underlying disk technology is dependent on the network. For example, the IBM Virtual Shared Disk subsystem is dependent on network connectivity, and if GPFS is utilizing that subsystem, a network failure recognized by IBM Virtual Shared Disk will ultimately affect GPFS. Symptoms included in an indirect failure would be inability to complete I/O or GPFS moving disks to the **down** state.

If a failure is first detected by an underlying disk subsystem such as the IBM Virtual Shared Disk, the error will usually be seen by GPFS as a series of disk errors.

The problem can also be first detected by the GPFS network communication layer. If network connectivity is lost between nodes or GPFS heart beating services cannot sustain communication to a node, GPFS will declare the node dead and perform recovery procedures. This problem will manifest itself by messages appearing in the **MMFS log** such as:

```
Mon Jun 25 22:23:36.298 2007: Close connection to 192.168.10.109 c5n109. Attempting reconnect.
Mon Jun 25 22:23:37.300 2007: Connecting to 192.168.10.109 c5n109
Mon Jun 25 22:23:37.398 2007: Close connection to 192.168.10.109 c5n109
Mon Jun 25 22:23:38.338 2007: Recovering nodes: 9.114.132.109
Mon Jun 25 22:23:38.722 2007: Recovered 1 nodes.
```

Nodes mounting file systems owned and served by other clusters may receive error messages similar to this:

```
Mon Jun 25 16:11:16 2007: Close connection to 89.116.94.81 k155n01
Mon Jun 25 16:11:21 2007: Lost membership in cluster remote.cluster. Unmounting file systems.
```

If sufficient numbers of nodes fail, GPFS will lose the quorum of nodes, which exhibits itself by messages appearing in the **MMFS log**, similar to this:

```
Mon Jun 25 11:08:10 2007: Close connection to 179.32.65.4 gpfs2
Mon Jun 25 11:08:10 2007: Lost membership in cluster gpfsxx.kgn.ibm.com. Unmounting file system.
```

When either of these cases occur, perform problem determination on your network connectivity. Failing components could be network hardware such as switches or host bus adapters.

Kernel panics with a 'GPFS dead man switch timer has expired, and there's still outstanding I/O requests' message

This problem can be detected by an error log with a label of KERNEL_PANIC, and the PANIC MESSAGES or a PANIC STRING.

For example:

```
GPFS Deadman Switch timer has expired, and there's still  outstanding I/O requests
```

GPFS is designed to tolerate node failures through per-node metadata logging (journaling). The log file is called the *recovery log*. In the event of a node failure, GPFS performs recovery by replaying the recovery log for the failed node, thus restoring the file system to a consistent state and allowing other nodes to continue working. Prior to replaying the recovery log, it is critical to ensure that the failed node has indeed failed, as opposed to being active but unable to communicate with the rest of the cluster.

In the latter case, if the failed node has direct access (as opposed to accessing the disk with an NSD server) to any disks that are a part of the GPFS file system, it is necessary to ensure that no I/O requests submitted from this node complete once the recovery log replay has started. To accomplish this, GPFS

uses the disk lease mechanism. The disk leasing mechanism guarantees that a node does not submit any more I/O requests once its disk lease has expired, and the surviving nodes use disk lease time out as a guideline for starting recovery.

This situation is complicated by the possibility of 'hung I/O'. If an I/O request is submitted prior to the disk lease expiration, but for some reason (for example, device driver malfunction) the I/O takes a long time to complete, it is possible that it may complete after the start of the recovery log replay during recovery. This situation would present a risk of file system corruption. In order to guard against such a contingency, when I/O requests are being issued directly to the underlying disk device, GPFS initiates a kernel timer, referred to as **dead man switch**. The **dead man switch** timer goes off in the event of disk lease expiration, and checks whether there is any outstanding I/O requests. If there is any I/O pending, a kernel panic is initiated to prevent possible file system corruption.

Such a kernel panic is not an indication of a software defect in GPFS or the operating system kernel, but rather it is a sign of

1. Network problems (the node is unable to renew its disk lease).
2. Problems accessing the disk device (I/O requests take an abnormally long time to complete). See "MMFS_LONGDISKIO" on page 4.

Quorum loss

Each GPFS cluster has a set of quorum nodes explicitly set by the cluster administrator.

These quorum nodes and the selected quorum algorithm determine the availability of file systems owned by the cluster. See *General Parallel File System: Concepts, Planning, and Installation Guide* and search for *quorum*.

When quorum loss or loss of connectivity occurs, any nodes still running GPFS suspend the use of file systems owned by the cluster experiencing the problem. This may result in GPFS access within the suspended file system receiving **ESTALE** errnos. Nodes continuing to function after suspending file system access will start contacting other nodes in the cluster in an attempt to rejoin or reform the quorum. If they succeed in forming a quorum, access to the file system is restarted.

Normally, quorum loss or loss of connectivity occurs if a node goes down or becomes isolated from its peers by a network failure. The expected response is to address the failing condition.

Delays and deadlocks

The first item to check when a file system appears hung is the condition of the networks including the network used to access the disks.

Look for increasing numbers of dropped packets on all nodes by issuing:

- **netstat -D** on an AIX node
- **ifconfig** *interfacename*, where *interfacename* is the name of the interface being used by GPFS for communication.

When using subnets (see *Using remote access with public and private IP addresses* in the *General Parallel File System: Advanced Administration Guide*), different interfaces may be in use for intra-cluster and intercluster communication. The presence of a hang or dropped packed condition indicates a network support issue that should be pursued first. Contact your local network administrator for problem determination for your specific network configuration.

If file system processes appear to stop making progress, there may be a system resource problem or an internal deadlock within GPFS.

1. First, check how full your file system is by issuing the **mmdf** command. If the **mmdf** command does not respond, contact the IBM Support Center. Otherwise, the system displays information similar to:

```
disk                disk size  failure holds    holds              free KB             free KB
name                    in KB    group metadata data        in full blocks        in fragments
--------------- ------------- -------- -------- ----- -------------------- -------------------
Disks in storage pool: system
sdcnsd              142028560        5 yes      yes       119856640 ( 84%)        264400 ( 0%)
sddnsd              142028560        5 yes      yes       119878144 ( 84%)        266512 ( 0%)
sdensd              142028560        5 yes      yes       119892480 ( 84%)        272160 ( 0%)
hd8vsdn97           139984896        5 yes      no        139461120 (100%)         52640 ( 0%)
hd7vsdn97           139984896        5 yes      yes       129242624 ( 92%)        196160 ( 0%) *
hd6n98              140095488        5 yes      yes       118995968 ( 85%)        247712 ( 0%)
hd5n98              140095488        5 yes      yes       118928896 ( 85%)        259312 ( 0%)
sdbnsd              142028560        5 no       yes       120285184 ( 85%)        217632 ( 0%)
hd3n97              140095488        5 yes      yes       118957568 ( 85%)        245056 ( 0%)
hd2n97              140095488        5 yes      yes       118931456 ( 85%)        253536 ( 0%)
sdlnsd              142028560        6 yes      yes       119882240 ( 84%)        282096 ( 0%)
sdknsd              142028560        6 yes      yes       119900160 ( 84%)        282112 ( 0%)
sdjnsd              142028560        6 yes      yes       131812352 ( 93%)         84704 ( 0%)
sdinsd              142028560        6 yes      yes       119889408 ( 84%)        281680 ( 0%)
sdhnsd              142028560        6 yes      yes       119862784 ( 84%)        287392 ( 0%)
sdgnsd              142028560        6 yes      yes       119861248 ( 84%)        286128 ( 0%)
sdfnsd              142028560        6 yes      yes       119860224 ( 84%)        278720 ( 0%)
hd9vsdn97           139984896        6 yes      yes       118980608 ( 85%)        268960 ( 0%)
hd10vsdn98          139984896        6 yes      yes       118963200 ( 85%)        267408 ( 0%)
hd11vsdn98          139984896        6 yes      yes       118989312 ( 85%)        250080 ( 0%)
hd12vsdn98          139984896        6 yes      yes       118991872 ( 85%)        264624 ( 0%)
hd13n97             140095488        7 yes      yes       118893056 ( 85%)        264336 ( 0%)
hd14n97             140095488        7 yes      yes       118862336 ( 85%)        251840 ( 0%)
hd15n97             140095488        7 yes      yes       118886400 ( 85%)        254944 ( 0%)
hd16n97             140095488        7 yes      yes       118878208 ( 85%)        263088 ( 0%)
hd17n97             140095488        7 yes      yes       118866432 ( 85%)        259328 ( 0%)
hd18n97             140095488        7 no       no                0 (  0%)             0 ( 0%)
hd19n98             140095488        7 yes      yes       118867456 ( 85%)        260448 ( 0%)
hd20n98             140095488        7 yes      yes       118886400 ( 85%)        263008 ( 0%)
hd21n98             140095488        7 yes      yes       118873600 ( 85%)        259168 ( 0%)
hd22n98             140095488        7 yes      yes       136262144 ( 97%)         85392 ( 0%)
hd23n98             140094992        7 yes      yes       118865408 ( 85%)        260240 ( 0%)
                -------------                         -------------------- -------------------
(pool total)       4363670464                            3628322304 ( 83%)       7334656 ( 0%)

                =============                         ==================== ===================
(data)             4083590080                            3488861184 ( 85%)       7282016 ( 0%)
(metadata)         4081546416                            3508037120 ( 86%)       7117024 ( 0%)
                =============                         ==================== ===================
(total)            4363670464                            3628322304 ( 83%)       7334656 ( 0%)

Inode Information
-----------------
Number of used inodes:           87789
Number of free inodes:         4447129
Number of allocated inodes:    4534918
Maximum number of inodes:      4534918
```

GPFS operations that involve allocation of data and metadata blocks (that is, file creation and writes) will slow down significantly if the number of free blocks drops below 5% of the total number. Free up some space by deleting some files or snapshots (keeping in mind that deleting a file will not necessarily result in any disk space being freed up when snapshots are present). Another possible cause of a performance loss is the lack of free inodes. Issue the **mmchfs** command to increase the number of inodes for the file system so there is at least a minimum of 5% free.

GPFS error messages for file system delays and deadlocks

6027-533

File system *file_system* is approaching the limit for the maximum number of inodes or files.

operating system error log entry

Jul 19 12:51:49 node1 mmfs: Error=MMFS_SYSTEM_WARNING, ID=0x4DC797C6, Tag=3690419: File system warning. Volume fs1. Reason: File system fs1 is approaching the limit for the maximum number of inodes/files.

2. If the file system is not nearing its maximum number of files, look at threads that are waiting. On all nodes of a live system, to determine which nodes have threads waiting longer than 10 seconds, issue this command on each node:

   ```
   /usr/lpp/mmfs/bin/mmfsadm dump waiters 10
   ```

   Determine which nodes have the long waiting threads.

   For all nodes that have threads waiting longer than 10 seconds, issue:

   ```
   mmfsadm dump all
   ```

   Note:

   a. Each node can potentially dump up to 200 MB of data.

   b. Run an **mmfsadm dump all** command only on nodes that you are sure the threads are really hung. An **mmfsadm dump all** command can follow pointers that are changing and cause the node to crash. To determine if a node is hung, determine the percentage of user time that is running (1 indicates quiet time), issue:

      ```
      vmstat 5 2
      ```

   Your next course of action depends on the information found in the dump of the waiters. Examine each long waiting thread and determine what it is waiting on:

   kxAwaitIOCompletion

   Indicates an abnormal wait time for I/O completion. Look into underlying disk subsystem problems. To determine which disk is being waited on, look in the **dump disk** section for:

   ```
   In progress:
   In progress: read operation ...
   In progress: write operation ...
   ```

   If the lines below the **In progress** say **none**, there is no outstanding I/O for that device.

   If your disks are virtual shared disk, go to the appropriate IBM Reliable Scalable Cluster Technology (RSCT) document at: publib.boulder.ibm.com/clresctr/windows/public/rsctbooks.html and search for *virtual shared disk*. Follow the problem determination and repair actions specified.

   The problem could be either a communication or a disk problem. This section will list the active server node for the hung I/O operation on that NSD. To determine if the hung operation is a communication or disk device problem, go to the **dump nsd** section in the output for the active server node. Then go to the **dump disk** section. If there are long *in progress* times for the NSD in the **dump disk** section, perform disk problem determination again. If there are no *in progress* I/O operations, but there are long communication waiters, see the long **Remote Procedure Call** section below.

   Remote Procedure Call (RPC) waits

   Determine why the RPC has not completed on some other node.

   In the dump, look in the *tscomm* section to find the same thread with **pending** messages to find the other node's IP address and the message name that was sent. If there are multiple nodes with **pending** replies, you have to check all of them to see the holdup. **success** means the reply has already returned.

   If the message is **tmMsgRevoke**, there might not be a thread on the other node to handle it. These can be queued on the lockable object for later handling.

   At the end of the *tscomm* section, there may also be messages that have not been given a handler thread. This happens quite often when there are hundreds of **sgmMsgMount** requests coming to a file system manager all at once. This is normal.

In some rare cases where there are TCP problems, we have seen messages only partially returned. In this case there will be a few extra lines after a node in the connection table that shows how many bytes have been received so far and how many it expects.

In the dump from the other node see why the message handler for that message name is waiting. In the **tscomm** section, if you see `destination ip@ pending`, contact the IBM Support Center.

Wait on a mutex

Follow the mutex pointer value in the *dump mutexes* section to see which thread is holding the mutex, and find that thread in the *dump threads* section to see why that thread is waiting. The *desc* pointer on a thread is what is recorded as a mutex holder.

If you see a *by kernel thread*, issue an **mmfsadm dump kthread**.

Wait on condition variable

These are harder than mutexes to diagnose, because it requires a knowledge of the semantics of state changes in internal structures. In these cases, the thread is waiting for another thread to make a state change and wake up the waiting thread (like a Linux event wait).

Wait on a lock

There are *lockable* objects such as OpenFiles, AllocSegments, IndBlocks, that have keys. The keys look like 4 hex words separated by colons. Finding the object that this thread is working on is usually a matter of :

a. Finding any object with the *writer* field holding the thread's *desc* value (if the thread has an exclusive type lock), or finding any object that already has a lock on it. For example, [lx]. These locks are in the *dump locks* section, but also show up in the *dump files* and *dump buffers* sections.

   If the node is also the file system manager, there is a *dump tokenmgr* section that has token information for all the nodes. Also for the file system manager in the *dump threads* section, look for a `KX`. If there is a `KX` call, this indicates a call to the kernel was being made. Issue an **mmfsadm dump kthread**.

b. It is usually difficult to find out why the threads that have the lock are not releasing it. Usually there is another thread, or kthread, to follow in this case.

Node cannot be added to the GPFS cluster

There is an indication leading you to the conclusion that a node cannot be added to a cluster and steps to follow to correct the problem.

That indication is:

- You issue the **mmcrcluster** or **mmaddnode** command and receive the message:

  6027-1598

  Node *name* was not added to the cluster. The node appears to already belong to a GPFS cluster.

Steps to follow if a node cannot be added to a cluster:

1. Run the **mmlscluster** command to verify that the node is not in the cluster.
2. If the node is not in the cluster, issue this command on the node that could not be added:

   ```
   mmdelnode -f
   ```

3. Reissue the **mmaddnode** command.

Remote node expelled after remote file system successfully mounted

This problem produces 'node expelled from cluster' messages.

One cause of this condition is when the **subnets** attribute of the **mmchconfig** command has been used to specify subnets to GPFS, and there is an incorrect netmask specification on one or more nodes of the clusters involved in the remote mount. Check to be sure that all netmasks are correct for the network interfaces used for GPFS communication.

Disaster recovery problems

As with any type of problem or failure, obtain the GPFS log files (**mmfs.log.***) from all nodes in the cluster and, if available, the content of the internal dumps.

For more information see:

- The topic titled *Establishing disaster recovery for your GPFS cluster* in the *General Parallel File System: Advanced Administration Guide* for detailed information about GPFS disaster recovery
- "Creating a master MMFS.log file" on page 2 and "Information to collect before contacting the IBM Support Center" on page 37

These two messages:

6027-489
The file system descriptor quorum has been overridden.

6027-490
The descriptor replica on disk *name* has been excluded.

might appear in the **MMFS log** for active/active disaster recovery scenarios with GPFS replication. The purpose of these messages is to record the fact that a quorum override decision has been made after the loss of a majority of disks. A message similar to the these will appear in the log on the file system manager node every time it reads the file system descriptor with an overridden quorum:

```
...
6027-489 The file system descriptor quorum has been overridden.
6027-490 The descriptor replica on disk gpfs23nsd has been excluded.
6027-490 The descriptor replica on disk gpfs24nsd has been excluded.
...
```

For more information on quorum override, see *General Parallel File System: Concepts, Planning, and Installation Guide* and search on *quorum*.

For PPRC and FlashCopy-based configurations, additional problem determination information may be collected from the ESS log file. This information and the appropriate ESS documentation should be consulted when dealing with various types disk subsystem-related failures. For instance, if users are unable to perform a PPRC failover (or failback) task successfully or unable to generate a FlashCopy® of a disk volume, they should consult the subsystem log and the appropriate ESS documentation. For additional information, refer to:

- *IBM Enterprise Storage Server* (http://www.redbooks.ibm.com/redbooks/pdfs/sg245465.pdf)
- *IBM Enterprise Storage Server Web Interface User's Guide* (http://publibfp.boulder.ibm.com/epubs/pdf/f2bui05.pdf).

Disaster recovery setup problems

These setup problems may impact your ability to use disaster recovery successfully:

1. Considerations of data integrity require proper setup of PPRC consistency groups in PPRC environments. Additionally, when using the FlashCopy facility, make sure to suspend all I/O activity before generating the FlashCopy image. See "Data integrity" on page 87.
2. In certain cases, it may not be possible to restore access to the file system even after relaxing the node and disk quorums. For example, in a three failure group configuration, GPFS will tolerate and recover from a complete loss of a single failure group (and the tiebreaker with a quorum override).

However, all disks in the remaining failure group must remain active and usable in order for the file system to continue its operation. A subsequent loss of at least one of the disks in the remaining failure group would render the file system unusable and trigger a forced unmount. In such situations, users may still be able perform a restricted mount (as described in “Restricted mode mount” on page 21) and attempt to recover parts of their data from the damaged file system.

3. When running **mmfsctl syncFSconfig**, you may get an error similar to this one:

   ```
   mmfsctl: None of the nodes in the peer cluster can be reached
   ```

   If this happens, check the network connectivity between the peer GPFS clusters and verify their remote shell setup. This command requires full TCP/IP connectivity between the two sites, and all nodes must be able to communicate using ssh or rsh without the use of a password.

Other problems with disaster recovery

1. Currently, users are advised to always specify the **all** option when invoking the **mmfsctl syncFSconfig** command, rather than the device name of one specific file system. This enables GPFS to detect and correctly resolve the configuration discrepancies that may have occurred as result of a manual administrative action in the target GPFS cluster (the one into which the configuration is being imported).
2. The optional **SpecFile** parameter to the **mmfsctl syncFSconfig** (specified with the **-S** flag) must be a fully-qualified path name defining the location of the spec data file on nodes in the target cluster. It is not the local path name to the file on the node, from which the **mmfsctl** command is being issued. A copy of this file must be available at the provided path name on all peer contact nodes (the ones defined in the **RemoteNodesFile**).

GPFS commands are unsuccessful

GPFS commands can be unsuccessful for various reasons.

Unsuccessful command results will be indicated by:

- Return codes indicating the GPFS daemon is no longer running.
- Command specific problems indicating you are unable to access the disks.
- A nonzero return code from the GPFS command.

Some reasons that GPFS commands can be unsuccessful include:

1. If all commands are generically unsuccessful, this may be due to a daemon failure. Verify that the GPFS daemon is active. Issue:

   ```
   mmgetstate
   ```

   If the daemon is not active, check **/var/adm/ras/mmfs.log.previous** on the local node and on the file system manager node. This file enumerates the failing sequence of the GPFS daemon.

   If there is a communication failure with the file system manager node, you will receive an error and the **errno** global variable may be set to EIO (I/O error).
2. Verify the GPFS cluster configuration data files are not locked and are accessible. To determine if the GPFS cluster configuration data files are locked, see “GPFS cluster configuration data files are locked” on page 42.
3. The **rsh** command is not functioning correctly. See “Authorization problems” on page 41.

   If **rsh** is not functioning properly on a node in the GPFS cluster, a GPFS administration command that needs to run work on that node will fail with a ’permission is denied’ error. The system displays information similar to:

   ```
   mmlscluster
   rshd: 0826-813 Permission is denied.
   mmdsh: 6027-1615 k145n02 rsh process had return code 1.
   mmlscluster: 6027-1591 Attention: Unable to retrieve GPFS cluster files from
   node k145n02
   ```

```
rshd: 0826-813 Permission is denied.
mmdsh: 6027-1615 k145n01 rsh process had return code 1.
mmlscluster: 6027-1592 Unable to retrieve GPFS cluster files from node k145n01
```

These messages indicate that **rsh** is not working properly on nodes **k145n01** and **k145n02**.

If you encounter this type of failure, determine why **rsh** is not working on the identified node. Then fix the problem.

4. Most problems encountered during file system creation fall into three classes:
   - You did not create network shared disks which are required to build the file system.
   - The creation operation cannot access the disk.

     Follow the procedures for checking access to the disk. This can result from a number of factors including those described in "NSD and underlying disk subsystem failures" on page 89.
   - Unsuccessful attempt to communicate with the file system manager.

     The file system creation runs on the file system manager node. If that node goes down, the **mmcrfs** command may not succeed.
5. If the **mmdelnode** command was unsuccessful and you plan to permanently de-install GPFS from a node, you should first remove the node from the cluster. If this is not done and you run the **mmdelnode** command after the **mmfs** code is removed, the command will fail and display a message similar to this example:

   ```
   Verifying GPFS is stopped on all affected nodes ...
   k145n05: ksh: /usr/lpp/mmfs/bin/mmremote: not found.
   ```

   If this happens, power off the node and run the **mmdelnode** command again.
6. If you have successfully installed and are operating with the latest level of GPFS, but cannot run the new functions available, it is probable that you have not issued the **mmchfs -V full** or **mmchfs -V compat** command to change the version of the file system. This command must be issued for *each* of your file systems.

   Note: Before issuing the **-V** option (with **full** or **compat**), see the ″Migration, coexistence and compatibility″ topic in the *General Parallel File System: Concepts, Planning, and Installation Guide*. You must ensure that all nodes in the cluster have been migrated to the latest level of GPFS code and that you have successfully run the **mmchconfig release=LATEST** command.

   Make sure you have operated with the new level of code for some time and are certain you want to migrate to the latest level of GPFS. Issue the **mmchfs -V full** command only after you have definitely decided to accept the latest level, as this will cause disk changes that are incompatible with previous levels of GPFS.

   For more information about the **mmchfs** command, see the *GPFS: Administration and Programming Reference*.

GPFS error messages for unsuccessful GPFS commands

If message **6027-538** is returned from the **mmcrfs** command, verify that the disk descriptors are specified correctly and that all named disks exist and are online. Issue the **mmlsnsd** command to check the disks.

6027-538
Error accessing disks.

If the daemon failed while running the command, you will see message **6027-663** Follow the procedures in "GPFS daemon went down" on page 48.

6027-663
Lost connection to file system daemon.

If the daemon was not running when you issued the command, you will see message **6027-665** Follow the procedures in “GPFS daemon will not come up” on page 45.

6027-665
Failed to connect to file system daemon: *errorString*

When GPFS commands are unsuccessful, the system may display information similar to these error messages:

6027-1626
Command is not supported in the *type* environment.

6027-1627
The following nodes are not aware of the configuration server change: *nodeList*. Do not start GPFS on the above nodes until the problem is resolved.

Application program errors

When receiving application program errors, there are various courses of action to take.

Follow these steps to help resolve application program errors:

1. Loss of file system access usually appears first as an error received by an application. Such errors are normally encountered when the application tries to access an unmounted file system.

   The most common reason for losing access to a single file system is a failure somewhere in the path to a large enough number of disks to jeopardize your data if operation continues. These errors may be reported in the the operating system error log on any node because they are logged in the first node to detect the error. Check all error logs for errors.

   The **mmlsmount** command can be used to determine the nodes that have successfully mounted a file system.
2. There are several cases where the state of a given disk subsystem will prevent access by GPFS. This will be seen by the application as I/O errors of various types and will be reported in the error logs as **MMFS_SYSTEM_UNMOUNT** or **MMFS_DISKFAIL** records. This state can be found by issuing the **mmlsdisk** command.

   For an IBM Virtual Shared Disk, this state can be found by running the **lsvsd -l** command. An IBM Virtual Shared Disk that is indicated as STP, or SUS is inaccessible to GPFS. Locate the appropriate IBM Reliable Scalable Cluster Technology (RSCT) document at: publib.boulder.ibm.com/clresctr/windows/public/rsctbooks.html and search for *restoration of access*. Follow the directions specified.
3. If allocation of data blocks or files (which quota limits should allow) fails, issue the **mmlsquota** command for the user, group or fileset.

   If filesets are involved, use these steps to determine which fileset was being accessed at the time of the failure:

   a. From the error messages generated, obtain the path name of the file being accessed.
   b. Use this **mmlsattr -L** command to obtain the fileset name.

      ```
      mmlsquota -e -j 'mmlsattr -L . | perl -ne 'print "$1\n" if /fileset name:\s*(.*)$/''
      ```

      The system produces output similar to:

      ```
            Block Limits                                    |     File Limits
      Filesystem type    KB  quota limit in_doubt grace |  files quota limit in_doubt grace  Remarks
      fs1        FILESET 2152    0      0     0    none |   250      0   250    0     none
      ```

   The **mmlsquota** output when checking the user and group quota is similar. If usage is equal to or approaching the hard limit, or if the grace period has expired, make sure that no quotas are lost by checking *in doubt* values.

   If quotas are exceeded in the *in doubt* category, run the **mmcheckquota** command. For more information, see “The mmcheckquota command” on page 30.

Note: There is no way to force GPFS nodes to relinquish all their local shares in order to check for lost quotas. This can only be determined by running the **mmcheckquota** command immediately after mounting the file system, and before any allocations are made. In this case, the value *in doubt* is the amount lost.

To display the latest quota usage information, use the **-e** option on either the **mmlsquota** or the **mmrepquota** commands. Remember that the **mmquotaon** and **mmquotaoff** commands do not enable and disable quota management. These commands merely control enforcement of quota limits. Usage continues to be counted and recorded in the quota files regardless of enforcement.

Reduce quota usage by deleting or compressing files or moving them out of the file system. Consider increasing quota limit.

GPFS error messages for application program errors

Application program errors can be associated with these GPFS message numbers:

6027-506
program: *loadFile* is already loaded at *address*.

6027-695
File system is read-only.

6027-1324
Unable to write new descriptor to the file *fileName*. Ensure the file system has enough space and retry.

Additionally, if you use IBM Virtual Shared Disks, you may receive these errors:

- `write failed: I/O error OR read failed: I/O error`

  A disk might be broken or an IBM Virtual Shared Disk might be stopped, or suspended. Check the disk state using the **lsvsd -l** command. A state of `ACT` (active) indicates a communicating disk.
- `write failed: Connection timed out OR read failed: Connection timed out`

  Consider tuning or evaluating the IBM Virtual Shared Disk configuration, including the number and size of buddy buffers. Locate the appropriate IBM Reliable Scalable Cluster Technology (RSCT) document at: publib.boulder.ibm.com/clresctr/windows/public/rsctbooks.html for more information.

Troubleshooting Windows

Use this information to help troubleshoot Windows problems.

The topics that follow apply to Windows 2003 R2. If you are having a problem not covered here, try using the Interop Community Forums at: http://www.interopsystems.com/community/default.aspx.

ssh-keysign error

There is a bug in OpenSSH (as of version 4.6.0.1.2 – a.k.a. OpenSSH_4.6p1) that causes the following error to be displayed for accounts other than the one that did the installation.

```
$ ssh bear101
/usr/local/libexec/ssh-keysign: error in loading shared libraries
libc.so.3.5: cannot open shared object file: No such file or directory
ssh_keysign: no reply
key_sign failed
Welcome to the SUA utilities.

DISPLAY=localhost:0.0
```

A work around is to comment out this line in **/usr/local/etc/ssh_config**:

```
#HostbasedAuthentication yes
```

Home and .ssh directory ownership and permissions

Make sure the user owns their home directory, which is not normally the case on Windows. They should also own ~/.ssh and the files it contains. Here is an example of file attributes that do work:

```
bash-3.00$ ls -l -d ~
drwx------  1 demyn  Domain Users  0 Dec  5 11:53 /dev/fs/D/Users/demyn
bash-3.00$ ls -l -d ~/.ssh
drwx------  1 demyn  Domain Users  0 Oct 26 13:37 /dev/fs/D/Users/demyn/.ssh
bash-3.00$ ls -l ~/.ssh
total 11
drwx------  1 demyn  Domain Users     0 Oct 26 13:37 .
drwx------  1 demyn  Domain Users     0 Dec  5 11:53 ..
-rw-r--r--  1 demyn  Domain Users   603 Oct 26 13:37 authorized_keys2
-rw-------  1 demyn  Domain Users   672 Oct 26 13:33 id_dsa
-rw-r--r--  1 demyn  Domain Users   603 Oct 26 13:33 id_dsa.pub
-rw-r--r--  1 demyn  Domain Users  2230 Nov 11 07:57 known_hosts
bash-3.00$
```

Problems running as Administrator

We have found there are problems using SSH when running as the domain **Administrator** user. These issues do not apply to other accounts, even if they are members of the **Administrators** group.

Chapter 7. GPFS file system problems

Suspect a GPFS file system problem when a file system will not mount or unmount.

You can also suspect a file system problem if a file system unmounts unexpectedly, or you receive an error message indicating that file system activity can no longer continue due to an error, and the file system is being unmounted to preserve its integrity. Record all error messages and log entries that you receive relative to the problem, making sure that you look on all affected nodes for this data.

These are some of the error encountered with GPFS file systems:

File system will not mount

There are indications leading you to the conclusion that your file system will not mount and courses of action you can take to correct the problem.

Some of those indications include:

- On performing a manual mount of the file system, you get errors from either the operating system or GPFS.
- If the file system was created with the option of an automatic mount, you will have failure return codes in the **MMFS log**.
- Your application cannot access the data it needs. Check the **MMFS log** for messages.
- Return codes or error messages from the **mmmount** command.
- The **mmlsmount** command indicates that the file system is not mounted on certain nodes.

If your file system will not mount, follow these steps:

1. On a quorum node in the cluster that owns the file system, verify that quorum includes this node. Check the **MMFS log** to see if a *mmfsd ready* message has been logged, and that no errors were reported on this or other nodes.
2. Verify that a conflicting command is not running. This applies only to the cluster that owns the file system. However, other clusters would be prevented from mounting the file system if a conflicting command is running in the cluster that owns the file system.

   For example, a **mount** command may not be issued while the **mmfsck** command is running. The **mount** command may not be issued until the conflicting command completes. Note that interrupting the **mmfsck** command is not a solution because the file system will not be mountable until the command completes. Try again after the conflicting command has completed.

3. Verify that sufficient disks are available to access the file system by issuing the **mmlsdisk** command. GPFS requires a minimum number of disks to find a current copy of the core metadata. If sufficient disks cannot be accessed, the mount will fail. The corrective action is to fix the path to the disk. See "NSD and underlying disk subsystem failures" on page 89.

   Missing disks can also cause GPFS to be unable to find critical metadata structures. The output of the **mmlsdisk** command will show any unavailable disks. If you have not specified metadata replication, the failure of one disk may result in your file system being unable to mount. If you have specified metadata replication, it will require two disks in different failure groups to disable the entire file system. If there are down disks, issue the **mmchdisk start** command to restart them and retry the mount.

   For a remote file system, **mmlsdisk** provides information about the disks of the file system. However **mmchdisk** must be run from the cluster that owns the file system.

   If there are no disks down, and your file system includes NSDs created on virtual shared disks, look on the virtual shared disk servers for operating system error log entries. Identify any virtual shared disk that has caused an error log entry to be generated. Locate the appropriate IBM Reliable Scalable Cluster Technology (RSCT) document at: publib.boulder.ibm.com/clresctr/windows/public/rsctbooks.html and search on *diagnosing IBM Virtual Shared Disk problems*. Follow the problem determination and repair actions specified. Look on the network shared disk servers for error log entries. Identify any network shared disk that has caused an error log entry to be generated.

   If there are no disks down, you can also look locally for error log reports, and follow the problem determination and repair actions specified in your disk vendor problem determination guide. If the disk has failed, follow the procedures in "NSD and underlying disk subsystem failures" on page 89.
4. Verify that communication paths to the other nodes are available. The lack of communication paths between all nodes in the cluster may impede contact with the file system manager.
5. Verify that the file system is not already mounted. Issue the **mount** command.
6. Verify that the GPFS daemon on the file system manager is available. Run the **mmlsmgr** command to determine which node is currently assigned as the file system manager. Run a trivial data access command such as an **ls** on the mount point directory. If the command fails, see "GPFS daemon went down" on page 48.
7. Check to see if the mount point directory exists and that there is an entry for the file system in the **/etc/fstab** file (for Linux) or **/etc/filesystems** file (for AIX). The device name for a file system mount point will be listed in column one of the **/etc/fstab** entry or as a **dev=** attribute in the **/etc/filesystems** stanza entry. A corresponding device name must also appear in the **/dev** file system.

   If any of these elements are missing, an update to the configuration information may not have been propagated to this node. Issue the **mmrefresh** command to rebuild the configuration information on the node and reissue the **mount** command.

   Do not add GPFS file system information to **/etc/filesystems** (for AIX) or **/etc/fstab** (for Linux) directly. If after running **mmrefresh -f** the file system information is still missing from **/etc/filesystems** (for AIX) or **/etc/fstab** (for Linux) contact the IBM Support Center.
8. Check the number of file systems that are already mounted. There is a maximum number of 256 mounted file systems for a GPFS cluster. Remote file systems are included in this number.
9. If you issue **mmchfs -V compat**, it enables backwardly-compatible format changes only. Nodes in remote clusters that were able to mount the file system before will still be able to do so.

   If you issue **mmchfs -V full**, it enables all new functions that require different on-disk data structures. Nodes in remote clusters running an older GPFS version will no longer be able to mount the file system. If there are any nodes running an older GPFS version that have the file system mounted at the time this command is issued, the **mmchfs** command will fail. For more information about completing the migration to a new level of GPFS, see the *General Parallel File System: Concepts, Planning, and Installation Guide*.

   All nodes that access the file system must be upgraded to the same level of GPFS. Check for the possibility that one or more of the nodes was accidently left out of an effort to upgrade a multi-node

system to a new GPFS release. If you need to return to the earlier level of GPFS, you must re-create the file system from the backup medium and restore the content in order to access it.

10. If DMAPI is enabled for the file system, ensure that a data management application is started and has set a disposition for the mount event. Refer to *General Parallel File System: Data Management API Guide* and the user's guide from your data management vendor.

    The data management application must be started in the cluster that owns the file system. If the application is not started, other clusters will not be able to mount the file system. Remote mounts of DMAPI managed file systems may take much longer to complete than those not managed by DMAPI.

11. Issue the **mmlsfs** command to check whether the automatic mount option has been specified. If automatic mount option is expected, check the **MMFS log** in the cluster that owns and serves the file system, for progress reports indicating:

    ```
    starting ...
    mounting ...
    mounted ....
    ```

12. If quotas are enabled, check if there was an error while reading quota files. See "MMFS_QUOTA" on page 4.
13. Verify the **maxblocksize** configuration parameter. If **maxblocksize** is less than the blocksize of the file system you are attempting to mount, you will not be able to mount it.

GPFS error messages for file system mount problems

6027-304
Remount failed for *device* after daemon restart.

6027-419
Failed to read a file system descriptor.

6027-549
Failed to open *name*.

6027-580
Unable to access vital system metadata, too many disks are unavailable.

6027-645
Attention: mmcommon getEFOptions *fileSystem* failed. Checking *fileName*.

Error numbers specific to GPFS application calls when a file system mount is not successful

When a mount of a file system is not successful, GPFS may report these error numbers in the operating system error log or return them to an application:

ENO_QUOTA_INST = 237, No Quota management enabled.
To enable quotas for the file system, issue the **mmchfs -Q yes** command. To disable quotas for the file system issue the **mmchfs -Q no** command.

Automount file system will not mount

If an automount fails when you **cd** into the mount point directory, first check that the file system in question is of automount type. Use the **mmlsfs** command for local file systems. Use the **mmremotefs show** command for remote file systems.

Steps to follow if automount fails to mount on Linux

On Linux, perform these steps:

1. Verify that the GPFS file system mount point is actually a symbolic link to a directory in the automountdir directory. If **automountdir=/gpfs/autmountdir** then the mount point **/gpfs/gpfs66** would be a symbolic link to **/gpfs/automountdir/gpfs66**.
   a. First, verify that GPFS is up and running.
   b. Use the **mmlsconfig** command to verify the automountdir directory. The default automountdir is named **/gpfs/automountdir**. If the GPFS file system mount point is not a symbolic link to the GPFS automountdir directory, then accessing the mount point will not cause the automounter to mount the file system.
   c. If the command **/bin/ls -ld** of the mount point shows a directory, then run the command **mmrefresh -f**. If the directory is empty, the command **mmrefresh -f** will remove the directory and create a symbolic link.

      If the directory is not empty, you need to move or remove the files contained in that directory, or change the mount point of the file system. For a local file system, use the **mmchfs** command. For a remote file system, use the **mmremotefs** command.
   d. Once the mount point directory is empty, run the **mmrefresh -f** command.
2. Verify that the **autofs** mount has been established. Issue this command:

   ```
   mount | grep automount
   ```

   Output should be similar to this:

   ```
   automount(pid20331) on /gpfs/automountdir type autofs (rw,fd=5,pgrp=20331,minproto=2,maxproto=3)
   ```

   For RHEL5, verify the following line is in the default master map file (/etc/auto.master):

   ```
   /gpfs/automountdir program:/usr/lpp/mmfs/bin/mmdynamicmap
   ```

   For example, issue:

   ```
   /root # grep mmdynamicmap /etc/auto.master
   ```

   Output should be similar to this:

   ```
   /gpfs/automountdir program:/usr/lpp/mmfs/bin/mmdynamicmap
   ```

   This is an **autofs** program map, and there will be a single mount entry for all GPFS automounted file systems. The symbolic link points to this directory, and access through the symbolic link triggers the mounting of the target GPFS file system. To create this GPFS **autofs** mount, issue the **mmcommon startAutomounter** command, or stop and restart GPFS using the **mmshutdown** and **mmstartup** commands.
3. Verify that the automount daemon is running. Issue this command:

   ```
   ps -ef | grep automount
   ```

   Output should be similar to this:

   ```
   root 5116 1 0 Jun25 pts/0 00:00:00 /usr/sbin/automount /gpfs/automountdir program
       /usr/lpp/mmfs/bin/mmdynamicmap
   ```

   For RHEL5, verify that the **autofs** daemon is running. Issue this command:

   ```
   ps -ef | grep automount
   ```

   Output should be similar to this:

   ```
   root 22646 1 0 01:21 ? 00:00:02 automount
   ```

   To start the automount daemon, issue the **mmcommon startAutomounter** command, or stop and restart GPFS using the **mmshutdown** and **mmstartup** commands.
4. Verify that the mount command was issued to GPFS by examining the **MMFS log**. You should see something like this:

   ```
   Mon Jun 25 11:33:03 2004: Command: mount gpfsx2.kgn.ibm.com:gpfs55 5182
   ```
5. Examine **/var/log/messages** for **autofs** error messages.

This is an example of what you might see if the remote file system name does not exist.

```
Jun 25 11:33:03 linux automount[20331]: attempting to mount entry /gpfs/automountdir/gpfs55
Jun 25 11:33:04 linux automount[28911]: >> Failed to open gpfs55.
Jun 25 11:33:04 linux automount[28911]: >> No such device
Jun 25 11:33:04 linux automount[28911]: >> mount: fs type gpfs not supported by kernel
Jun 25 11:33:04 linux automount[28911]: mount(generic): failed to mount /dev/gpfs55 (type gpfs)
    on /gpfs/automountdir/gpfs55
```

6. After you have established that GPFS has received a mount request from **autofs** (Step 4 on page 64) and that mount request failed (Step 5 on page 64), issue a mount command for the GPFS file system and follow the directions in "File system will not mount" on page 61.

Steps to follow if automount fails to mount on AIX

On AIX, perform these steps:

1. First, verify that GPFS is up and running.
2. Verify that GPFS has established **autofs** mounts for each automount file system. Issue the following command

   ```
   mount | grep autofs
   ```

   The output is similar to this:

   ```
   /var/mmfs/gen/mmDirectMap /gpfs/gpfs55 autofs Jun 25 15:03 ignore
   /var/mmfs/gen/mmDirectMap /gpfs/gpfs88 autofs Jun 25 15:03 ignore
   ```

   These are direct mount **autofs** mount entries. Each GPFS automount file system will have an **autofs** mount entry. These **autofs** direct mounts allow GPFS to mount on the GPFS mount point. To create any missing GPFS **autofs** mounts, issue the **mmcommon startAutomounter** command, or stop and restart GPFS using the **mmshutdown** and **mmstartup** commands.
3. Verify that the **autofs** daemon is running. Issue this command:

   ```
   ps -ef | grep automount
   ```

   Output is similar to this:

   ```
   root 9820 4240 0 15:02:50 - 0:00 /usr/sbin/automountd
   ```

   To start the automount daemon, issue the **mmcommon startAutomounter** command, or stop and restart GPFS using the **mmshutdown** and **mmstartup** commands.
4. Verify that the mount command was issued to GPFS by examining the **MMFS log**. You should see something like this:

   ```
   Mon Jun 25 11:33:03 2007: Command: mount gpfsx2.kgn.ibm.com:gpfs55 5182
   ```
5. Since the **autofs** daemon logs status using **syslogd**, examine the **syslogd** log file for status information from **automountd**. Here is an example of a failed automount request:

   ```
   Jun 25 15:55:25 gpfsa1 automountd [9820 ] :mount of /gpfs/gpfs55:status 13
   ```
6. After you have established that GPFS has received a mount request from **autofs** (Step 4) and that mount request failed (Step 5), issue a mount command for the GPFS file system and follow the directions in "File system will not mount" on page 61.
7. If automount fails for a non-GPFS file system and you are using file **/etc/auto.master**, use file **/etc/auto_master** instead. Add the entries from **/etc/auto.master** to **/etc/auto_master** and restart the automount daemon.

Remote file system will not mount

These are some of the errors encountered when mounting remote file systems:

- "Remote file system will not mount due to differing GPFS cluster security configurations" on page 66
- "Remote file system I/O fails with the "Function not implemented" error message when UID mapping is enabled" on page 66
- "Cannot resolve contact node address" on page 67

- "The remote cluster name does not match the cluster name supplied by the mmremotecluster command" on page 67
- "Contact nodes down or GPFS down on contact nodes" on page 67
- "GPFS is not running on the local node" on page 68
- "The NSD disk does not have a NSD server specified and the mounting cluster does not have direct access to the disks" on page 68
- "The cipherList option has not been set properly" on page 68

Remote file system I/O fails with the "Function not implemented" error message when UID mapping is enabled

When user ID (UID) mapping in a multi-cluster environment is enabled, certain kinds of mapping infrastructure configuration problems might result in I/O requests on a remote file system failing:

```
ls -l /fs1/testfile
ls: /fs1/testfile: Function not implemented
```

To troubleshoot this error, verify the following configuration details:

1. That **/var/mmfs/etc/mmuid2name** and **/var/mmfs/etc/mmname2uid** helper scripts are present and executable on all nodes in the local cluster and on all quorum nodes in the file system home cluster, along with any data files needed by the helper scripts.
2. That UID mapping is enabled in both local cluster and remote file system home cluster configuration.
3. That UID mapping helper scripts are working correctly.

For more information about configuring UID mapping, see the "UID Mapping for GPFS in a Multi-cluster Environment" white paper at: http://www-03.ibm.com/systems/clusters/software/whitepapers/uid_gpfs.html.

Remote file system will not mount due to differing GPFS cluster security configurations

A mount command fails with a message similar to this:

```
Cannot mount gpfsxx2.ibm.com:gpfs66: Host is down.
```

The **MMFS log** on the cluster issuing the mount command should have entries similar to these:

```
There is more information in the log file /var/adm/ras/mmfs.log.latest
Mon Jun 25 16:39:27 2007: Waiting to join remote cluster gpfsxx2.ibm.com
Mon Jun 25 16:39:27 2007: Command: mount gpfsxx2.ibm.com:gpfs66 30291
Mon Jun 25 16:39:27 2007: The administrator of 199.13.68.12 gpfslx2 requires
secure connections.  Contact the administrator to obtain the target clusters
key and register the key using "mmremotecluster update".
Mon Jun 25 16:39:27 2007: A node join was rejected.  This could be due to
incompatible daemon versions, failure to find the node
in the configuration database, or no configuration manager found.
Mon Jun 25 16:39:27 2007: Failed to join remote cluster gpfsxx2.ibm.com
Mon Jun 25 16:39:27 2007: Command err 693: mount gpfsxx2.ibm.com:gpfs66 30291
```

The **MMFS log** file on the cluster that owns and serves the file system will have an entry indicating the problem as well, similar to this:

```
Mon Jun 25 16:32:21 2007: Kill accepted connection from 199.13.68.12 because security is required, err 74
```

To resolve this problem, contact the administrator of the cluster that owns and serves the file system to obtain the key and register the key using **mmremotecluster** command.

The SHA digest field of the **mmauth show** and **mmremotecluster** commands may be used to determine if there is a key mismatch, and on which cluster the key should be updated. For more information on the SHA digest, see "The SHA digest" on page 33.

Cannot resolve contact node address

The following error may occur if the contact nodes for **gpfsyy2.ibm.com** could not be resolved. You would expect to see this if your DNS server was down, or the contact address has been deleted.

```
Mon Jun 25 15:24:14 2007: Command: mount gpfsyy2.ibm.com:gpfs14 20124
Mon Jun 25 15:24:14 2007: Host 'gpfs123.ibm.com' in gpfsyy2.ibm.com is not valid.
Mon Jun 25 15:24:14 2007: Command err 2: mount gpfsyy2.ibm.com:gpfs14 20124
```

To resolve the problem, correct the contact list and try the mount again.

The remote cluster name does not match the cluster name supplied by the mmremotecluster command

A mount command fails with a message similar to this:

```
Cannot mount gpfslx2:gpfs66: Network is unreachable
```

and the **MMFS log** contains message similar to this:

```
Mon Jun 25 12:47:18 2007: Waiting to join remote cluster gpfslx2
Mon Jun 25 12:47:18 2007: Command: mount gpfslx2:gpfs66 27226
Mon Jun 25 12:47:18 2007: Failed to join remote cluster gpfslx2
Mon Jun 25 12:47:18 2007: Command err 719: mount gpfslx2:gpfs66 27226
```

Perform these steps:

1. Verify that the remote cluster name reported by the **mmremotefs show** command is the same name as reported by the **mmlscluster** command from one of the contact nodes.
2. Verify the list of contact nodes against the list of nodes as shown by the **mmlscluster** command from the remote cluster.

In this example, the correct cluster name is **gpfslx2.ibm.com** and not **gpfslx2**

```
mmlscluster
```

Output is similar to this:

```
GPFS cluster information
========================
  GPFS cluster name:         gpfslx2.ibm.com
  GPFS cluster id:           649437685184692490
  GPFS UID domain:           gpfslx2.ibm.com
  Remote shell command:      /usr/bin/ssh
  Remote file copy command:  /usr/bin/scp

GPFS cluster configuration servers:
-----------------------------------
  Primary server:    gpfslx2.ibm.com
  Secondary server:  (none)

 Node  Daemon node name   IP address      Admin node name       Designation
--------------------------------------------------------------------------
   1   gpfslx2            198.117.68.68   gpfslx2.ibm.com       quorum
```

Contact nodes down or GPFS down on contact nodes

A mount command fails with a message similar to this:

```
GPFS: 6027-510 Cannot mount /dev/gpfs22 on /gpfs22: A remote host did not respond
    within the timeout period.
```

On AIX, the **MMFS log** will have entries similar to this:

```
Mon Jun 25 13:11:14 2007: Command: mount gpfslx22:gpfs22 19004
Mon Jun 25 13:11:14 2007: Waiting to join remote cluster gpfslx22
Mon Jun 25 13:11:15 2007: Connecting to 199.13.68.4 gpfslx22
Mon Jun 25 13:16:36 2007: Failed to join remote cluster gpfslx22
Mon Jun 25 13:16:36 2007: Command err 78: mount gpfslx22:gpfs22 19004
```

On Linux, the **MMFS log** will have entries similar to this:

```
Mon Jun 25 13:11:15 2007: Command: mount gpfslx22:gpfs22 19004
Mon Jun 25 13:11:15 2007: Waiting to join remote cluster gpfslx22
Mon Jun 25 13:11:16 2007: Connecting to 199.13.67.8 gpfslx22
Mon Jun 25 13:16:37 2007: Failed to join remote cluster gpfslx22
Mon Jun 25 13:16:37 2007: Command err 110: mount gpfslx22:gpfs22 19004
```

To resolve the problem, use the **mmremotecluster show** command and verify that the cluster name matches the remote cluster and the contact nodes are valid nodes in the remote cluster. Verify that GPFS is active on the contact nodes in the remote cluster. Another way to resolve this problem is to change the contact nodes using the **mmremotecluster update** command.

GPFS is not running on the local node

A mount command fails with a message similar to this:

```
mount: fs type gpfs not supported by kernel
```

Follow your procedures for starting GPFS on the local node.

The NSD disk does not have a NSD server specified and the mounting cluster does not have direct access to the disks

A file system mount fails with a message similar to this:

```
Failed to open gpfs66.
No such device
mount: Stale NFS file handle
Some file system data are inaccessible at this time.
Check error log for additional information.
Cannot mount gpfslx2.ibm.com:gpfs66: Stale NFS file handle
```

The **MMFS log** will contain information similar to this:

```
Mon Jun 25 14:10:46 2007: Command: mount gpfslx2.ibm.com:gpfs66 28147
Mon Jun 25 14:10:47 2007: Waiting to join remote cluster gpfslx2.ibm.com
Mon Jun 25 14:10:47 2007: Connecting to 199.13.68.4 gpfslx2
Mon Jun 25 14:10:47 2007: Connected to 199.13.68.4 gpfslx2
Mon Jun 25 14:10:47 2007: Joined remote cluster gpfslx2.ibm.com
Mon Jun 25 14:10:48 2007: Global NSD disk, gpfs1nsd, not found.
Mon Jun 25 14:10:48 2007: Disk failure.  Volume gpfs66. rc = 19. Physical volume gpfs1nsd.
Mon Jun 25 14:10:48 2007: File System gpfs66 unmounted by the system with return code 19 reason code 0
Mon Jun 25 14:10:48 2007: No such device
Mon Jun 25 14:10:48 2007: Command err 666: mount gpfslx2.ibm.com:gpfs66 28147
```

To resolve the problem, the cluster that owns and serves the file system must define one or more NSD servers.

The cipherList option has not been set properly

Another reason for remote mount to fail is if **cipherList** is not set to a valid value. A mount command would fail with messages similar to this:

```
6027-510 Cannot mount /dev/dqfs1 on /dqfs1: A remote host is not available.
```

The **MMFS log** would contain messages similar to this:

```
Wed Jul 18 16:11:20.496 2007: Command: mount remote.cluster:fs3 655494
Wed Jul 18 16:11:20.497 2007: Waiting to join remote cluster remote.cluster
Wed Jul 18 16:11:20.997 2007: Remote mounts are not enabled within this cluster. \
See the Advanced Administration Guide for instructions. In particular ensure keys have been \
generated and a cipherlist has been set.
Wed Jul 18 16:11:20.998 2007: A node join was rejected.  This could be due to
incompatible daemon versions, failure to find the node
in the configuration database, or no configuration manager found.
Wed Jul 18 16:11:20.999 2007: Failed to join remote cluster remote.cluster
Wed Jul 18 16:11:20.998 2007: Command: err 693: mount remote.cluster:fs3 655494
Wed Jul 18 16:11:20.999 2007: Message failed because the destination node refused the connection.
```

The **mmchconfig cipherlist=AUTHONLY** command must be run on both the cluster that owns and controls the file system, and the cluster that is attempting to mount the file system.

Mount failure due to client nodes joining before NSD servers are online

If a client node joins the GPFS cluster and attempts file system access prior to the file system's NSD servers being active, the mount fails. This is especially true when automount is used. This situation can occur during cluster startup, or any time that an NSD server is brought online with client nodes already active and attempting to mount a file system served by the NSD server.

The file system mount failure produces a message similar to this:

```
Mon Jun 25 11:23:34 EST 2007: mmmount: Mounting file systems ...
No such device
Some file system data are inaccessible at this time.
Check error log for additional information.
After correcting the problem, the file system must be unmounted and then
mounted again to restore normal data access.
Failed to open fs1.
No such device
Some file system data are inaccessible at this time.
Cannot mount /dev/fs1 on /fs1: Missing file or filesystem
```

The **MMFS log** contains information similar to this:

```
Mon Jun 25 11:23:54 2007: Command: mount fs1 32414
Mon Jun 25 11:23:58 2007: Disk failure. Volume fs1. rc = 19. Physical volume sdcnsd.
Mon Jun 25 11:23:58 2007: Disk failure. Volume fs1. rc = 19. Physical volume sddnsd.
Mon Jun 25 11:23:58 2007: Disk failure. Volume fs1. rc = 19. Physical volume sdensd.
Mon Jun 25 11:23:58 2007: Disk failure. Volume fs1. rc = 19. Physical volume sdgnsd.
Mon Jun 25 11:23:58 2007: Disk failure. Volume fs1. rc = 19. Physical volume sdhnsd.
Mon Jun 25 11:23:58 2007: Disk failure. Volume fs1. rc = 19. Physical volume sdinsd.
Mon Jun 25 11:23:58 2007: File System fs1 unmounted by the system with return code 19
reason code 0
Mon Jun 25 11:23:58 2007: No such device
Mon Jun 25 11:23:58 2007: File system manager takeover failed.
Mon Jun 25 11:23:58 2007: No such device
Mon Jun 25 11:23:58 2007: Command: err 52: mount fs1 32414
Mon Jun 25 11:23:58 2007: Missing file or filesystem
```

Two **mmchconfig** command options are used to specify the amount of time for GPFS mount requests to wait for an NSD server to join the cluster:

nsdServerWaitTimeForMount
: Specifies the number of seconds to wait for an NSD server to come up at GPFS cluster startup time, after a quorum loss, or after an NSD server failure.

  Valid values are between 0 and 1200 seconds. The default is 300. The interval for checking is 10 seconds. If **nsdServerWaitTimeForMount** is 0, **nsdServerWaitTimeWindowOnMount** has no effect.

nsdServerWaitTimeWindowOnMount
: Specifies a time window to determine if quorum is to be considered *recently formed*.

  Valid values are between 1 and 1200 seconds. The default is 600. If **nsdServerWaitTimeForMount** is 0, **nsdServerWaitTimeWindowOnMount** has no effect.

The GPFS daemon need not be restarted in order to change these values. The scope of these two operands is the GPFS cluster. The **-N** flag can be used to set different values on different nodes. In this case, the settings on the file system manager node take precedence over the settings of nodes trying to access the file system.

When a node rejoins the cluster (after it was expelled, experienced a communications problem, lost quorum, or other reason for which it dropped connection and rejoined), that node resets all the failure times that it knows about. Therefore, when a node rejoins it sees the NSD servers as never having failed. From the node's point of view, it has rejoined the cluster and old failure information is no longer relevant.

GPFS checks the cluster formation criteria first. If that check falls outside the window, GPFS then checks for NSD server fail times being within the window.

File system will not unmount

There are indications leading you to the conclusion that your file system will not unmount and a course of action to correct the problem.

Those indications include:

- Return codes or error messages indicate the file system will not unmount.
- The **mmlsmount** command indicates that the file system is still mounted on one or more nodes.
- Return codes or error messages from the **mmumount** command.

If your file system will not unmount, follow these steps:

1. If you get an error message similar to:

   ```
   umount: /gpfs1: device is busy
   ```

   the file system will not unmount until all processes are finished accessing it. If **mmfsd** is up, the processes accessing the file system can be determined. See "The lsof command" on page 22. These processes can be killed with the command:

   ```
   lsof filesystem | grep -v COMMAND | awk '{print $2}' | xargs kill -9
   ```

   If **mmfsd** is not operational, the **lsof** command will not be able to determine which processes are still accessing the file system.

   For Linux nodes it is possible to use the **/proc** pseudo file system to determine current file access. For each process currently running on the system, there is a subdirectory **/proc/***pid***/fd**, where *pid* is the numeric process ID number. This subdirectory is populated with symbolic links pointing to the files that this process has open. You can examine the contents of the **fd** subdirectory for all running processes, manually or with the help of a simple script, to identify the processes that have open files in GPFS file systems. Terminating all of these processes may allow the file system to unmount successfully.
2. Verify that there are no disk media failures.

   If using NSDs created on virtual shared disks, look on the virtual shared disk servers for error logs. Identify any virtual shared disk that has generated an error log. Locate the appropriate IBM Reliable Scalable Cluster Technology (RSCT) document at: publib.boulder.ibm.com/clresctr/windows/public/rsctbooks.html and search on *diagnosing IBM Virtual Shared Disk problems*. Follow the problem determination and repair actions specified.

   Look on the NSD server node for error log entries. Identify any NSD server node that has generated an error log entry. See "Disk media failure" on page 94 for problem determination and repair actions to follow.
3. If the file system *must* be unmounted, you can force the unmount by issuing the **mmumount -f** command:

   Note:

   a. See "File system forced unmount" on page 71 for the consequences of doing this.
   b. Before forcing the unmount of the file system, issue the **lsof** command and close any files that are open.

c. On Linux, you might encounter a situation where a GPFS file system cannot be unmounted, even if you issue the **umount -f** command. In this case, you must reboot the node to clear the condition. You can also try the system **umount** command before you reboot. For example:

```
umount -f /filesystem
```

4. If a file system that is mounted by a remote cluster needs to be unmounted, you can force the unmount by issuing the command:

```
mmunmount filesystemname -f -C remoteclustername
```

File system forced unmount

There are indications that lead you to the conclusion that your file system has been forced to unmount and various courses of action that you can take to correct the problem.

Those indications are:

- Forced unmount messages in the **MMFS log**.
- Your application no longer has access to data.
- Your application is getting `ESTALE` or `ENOENT` return codes.
- Multiple unsuccessful attempts to appoint a file system manager may cause the cluster manager to unmount the file system everywhere.

  Such situations involve the failure of paths to disk resources from many, if not all, nodes. The underlying problem may be at the disk subsystem level, or lower. The error logs for each node that unsuccessfully attempted to appoint a file system manager will contain records of a file system unmount with an error that are either coded **212**, or that occurred when attempting to assume management of the file system. Note that these errors apply to a specific file system although it is possible that shared disk communication paths will cause the unmount of multiple file systems.
- File system unmounts with an error indicating too many disks are unavailable.

  The **mmlsmount** command can be used to determine which nodes currently have a given file system mounted.

If your file system has been forced to unmount, follow these steps:

1. With the failure of a single disk, if you have not specified multiple failure groups and replication of metadata, GPFS will not be able to continue because it cannot write logs or other critical metadata. If you have specified multiple failure groups and replication of metadata, the failure of multiple disks in the same failure group will put you in the same position. In either of these situations, GPFS will forcibly unmount the file system. This will be indicated in the error log by records indicating exactly which access failed, with a **MMFS_SYSTEM_UNMOUNT** record indicating the forced unmount.

   The user response to this is to take the needed actions to restore the disk access and issue the **mmchdisk** command to disks that are shown as down in the information displayed by the **mmlsdisk** command.
2. Internal errors in processing data on a single file system may cause loss of file system access. These errors may clear with the invocation of the **umount** command, followed by a remount of the file system, but they should be reported as problems to IBM.
3. If an **MMFS_QUOTA** error log entry containing `Error writing quota file...` is generated, the quota manager continues operation if the next write for the user, group, or fileset is successful. If not, further allocations to the file system will fail. Check the error code in the log and make sure that the disks containing the quota file are accessible. Run the **mmcheckquota** command. For more information, see "The mmcheckquota command" on page 30.

   If the file system must be repaired without quotas:

   a. Disable quota management by issuing the command:

   ```
   mmchfs Device -Q no
   ```

   b. Issue the **mmmount** command for the file system.

c. Make any necessary repairs and install the backup quota files.
d. Issue the **mmumount -a** command for the file system.
e. Restore quota management by issuing the **mmchfs** *Device* **-Q yes** command.
f. Run the **mmcheckquota** command with the **-u**, **-g**, and **-j** options. For more information, see "The mmcheckquota command" on page 30.
g. Issue the **mmmount** command for the file system.

4. If errors indicate that too many disks are unavailable, see "Additional failure group considerations."

Additional failure group considerations

There is a structure in GPFS called the *file system descriptor* that is initially written to every disk in the file system, but is replicated on a subset of the disks as changes to the file system occur, such as adding or deleting disks. Based on the number of failure groups and disks, GPFS creates between one and five replicas of the descriptor:

- If there are at least five different failure groups, five replicas are created.
- If there are at least three different disks, three replicas are created.
- If there are only one or two disks, a replica is created on each disk.

Once it is decided how many replicas to create, GPFS picks disks to hold the replicas, so that all replicas will be in different failure groups, if possible, to reduce the risk of multiple failures. In picking replica locations, the current state of the disks is taken into account. Stopped or suspended disks are avoided. Similarly, when a failed disk is brought back online, GPFS may modify the subset to rebalance the file system descriptors across the failure groups. The subset can be found by issuing the **mmlsdisk -L** command.

GPFS requires a majority of the replicas on the subset of disks to remain available to sustain file system operations:

- If there are at least five different failure groups, GPFS will be able to tolerate a loss of two of the five groups. If disks out of three different failure groups are lost, the file system descriptor may become inaccessible due to the loss of the majority of the replicas.
- If there are at least three different failure groups, GPFS will be able to tolerate a loss of one of the three groups. If disks out of two different failure groups are lost, the file system descriptor may become inaccessible due to the loss of the majority of the replicas.
- If there are fewer than three failure groups, a loss of one failure group may make the descriptor inaccessible.

  If the subset consists of three disks and there are only two failure groups, one failure group must have two disks and the other failure group has one. In a scenario that causes one entire failure group to disappear all at once, if the half of the disks that are unavailable contain the single disk that is part of the subset, everything stays up. The file system descriptor is moved to a new subset by updating the remaining two copies and writing the update to a new disk added to the subset. But if the downed failure group contains a majority of the subset, the file system descriptor cannot be updated and the file system has to be force unmounted.

  Introducing a third failure group consisting of a single disk that is used solely for the purpose of maintaining a copy of the file system descriptor can help prevent such a scenario. You can designate this disk by using the **descOnly** designation for disk usage on the disk descriptor. With the **descOnly** designation, the disk does not hold any of the other file system data or metadata and can be as small as 4 MB. See *NSD creation considerations* in *GPFS: Concepts, Planning, and Installation Guide* and *Establishing disaster recovery for your GPFS cluster* in *GPFS: Advanced Administration Guide*.

GPFS error messages for file system forced unmount problems

Indications there are not enough disks available:

6027-418
Inconsistent file system quorum. readQuorum=*value* writeQuorum=*value* quorumSize=*value*.

6027-419
Failed to read a file system descriptor.

Indications the file system has been forced to unmount:

6027-473
File System *fileSystem* unmounted by the system with return code *value* and reason code *value*.

6027-474
Recovery Log I/O Failed, unmounting *fileSystem*.

If your disks are IBM Virtual Shared Disks, locate the appropriate IBM Reliable Scalable Cluster Technology (RSCT) document at: publib.boulder.ibm.com/clresctr/windows/public/rsctbooks.html and search on *diagnosing IBM Virtual Shared Disk problems*. Follow the problem determination and repair actions specified.

Error numbers specific to GPFS application calls when a file system has been forced to unmount

When a file system has been forced to unmount, GPFS may report these error numbers in The operating system error log or return them to an application:

EPANIC = 666, A file system has been forcibly unmounted because of an error. Most likely due to the failure of one or more disks containing the last copy of metadata.
See “The operating system error log facility” on page 2 for details.

EALL_UNAVAIL = 218, A replicated read or write failed because none of the replicas were available.
Multiple disks in multiple failure groups are unavailable. Follow the procedures in Chapter 8, “GPFS disk problems,” on page 89 for unavailable disks.

Unable to determine whether a file system is mounted

Certain GPFS file system commands cannot be performed when the file system in question is mounted.

In certain failure situations, GPFS cannot determine whether the file system in question is mounted or not, and so cannot perform the requested command. In such cases, message **6027-1996** (Command was unable to determine whether the file system is mounted) is issued.

If you encounter this message, investigate why GPFS could not determine whether the file system was mounted, resolve the problem, and reissue the command. If you cannot determine or resolve the problem, you may be able to successfully run the command by first shutting down the GPFS daemon on all nodes of the cluster (using **mmshutdown -a**), thus ensuring that the file system is not mounted.

GPFS error messages for file system mount status

6027-1996
Command was unable to determine whether the file system is mounted.

Multiple file system manager failures

The correct operation of GPFS requires that one node per file system function as the file system manager at all times. This instance of GPFS has additional responsibilities for coordinating usage of the file system.

When the file system manager node fails, another file system manager is appointed in a manner that is not visible to applications except for the time required to switch over.

There are situations where it may be impossible to appoint a file system manager. Such situations involve the failure of paths to disk resources from many, if not all, nodes. In this event, the cluster manager nominates several host names to successively try to become the file system manager. If none succeed, the cluster manager unmounts the file system everywhere. See “NSD and underlying disk subsystem failures” on page 89.

The required action here is to address the underlying condition which caused the forced unmounts and then remount the file system. In most cases, this means correcting the path to the disks required by GPFS. If IBM Virtual Shared Disk or NSD disk servers are being used, the most common failure is the loss of access through the switch or communications networks. If SAN access is being used to all disks, the most common failure is the loss of connectivity through the SAN.

GPFS error messages for multiple file system manager failures

The inability to successfully appoint a file system manager after multiple attempts can be associated with both the error messages listed in “File system forced unmount” on page 71, as well as these additional messages:

- When a forced unmount occurred on all nodes:

  6027-635
  The current file system manager failed and no new manager will be appointed.
- If message **6027-636** is displayed, it means that there may be a disk failure. If you are using NSDs created on IBM Virtual Shared Disks, locate the appropriate IBM Reliable Scalable Cluster Technology (RSCT) document at: publib.boulder.ibm.com/clresctr/windows/public/rsctbooks.html and search on *diagnosing IBM Virtual Shared Disk problems*. Follow the problem determination and repair actions specified. See “NSD and underlying disk subsystem failures” on page 89 for NSD problem determination and repair procedures.

  6027-636
  Disks marked as stopped or offline.
- Message **6027-632** is the last message in this series of messages. See the accompanying messages:

  6027-632
  Failed to appoint a new manager for *file_system*.
- Message **6027-631** occurs on each attempt to appoint a new manager (see the messages on the referenced node for the specific reason as to why it failed):

  6027-631
  Failed to appoint node *nodeName* as manager for *fileSystem*.
- Message **6027-638** indicates which node had the original error (probably the original file system manager node):

  6027-638
  File system *fileSystem* unmounted by node *nodeName*.

Error numbers specific to GPFS application calls when file system manager appointment fails

When the appointment of a file system manager is unsuccessful after multiple attempts, GPFS may report these error numbers in error logs, or return them to an application:

ENO_MGR = 212, The current file system manager failed and no new manager could be appointed.
This usually occurs when a large number of disks are unavailable or when there has been a major network failure. Run **mmlsdisk** to determine whether disks have failed and take corrective action if they have by issuing the **mmchdisk** command.

Discrepancy between GPFS configuration data and the on-disk data for a file system

There is an indication leading you to the conclusion that there may be a discrepancy between the GPFS configuration data and the on-disk data for a file system.

You issue a disk command (for example, **mmadddisk**, **mmdeldisk**, or **mmrpldisk**) and receive the message:

6027-1290
GPFS configuration data for file system *fileSystem* may not be in agreement with the on-disk data for the file system. Issue the command:

```
mmcommon recoverfs fileSystem
```

Before a disk is added to or removed from a file system, a check is made that the GPFS configuration data for the file system is in agreement with the on-disk data for the file system. The above message is issued if this check was not successful. This may occur if an earlier GPFS disk command was unable to complete successfully for some reason. Issue the **mmcommon recoverfs** command to bring the GPFS configuration data into agreement with the on-disk data for the file system.

If running **mmcommon recoverfs** does not resolve the problem, contact the IBM Support Center.

Errors associated with storage pools, filesets and policies

When and error is suspected while working with storage pools, policies and filesets, check the relevant section in *General Parallel File System: Advanced Administration Guide* to ensure that your setup is correct.

When you are sure that your setup is correct, see if your problem falls into one of these categories:

A NO_SPACE error occurs when a file system is known to have adequate free space

A **ENOSPC** (**NO_SPACE**) message can be returned even if a file system has remaining space. The **NO_SPACE** error might occur even if the **df** command shows that the file system is not full.

The user might have a policy that writes data into a specific storage pool. When the user tries to create a file in that storage pool, it returns the **ENOSPC** error if the storage pool is full. The user next issues the **df**

command, which indicates that the file system is not full, because the problem is limited to the one storage pool in the user's policy. In order to see if a particular storage pool is full, the user must issue the **mmdf** command.

Here is a sample scenario:

1. The user has a policy rule that says files whose name contains the word 'tmp' should be put into storage pool **sp1** in the file system **fs1**. This command displays the rule:

   ```
   mmlspolicy fs1 -L
   ```

   The system produces output similar to this:

   ```
   /* This is a policy for GPFS file system fs1 */

   /* File Placement Rules */
   RULE SET POOL 'sp1' WHERE name like '%tmp%'
   RULE 'default' SET POOL 'system'
   /* End of Policy */
   ```

2. The user moves a file from the **/tmp** directory to **fs1** that has the word 'tmp' in the file name, meaning data of **tmpfile** should be placed in storage pool **sp1**:

   ```
   mv /tmp/tmpfile /fs1/
   ```

   The system produces output similar to this:

   ```
   mv: writing `/fs1/tmpfile': No space left on device
   ```

   This is an out-of-space error.

3. This command shows storage information for the file system:

   ```
   df |grep fs1
   ```

   The system produces output similar to this:

   ```
   /dev/fs1             280190976 140350976 139840000  51% /fs1
   ```

   This output indicates that the file system is only 51% full.

4. To query the storage usage for an individual storage pool, the user must issue the **mmdf** command.

   ```
   mmdf fs1
   ```

   The system produces output similar to this:

   ```
   disk                disk size  failure holds    holds              free KB             free KB
   name                    in KB    group metadata data        in full blocks        in fragments
   --------------- ------------- -------- -------- ----- -------------------- -------------------
   Disks in storage pool: system
   gpfs1nsd            140095488     4001 yes      yes       139840000 (100%)         19936 ( 0%)
                   -------------                         -------------------- -------------------
   (pool total)        140095488                               139840000 (100%)         19936 ( 0%)

   Disks in storage pool: sp1
   gpfs2nsd            140095488     4001 no       yes               0 (  0%)           248 ( 0%)
                   -------------                         -------------------- -------------------
   (pool total)        140095488                                     0 (  0%)           248 ( 0%)

                   =============                         ==================== ===================
   (data)              280190976                               139840000 ( 50%)         20184 ( 0%)
   (metadata)          140095488                               139840000 (100%)         19936 ( 0%)
                   =============                         ==================== ===================
   (total)             280190976                               139840000 ( 50%)         20184 ( 0%)

   Inode Information
   -----------------
   ```

```
Number of used inodes: 74
Number of free inodes: 137142
Number of allocated inodes: 137216
Maximum number of inodes: 150016
```

In this case, the user sees that storage pool **sp1** has 0% free space left and that is the reason for the **NO_SPACE** error message.

5. To resolve the problem, the user must change the placement policy file to avoid putting data in a full storage pool, delete some files in storage pool **sp1**, or add more space to the storage pool.

Negative values occur in the 'predicted pool utilizations', when some files are 'ill-placed'

This is a hypothetical situation where ill-placed files can cause GPFS to produce a 'Predicted Pool Utilization' of a negative value.

Suppose that 2 GB of data from a 5 GB file named **abc**, that is supposed to be in the **system** storage pool, are actually located in another pool. This 2 GB of data is said to be 'ill-placed'. Also, suppose that 3 GB of this file are in the **system** storage pool, and no other file is assigned to the **system** storage pool.

If you run the **mmapplypolicy** command to schedule file **abc** to be moved from the **system** storage pool to a storage pool named **YYY**, the **mmapplypolicy** command does the following:

1. Starts with the 'Current pool utilization' for the **system** storage pool, which is 3 GB.
2. Subtracts 5 GB, the size of file **abc**.
3. Arrives at a 'Predicted Pool Utilization' of negative 2 GB.

The **mmapplypolicy** command does not know how much of an 'ill-placed' file is currently in the wrong storage pool and how much is in the correct storage pool.

When there are ill-placed files in the **system** storage pool, the 'Predicted Pool Utilization' can be any positive or negative value. The positive value can be capped by the **LIMIT** clause of the **MIGRATE** rule. The 'Current Pool Utilizations' should always be between 0% and 100%.

Policies - Usage errors

These are common mistakes and misunderstandings encountered when dealing with policies:

1. You are advised to test your policy rules using the **mmapplypolicy** command with the **-I test** option.

   Also consider specifying a test-subdirectory within your file system. Do not apply a policy to an entire file system of vital files until you are confident that the rules correctly express your intentions. Even then, you are advised to do a sample run with the **mmapplypolicy -I test** command using the option **-L 3** or higher, to better understand which files are selected as candidates, and which candidates are chosen.

   The **-L** flag of the **mmapplypolicy** command can be used to check a policy before it is applied. For examples and more information on this flag, see "The mmapplypolicy -L command" on page 23.
2. There is a 1 MB limit on the total size of the policy file installed in GPFS.
3. GPFS is POSIX compatible and supports the POSIX file attributes: **mtime**, **ctime** and **atime**. However only **mtime** and **atime** may be referenced from the policy rules with the **MODIFICATION_TIME** and **ACCESS_TIME** attributes respectively. The **CREATION_TIME** attribute for a file is not supported by GPFS, and thus cannot be referenced from a policy.
4. Ensure that all clocks on all nodes of the GPFS cluster are synchronized. Depending on the policies in effect, variations in the clock times can cause unexpected behavior.

   The **mmapplypolicy** command uses the time on the node on which it is run as the current time. Policy rules may refer to a file's last access time or modification time, which is set by the node which last

accessed or modified the file. If the clocks are not synchronized, files may be treated as older or younger than their actual age, and this could cause files to be migrated or deleted prematurely, or not at all.

A suggested solution is to use NTP to keep the clocks synchronized on all nodes in the cluster.

5. The rules of a policy file are evaluated in order.

   A new file is assigned to the storage pool of the first rule that it matches. If the file fails to match any rule, the file creation fails with an **EINVAL** error code. A suggested solution is to put a **DEFAULT** clause as the last entry of the policy file.

6. When a policy file is installed, GPFS verifies that the named objects (storage pools) exist.

   However, GPFS allows an administrator to delete pools that are mentioned in the policy file. This allows more freedom for recovery from hardware errors. Consequently, the administrator must be careful when deleting objects referenced in the policy.

Filesets - Usage errors

These are common mistakes and misunderstandings encountered when dealing with filesets:

1. The maximum number of filesets that GPFS supports is 1000 filesets per file system.
2. Fileset junctions may not be deleted by the usual commands such as **rm -r** or **rmdir** commands, even though they look very much like ordinary directories.

   As a consequence these commands may fail when applied to a directory that is a fileset junction. Similarly, when **rm -r** is applied to a directory that contains a fileset junction, it will fail as well.

   On the other hand, **rm -r** will delete all the files contained in the filesets linked under the specified directory. Use the **mmunlinkfileset** command to remove fileset junctions.

3. Files and directories may not be moved from one fileset to another, nor may a hard link cross fileset boundaries.

   If the user is unaware of the locations of fileset junctions, **mv** and **ln** commands may fail unexpectedly. In most cases, the **mv** command will automatically compensate for this failure and use a combination of **cp** and **rm** to accomplish the desired result. Use the **mmlsfileset** command to view the locations of fileset junctions. Use the **mmlsattr -L** command to determine the fileset for any given file.

4. Because a snapshot saves the contents of a fileset, deleting a fileset included in a snapshot cannot completely remove the fileset.

   The fileset is put into a 'deleted' state and continues to appear in **mmlsfileset** output. Once the last snapshot containing the fileset is deleted, the fileset will be completely removed automatically. The **mmlsfileset** command indicates deleted fileset by putting their names in parentheses.

5. Deleting a large fileset may take some time and may be interrupted by other failures, such as disk errors or system crashes.

   When this occurs, the recovery action leaves the fileset in a 'being deleted' state. Such a fileset may not be linked into the namespace. The corrective action it to finish the deletion by reissuing the fileset delete command:

   ```
   mmdelfileset fs1 fsname1 -f
   ```

   The **mmlsfileset** command identifies filesets in this state by displaying a status of 'Deleting'.

6. If you unlink a fileset that has other filesets linked below it, any filesets linked to it (that is, child filesets) become inaccessible. The child filesets remain linked to the parent and will become accessible again when the parent is re-linked.
7. By default, the **mmdelfileset** command will not delete a fileset that is not empty.

   To empty a fileset, first unlink all its immediate child filesets, to remove their junctions from the fileset to be deleted. Then, while the fileset itself is still linked, use **rm -rf** or a similar command, to remove the rest of the contents of the fileset. Now the fileset may be unlinked and deleted. Alternatively, the fileset to be deleted can be unlinked first and then **mmdelfileset** can be used with the **-f** (force) option. This will unlink its child filesets, then destroy the files and directories contained in the fileset.

8. Deleting a fileset with the **-f** option requires traversing all the inodes in the file system. Although this operation is done in parallel, if the file system is large and the fileset is small, it may be faster to link the fileset and use **rm -rf** to empty the fileset before deleting it.

Storage pools - Usage errors

These are common mistakes and misunderstandings encountered when dealing with storage pools:

1. Only the **system** storage pool is allowed to store metadata. All other pools must have the **dataOnly** attribute.
2. Take care to create your storage pools with sufficient numbers of failure groups to enable GPFS replication.

   When the file system is created, GPFS requires all of the initial pools to have at least as many failure groups as defined by the default replication (**-m** and **-r** flags on the **mmcrfs** command). However, once the file system has been created, the user can create a storage pool with fewer failure groups than the default replication.

   The **mmadddisk** command issues a warning, but it allows the disks to be added and the storage pool defined. To use the new pool, the user must define a policy rule to create or migrate files into the new pool. This rule should be defined to set an appropriate replication level for each file assigned to the pool. If the replication level exceeds the number of failure groups in the storage pool, all files assigned to the pool incur added overhead on each write to the file, in order to mark the file as ill-replicated.

   To correct the problem, add additional disks to the storage pool, defining a different failure group, or insure that all policy rules that assign files to the pool also set the replication appropriately.
3. Do not use the **mmchdisk** or **mmrpldisk** command to change a disk's storage pool assignment.

   GPFS does not permit the **mmchdisk** or **mmrpldisk** command to change a disk's storage pool assignment. Changing the pool assignment requires all data residing on the disk to be moved to another disk before the disk can be reassigned. Moving the data is a costly and time-consuming operation; therefore GPFS requires an explicit **mmdeldisk** command to move it, rather than moving it as a side effect of another command.
4. Some storage pools allow larger disks to be added than do other storage pools.

   When the file system is created, GPFS defines the maximum size disk that can be supported using the on-disk data structures to represent it. Likewise, when defining a new storage pool, the newly created on-disk structures establish a limit on the maximum size disk that can be added to that pool.

   To add disks that exceed the maximum size allowed by a storage pool, simply create a new pool using the larger disks.
5. If you try to delete a storage pool when there are files still assigned to the pool, consider this:

   A storage pool is deleted when all disks assigned to the pool are deleted. To delete the last disk, all data residing in the pool must be moved to another pool. Likewise, any files assigned to the pool, whether or not they contain data, must be reassigned to another pool. The easiest method for reassigning all files and migrating all data is to use the **mmapplypolicy** command with a single rule to move all data from one pool to another. You should also install a new placement policy that does not assign new files to the old pool. Once all files have been migrated, reissue the **mmdeldisk** command to delete the disk and the storage pool.

   If all else fails, and you have a disk that has failed and cannot be recovered, Contact the IBM Support Center for commands to allow the disk to be deleted without migrating all data from it. Files with data left on the failed device will lose data. If the entire pool is deleted, any existing files assigned to that pool are reassigning to a 'broken' pool, which prevents writes to the file until the file is reassigned to a valid pool.
6. Ill-placed files - understanding and correcting them.

   The **mmapplypolicy** command migrates a file between pools by first assigning it to a new pool, then moving the file's data. Until the existing data is moved, the file is marked as 'ill-placed' to indicate that some of its data resides in its previous pool. In practice, **mmapplypolicy** assigns all files to be migrated to their new pools, then it migrates all of the data in parallel. Ill-placed files indicate that the **mmapplypolicy** command did not complete its last migration.

To correct the placement of the ill-placed files, the file data needs to be migrated to the assigned pools. You can use the **mmrestripefs**, or **mmrestripefile** commands to move the data.

7. Using the **-P** *PoolName* option on the **mmrestripefs**, command:

   This option restricts the restripe operation to a single storage pool. For example, after adding a disk to a pool, only the data in that pool needs to be restriped. In practice, **-P** *PoolName* simply restricts the operation to the files assigned to the specified pool. Files assigned to other pools are not included in the operation, even if the file is ill-placed and has data in the specified pool.

Errors encountered with policies

These are errors encountered with policies and how to analyze them:

1. Policy file never finishes, appears to be looping.

   The **mmapplypolicy** command runs by making two passes over the file system - one over the inodes and one over the directory structure. The policy rules are applied to each file to determine a list of candidate files. The list is sorted by the weighting specified in the rules, then applied to the file system. No file is ever moved more than once. However, due to the quantity of data involved, this operation may take a long time and appear to be hung or looping.

   The time required to run **mmapplypolicy** is a function of the number of files in the file system, the current load on the file system, and on the node in which **mmapplypolicy** is run. If this function appears to not finish, you may need to reduce the load on the file system or run **mmapplypolicy** on a less loaded node in the cluster.

2. Initial file placement is not correct.

   The placement rules specify a single pool for initial placement. The first rule that matches the file's attributes selects the initial pool. If that pool is incorrect, then the placement rules must be updated to select a different pool. You may see current placement rules by running **mmlspolicy -L**. For existing files, the file can be moved to its desired pool using the **mmrestripefile** or **mmchattr** commands.

   For examples and more information on **mmlspolicy -L**, see "The mmapplypolicy -L command" on page 23.

3. Data migration, deletion or exclusion not working properly.

   The **mmapplypolicy** command selects a list of candidate files to be migrated or deleted. The list is sorted by the weighting factor specified in the rules, then applied to a sufficient number of files on the candidate list to achieve the utilization thresholds specified by the pools. The actual migration and deletion are done in parallel.

   These are some reasons for apparently incorrect operation:

   - The file was not selected as a candidate for the expected rule. Each file is selected as a candidate for only the first rule that matched its attributes. If the matched rule specifies an invalid storage pool, the file is not moved. The **-L 4** option on **mmapplypolicy** displays the details for candidate selection and file exclusion.
   - The file was a candidate, but was not operated on. Only the candidates necessary to achieve the desired pool utilizations are migrated. Using the **-L 3** option displays more information on candidate selection and files chosen for migration.

     For more information on **mmlspolicy -L**, see "The mmapplypolicy -L command" on page 23.
   - The file was scheduled for migration but was not moved. In this case, the file will be shown as 'ill-placed' by the **mmlsattr -L** command, indicating that the migration did not succeed. This occurs if the new storage pool assigned to the file did not have sufficient free space for the file when the actual migration was attempted. Since migrations are done in parallel, it is possible that the target pool had files which were also migrating, but had not yet been moved. If the target pool now has sufficient free space, the files can be moved using the commands: **mmrestripefs**, **mmrestripefile**, **mmchattr**.

4. Asserts or error messages indicating a problem.

The policy rule language can only check for some errors at runtime. For example, a rule that causes a divide by zero cannot be checked when the policy file is installed. Errors of this type generate an error message and stop the policy evaluation for that file.

Note: I/O errors while migrating files indicate failing storage devices and must be addressed like any other I/O error. The same is true for any file system error or panic encountered while migrating files.

Errors encountered with filesets

These are error encountered with filesets and how to analyze them:

1. Problems can arise when running backup and archive utilities against a file system with unlinked filesets. See the section: *Filesets and backup* in *General Parallel File System: Advanced Administration Guide* for details.
2. In the rare case that the **mmfsck** command encounters a serious error checking the file system's fileset metadata, it may not be possible to reconstruct the fileset name and comment. These cannot be inferred from information elsewhere in the file system. If this happens, **mmfsck** will create a dummy name for the fileset, such as 'Fileset911' and the comment will be set to the empty string.
3. Sometimes **mmfsck** encounters orphaned files or directories (those without a parent directory), and traditionally these are reattached in a special directory called 'lost+found' in the file system root. When a file system contains multiple filesets, however, orphaned files and directories are reattached in the 'lost+found' directory in the root of the fileset to which they belong. For the root fileset, this directory appears in the usual place, but other filesets may each have their own 'lost+found' directory.

Errors encountered with storage pools

These are error encountered with policies and how to analyze them:

1. Access time to one pool appears slower than the others.

   A consequence of striping data across the disks is that the I/O throughput is limited by the slowest device. A device encountering hardware errors or recovering from hardware errors may effectively limit the throughput to all devices. However using storage pools, striping is done only across the disks assigned to the pool. Thus a slow disk impacts only its own pool; all other pools are not impeded.

   To correct the problem, check the connectivity and error logs for all disks in the slow pool.
2. Other storage pool problems might really be disk problems and should be pursued from the standpoint of making sure that your disks are properly configured and operational. See Chapter 8, "GPFS disk problems," on page 89.

Failures using the mmbackup command

Use the **mmbackup** command to back up the files in a GPFS file system to storage on a Tivoli® Storage Manager (TSM) server. A number of factors can cause **mmbackup** to fail.

The most common of these are:

- The file system is not mounted on the node issuing the **mmbackup** command.
- The file system is not mounted on the TSM client nodes.
- The TSM clients are not able to communicate with the TSM server due to authorization problems.
- The TSM server is down or out of storage space.
- When the target of the backup is tape, the TSM server may be unable to handle all of the backup client processes because the value of the TSM server's MAXNUMMP parameter is set lower than the number of client processes. This failure is indicated by message ANS1312E from TSM.
- The **mmbackup** command is unable to create a snapshot to perform the backup because a subdirectory with the name of the snapshot subdirectory already exists. This is usually caused by the

user doing a TSM restore of the backup without specifying a different name for receiving the restored contents of the file system than the name they were stored under in TSM, namely the snapshot subdirectory name.

The errors from **mmbackup** normally indicate the underlying problem.

GPFS error messages for mmbackup errors

6027-992
Subdirectory *name* already exists. Unable to create snapshot.

6027-1995
Device *deviceName* is not mounted on node *nodeName*.

6027-2101
Insufficient free space in *fileSystem* (*storage* minimum required).

TSM error messages

ANS1312E
Server media mount not possible.

Snapshot problems

Use the **mmlssnapshot** command as a general hint for snapshot-related problems, and the **mmsnapdir** command with the **-q** option to find out what snapshots exist, what state they are in, and the snapshot directory name used to permit access.

The **mmlssnapshot** command displays the list of **all** snapshots of a file system. This command lists the snapshot name, some attributes of the snapshot, as well as the snapshot's status. The **mmlssnapshot** command does not require the file system to be mounted.

Problems with locating a snapshot

Two commands are provided to assist in locating the snapshots in the file system directory structure. Only valid snapshots are visible in the file system directory structure. They appear in a hidden subdirectory of the file system's root directory. By default the subdirectory is named **.snapshots**. The valid snapshots appear as entries in the snapshot directory and may be traversed like any other directory. The **mmsnapdir** command can be used to display the assigned snapshot directory name.

Problems not directly related to snapshots

Many errors returned from the snapshot commands are not specifically related to the snapshot. For example, disk failures, IBM Virtual Shared Disk problems, or node failures could cause a snapshot command to fail. The response to these types of errors is always the same - fix the underlying problem and try the snapshot command again.

GPFS error messages for indirect snapshot errors

The error messages for this type of problem do not have message numbers, but can be recognized by their message text:

- 'Unable to sync all nodes, rc=*errorCode*.'
- 'Unable to get permission to create snapshot, rc=*errorCode*.'
- 'Unable to quiesce all nodes, rc=*errorCode*.'
- 'Unable to resume all nodes, rc=*errorCode*.'
- 'Unable to delete snapshot *filesystemName* from file system *snapshotName*, rc=*errorCode*.'
- 'Error restoring inode *number*, error *errorCode*.'
- 'Error deleting snapshot *snapshotName* in file system *filesystemName*, error *errorCode*.'

- '*commandString* failed, error *errorCode*.'
- 'None of the nodes in the cluster is reachable, or GPFS is down on all of the nodes.'
- 'File system *filesystemName* is not known to the GPFS cluster.'

Snapshot usage errors

Many errors returned from the snapshot commands are related to an incorrect snapshot name or other simple usage restrictions. Examples of incorrect snapshot name errors are trying to delete a snapshot that does not exist or trying to create a snapshot using the same name as one that already exists. An example of a snapshot restriction error is exceeding the maximum number of snapshots allowed at one time. For simple errors of these types, the user can determine the source of the error by reading the error message or by reading the description of the command. You may also run the **mmlssnapshot** command to see the complete list of existing snapshots.

GPFS error messages for snapshot usage errors

The error messages for this type of problem do not have message numbers, but can be recognized by their message text:

- 'File system *filesystemName* does not contain a snapshot *snapshotName*, rc=*errorCode*.'
- 'Cannot create a new snapshot until an existing one is deleted. File system *filesystemName* has a limit of *number* online snapshots.'
- 'Cannot restore snapshot. *snapshotName* is mounted on *number* nodes and in use on *number* nodes.'
- 'Cannot create a snapshot in a DM enabled file system, rc=*errorCode*.'

Snapshot status errors

Some snapshot commands like **mmdelsnapshot** and **mmrestorefs** may require a substantial amount of time to complete. If the command is interrupted, say by the user or due to a failure, the snapshot may be left in an invalid state. In many cases, the command must be completed before other snapshot commands are allowed to run. The source of the error may be determined from the error message, the command description, or the snapshot status available from **mmlssnapshot**.

GPFS error messages for snapshot status errors

The error messages for this type of problem do not have message numbers, but can be recognized by their message text:

- 'Cannot delete snapshot *snapshotName* which is *snapshotState*, error = *errorCode*.'
- 'Cannot restore snapshot *snapshotName* which is *snapshotState*, error = *errorCode*.'
- 'Previous snapshot *snapshotName* is invalid and must be deleted before a new snapshot may be created.'
- 'Previous snapshot *snapshotName* must be restored before a new snapshot may be created.'
- 'Previous snapshot *snapshotName* is invalid and must be deleted before another snapshot may be deleted.'
- 'Previous snapshot *snapshotName* is invalid and must be deleted before another snapshot may be restored.'
- 'More than one snapshot is marked for restore.'
- 'Offline snapshot being restored.'

Errors encountered when restoring a snapshot

If the **mmrestorefs** command is interrupted, the file system may not be consistent, and GPFS will not allow it to be mounted until the restore command completes. The error message for this case is:

'Mount of *filesystemName* failed: snapshot *snapshotName* must be restored before it can be mounted.'

If the **mmrestorefs** command fails due to error:

'Error restoring inode *number*, error *errorCode*.'

the user should fix the underlying problem and reissue the **mmrestorefs** command. If the user cannot fix the underlying problem, these steps can be taken to complete the restore command and recover the user data:

1. If there are other snapshots available, the user can restore a different snapshot.
2. If the error code in the message is **ENOSPC**, there are not enough free blocks in the file system to restore the selected snapshot. The user may add space to the file system by adding a new disk. As an alternative, the user may delete a different snapshot from the file system to free some existing space. The user is not allowed to delete the snapshot that is being restored. Once there is additional free space, reissue the **mmrestorefs** command.
3. The **mmrestorefs** command can be forced to continue, even if it encounters an error, by using the **-c** option. The command will restore as many files as possible, but may leave the file system in an inconsistent state. Some files may not have been restored or may no longer be accessible. The user should run **mmfsck** after the restore completes to make the file system consistent again.
4. If the above steps fail, the file system may be mounted in restricted mode, allowing the user to copy as many files as possible into a newly created file system, or one created from an offline backup of the data. See "Restricted mode mount" on page 21.

Note: In both steps 3 and 4, user data is lost. These steps are provided to allow as much user data as possible to be recovered.

Snapshot directory name conflicts

By default, all snapshots appear in a directory named **.snapshots** in the root directory of the file system. This directory is dynamically generated when the first snapshot is created and continues to exist even after the last snapshot is deleted. If the user tries to create the first snapshot, and a normal file or directory named **.snapshots** already exists, the **mmcrsnapshot** command will fail with a message:

'Snapshot directory *snapdirName* already exists.'

There are two ways to fix this problem:

1. Delete or rename the existing file or directory
2. Tell GPFS to use a different name for the dynamically-generated directory of snapshots by running the **mmsnapdir** command.

It is also possible to get a name conflict as a result of issuing the **mmrestorefs** command. Since **mmsnapdir** allows changing the name of the dynamically-generated snapshot directory, it is possible that an older snapshot contains a normal file or directory that conflicts with the current name of the snapshot directory. When this older snapshot is restored, the **mmrestorefs** command will recreate the old, normal file or directory in the file system root directory. The **mmrestorefs** command will not fail in this case, but the restored file or directory will hide the existing snapshots. After invoking **mmrestorefs** it may therefore appear as if the existing snapshots have disappeared. However, **mmlssnapshot** should still show all existing snapshots.

The fix is the similar to the one mentioned before. Perform one of these two steps:

1. After the **mmrestorefs** command completes, rename the conflicting file or directory that was restored in the root directory.
2. Run the **mmsnapdir** command to select a different name for the dynamically-generated snapshot directory.

Finally, the **mmsnapdir -a** option enables a dynamically-generated snapshot directory in every directory, not just the file system root. This allows each user quick access to snapshots of their own files by going into **.snapshots** in their home directory or any other of their directories.

Unlike **.snapshots** in the file system root, **.snapshots** in other directories is invisible, that is, an **ls -a** command will not list **.snapshots**. This is intentional because recursive file system utilities such as **find**, **du** or **ls -R** would otherwise either fail or produce incorrect or undesirable results. To access snapshots, the user must explicitly specify the name of the snapshot directory, for example: **ls ~/.snapshots**. If there is a name conflict (that is, a normal file or directory named **.snapshots** already exists in the user's home directory), the user must rename the existing file or directory.

Failures using the mmpmon command

The **mmpmon** command is thoroughly documented in the topic titled *Monitoring GPFS I/O performance with the **mmpmon** command* in the *General Parallel File System: Advanced Administration Guide*. Before proceeding with problem determination, review this material, as well as the **mmpmon** man page, to ensure that you are using **mmpmon** correctly.

Setup problems using mmpmon

Remember these points when using the **mmpmon** command:

- You must have root authority.
- The GPFS daemon must be active.
- The input file must contain valid input requests, one per line. When an incorrect request is detected by **mmpmon**, it issues an error message and terminates.

  Input requests that appear in the input file before the first incorrect request are processed by **mmpmon**.
- Do not alter the input file while **mmpmon** is running.
- Output from **mmpmon** is sent to standard output (STDOUT) and errors are sent to standard (STDERR).
- Up to five instances of **mmpmon** may run on a given node concurrently. See *Monitoring GPFS I/O performance with the **mmpmon** command* in the *General Parallel File System: Advanced Administration Guide* and search on *Running mmpmon concurrently from multiple users* for the limitations regarding concurrent usage of **mmpmon**.
- The **mmpmon** command does *not* support:
  - Monitoring read requests without monitoring writes, or the other way around.
  - Choosing which file systems to monitor.
  - Monitoring on a per-disk basis.
  - Specifying different size or latency ranges for reads and writes.
  - Specifying different latency values for a given size range.

Incorrect output from mmpmon

If the output from **mmpmon** is incorrect, such as zero counters when you know that I/O activity is taking place, consider these points:

1. Someone may have issued the `reset` or `rhist reset` requests.
2. Counters may have wrapped due to a large amount of I/O activity, or running **mmpmon** for an extended period of time. For a discussion of counter sizes and counter wrapping, see *Monitoring GPFS I/O performance with the **mmpmon** command* in the *General Parallel File System: Advanced Administration Guide* and search for *Counter sizes and counter wrapping*.
3. See *Monitoring GPFS I/O performance with the **mmpmon** command* in the *General Parallel File System: Advanced Administration Guide* and search for *General considerations about mmpmon output*, which gives specific instances where **mmpmon** output may be different than what was expected.

Abnormal termination or hang in mmpmon

If **mmpmon** hangs, perform these steps:

1. Ensure that sufficient time has elapsed to cover the **mmpmon** timeout value. It is controlled using the **-t** flag on the **mmpmon** command.
2. Issue the **ps** command to find the PID for **mmpmon**.
3. Issue the **kill** command to terminate this PID.
4. Try the function again.
5. If the problem persists, issue this command:

   ```
   mmfsadm dump eventsExporter
   ```
6. Copy the output of **mmfsadm** to a safe location.
7. Contact the IBM Support Center.

If **mmpmon** terminates abnormally, perform these steps:

1. Determine if the GPFS daemon has failed, and if so restart it.
2. Review your invocation of **mmpmon**, and verify the input.
3. Try the function again.
4. If the problem persists, contact the IBM Support Center.

Tracing the mmpmon command

When **mmpmon** does not work properly, there are two trace classes used to determine the cause of the problem. Use these only when requested by the IBM Support Center.

eventsExporter
: Reports attempts to connect and whether or not they were successful.

mmpmon
: Shows the command string that came in to **mmpmon**, and whether it was successful or not.

Note: Do not use the **perfmon** trace class of the GPFS trace to diagnose **mmpmon** problems. This trace event does not provide the necessary data.

NFS problems

There are some problems that can be encountered when GPFS interacts with NFS.

For details on how GPFS and NFS interact, see *General Parallel File System: Administration and Programming Reference* and search on *NFS and GPFS*.

These are some of the problems encountered when GPFS interacts with NFS:

- "NFS client with stale inode data"
- "NFS V4 problems" on page 87

NFS client with stale inode data

For performance reasons, some NFS implementations cache file information on the client. Some of the information (for example, file state information such as file size and timestamps) is not kept up-to-date in this cache. The client may view stale inode data (on **ls -l**, for example) if exporting a GPFS file system with NFS. If this is not acceptable for a given installation, caching can be turned off by mounting the file system on the client using the appropriate operating system **mount** command option (for example, **-o noac** on Linux NFS clients).

Turning off NFS caching will result in extra file systems operations to GPFS, and negatively affect its performance.

The clocks of all nodes in the GPFS cluster must be synchronized. If this is not done, NFS access to the data, as well as other GPFS file system operations, may be disrupted. NFS relies on metadata timestamps to validate the local operating system cache. If the same directory is either NFS-exported from more than one node, or is accessed with both the NFS and GPFS mount point, it is critical that clocks on all nodes that access the file system (GPFS nodes and NFS clients) are constantly synchronized using appropriate software (for example, NTP). Failure to do so may result in stale information seen on the NFS clients.

NFS V4 problems

Before analyzing an NFS V4 problem, review this documentation to determine if you are using NFS V4 ACLs and GPFS correctly.

1. The paper, *NFS Version 4 Protocol* and other information found at: www.nfsv4.org.
2. The unit titled, *Managing GPFS access control lists and NFS export* in *General Parallel File System: Administration and Programming Reference*.
3. The topic *GPFS exceptions and limitations to NFS V4 ACLs* in the *General Parallel File System: Administration and Programming Reference*.

The commands **mmdelacl** and **mmputacl** can be used to revert an NFS V4 ACL to a traditional ACL. Use the **mmdelacl** command to remove the ACL, leaving access controlled entirely by the permission bits in the mode. Then use the **chmod** command to modify the permissions, or the **mmputacl** and **mmeditacl** commands to assign a new ACL.

For files, the **mmputacl** and **mmeditacl** commands can be used at any time (without first issuing the **mmdelacl** command) to assign any type of ACL. The command **mmeditacl -k posix** provides a translation of the current ACL into traditional POSIX form and can be used to more easily create an ACL to edit, instead of having to create one from scratch.

Problems working with Samba

If Windows (Samba) clients fail to access files with messages indicating file sharing conflicts, and no such conflicts exist, there may be a mismatch with file locking rules.

File systems being exported with Samba may (depending on which version of Samba you are using) require the **-D nfs4** flag on the **mmchfs** or **mmcrfs** commands. This setting enables NFS V4 and CIFS (Samba) sharing rules. Some versions of Samba will fail share requests if the file system has not been configured to support them.

Data integrity

GPFS takes extraordinary care to maintain the integrity of customer data. However, certain hardware failures, or in extremely unusual circumstances, the occurrence of a programming error can cause the loss of data in a file system.

GPFS performs extensive checking to validate metadata and ceases using the file system if metadata becomes inconsistent. This can appear in two ways:

1. The file system will be unmounted and applications will begin seeing `ESTALE` return codes to file operations.
2. Error log entries indicating a `MMFS_SYSTEM_UNMOUNT` and a corruption error are generated.

If actual disk data corruption occurs, this error will appear on each node in succession. Contact the IBM Support Center before proceeding with the following steps.

1. Examine the error logs on the NSD servers for any indication of a disk error that has been reported.
2. Take appropriate disk problem determination and repair actions prior to continuing.

3. After completing any required disk repair actions, run the offline version of the **mmfsck** command on the file system.
4. If your error log or disk analysis tool indicates that specific disk blocks are in error, use the **mmfileid** command to determine which files are located on damaged areas of the disk, and then restore these files. See “The mmfileid command” on page 31 for more information.
5. If data corruption errors occur in only one node, it is probable that memory structures within the node have been corrupted. In this case, the file system is probably good but a program error exists in GPFS or another authorized program with access to GPFS data structures.

   Follow the directions in “Data integrity” on page 87 and then reboot the node. This should clear the problem. If the problem repeats on one node without affecting other nodes check the programming specifications code levels to determine that they are current and compatible and that no hardware errors were reported. Refer to *General Parallel File System: Concepts, Planning, and Installation Guide* for correct software levels.

Error numbers specific to GPFS application calls when data integrity may be corrupted

When there is the possibility of data corruption, GPFS may report these error numbers in the operating system error log, or return them to an application:

EVALIDATE=214, Invalid checksum or other consistency check failure on disk data structure.
: This indicates that internal checking has found an error in a metadata structure. The severity of the error depends on which data structure is involved. The cause of this is usually GPFS software, disk hardware or other software between GPFS and the disk. Running **mmfsck** should repair the error. The urgency of this depends on whether the error prevents access to some file or whether basic metadata structures are involved.

Chapter 8. GPFS disk problems

GPFS uses only disk devices prepared as Network Shared Disks (NSDs). However NSDs might exist on top of a number of underlying disk technologies.

NSDs, for example might be defined on top of Fibre Channel SAN connected disks, or defined on top of IBM Virtual Shared Disks. This information provides detail on the creation, use, and failure of NSDs and their underlying disk technologies

These are some of the errors encountered with GPFS disks and NSDs:

NSD and underlying disk subsystem failures

There are indications that will lead you to the conclusion that your file system has disk failures.

Some of those indications include:

- Your file system has been forced to unmount. See "File system forced unmount" on page 71
- The **mmlsmount** command indicates that the file system is not mounted on certain nodes.
- Your application is getting EIO errors.
- Operating system error logs indicate you have stopped using a disk in a replicated system, but your replication continues to operate.
- The **mmlsdisk** command shows that disks are down.

Error encountered while creating and using NSD disks

GPFS requires that disk devices be prepared as NSDs. This is done using the **mmcrnsd** command. The input to the **mmcrnsd** command is given in the form of disk descriptors. For a complete explanation of disk descriptors, refer to *General Parallel File System: Concepts, Planning, and Installation Guide* and search for *disk descriptors*.

For disks that are SAN-attached to all nodes in the cluster, *device_name* should refer to the disk device name in **/dev** on the node where the **mmcrnsd** command is issued. If *Server_List* is specified, *device_name* must refer to the name of the disk on the first server node. The same disk can have different local names on different nodes.

When you specify an NSD server node, that node performs all disk I/O operations on behalf of nodes in the cluster that do not have connectivity to the disk. You can also specify up to eight additional NSD server nodes. These additional NSD servers will become active if the first NSD server node fails or is unavailable.

When the **mmcrnsd** command encounters an error condition, this message is displayed:

6027-1636
Error found while checking disk descriptor: *descriptor*

Here, *descriptor* is the disk descriptor that encountered the error. Usually, this message is preceded by one or more messages describing the error more specifically.

Another error from **mmcrnsd** is:

6027-1661
Failed while processing disk descriptor *descriptor* on node *nodeName*.

This error can occur if an NSD server node does not have read and write access to the disk. The NSD server node needs to write an NSD volume ID to the raw disk. If an additional NSD server node is specified, that NSD server node will scan its disks to find this NSD volume ID string. If the disk is SAN-attached to all nodes in the cluster, the NSD volume ID is written to the disk by the node on which the **mmcrnsd** command is running.

Displaying NSD information

Use the **mmlsnsd** command to display information about the currently defined NSDs in the cluster. For example, if you issue **mmlsnsd**, your output may be similar to this:

```
File system Disk name NSD servers
------------------------------------------------------------------------

fs1 t65nsd4b (directly attached)
fs5 t65nsd12b c26f4gp01.ppd.pok.ibm.com,c26f4gp02.ppd.pok.ibm.com
fs6 t65nsd13b c26f4gp01.ppd.pok.ibm.com,c26f4gp02.ppd.pok.ibm.com,c26f4gp03.ppd.pok.ibm.com
```

This output shows that:

- There are three NSDs in this cluster: **t65nsd4b**, **t65nsd12b**, and **t65nsd13b**.
- NSD disk **t65nsd4b** of filesystem **fs1** is SAN-attached to all nodes in the cluster.
- NSD disk **t65nsd12b** of file system **fs5** has 2 NSD server nodes.
- NSD disk **t65nsd13b** of file system **fs6** has 3 NSD server nodes.

If you need to find out the local device names for these disks, you could use the **-m** option on the **mmlsnsd** command. For example, issuing:

```
mmlsnsd -m
```

produces output similar to this example:

```
 Disk name NSD volume ID Device Node name Remarks
---------------------------------------------------------------------------------------
t65nsd12b 0972364D45EF7B78 /dev/hdisk34 c26f4gp01.ppd.pok.ibm.com server node
t65nsd12b 0972364D45EF7B78 /dev/hdisk34 c26f4gp02.ppd.pok.ibm.com server node
t65nsd12b 0972364D45EF7B78 /dev/hdisk34 c26f4gp04.ppd.pok.ibm.com
t65nsd13b 0972364D00000001 /dev/hdisk35 c26f4gp01.ppd.pok.ibm.com server node
t65nsd13b 0972364D00000001 /dev/hdisk35 c26f4gp02.ppd.pok.ibm.com server node
t65nsd13b 0972364D00000001 - c26f4gp03.ppd.pok.ibm.com (not found) server node
t65nsd4b 0972364D45EF7614 /dev/hdisk26 c26f4gp04.ppd.pok.ibm.com
```

From this output we can tell that:

- The local disk name for **t65nsd12b** on NSD server **c26f4gp01** is **hdisk34**.
- NSD disk **t65nsd13b** is not attached to node on which the **mmlsnsd** command was issued, node**c26f4gp04**.
- The **mmlsnsd** command was not able to determine the local device for NSD disk **t65nsd13b** on **c26f4gp03** server.

To find the nodes to which disk **t65nsd4b** is attached and the corresponding local device for that disk, issue:

```
mmlsnsd -d t65nsd4b -M
```

Output is similar to this example:

```
Disk name    NSD volume ID      Device       Node name  Remarks
---------------------------------------------------------------------------------
t65nsd4b 0972364D45EF7614 /dev/hdisk92 c26f4gp01.ppd.pok.ibm.com
t65nsd4b 0972364D45EF7614 /dev/hdisk92 c26f4gp02.ppd.pok.ibm.com
t65nsd4b 0972364D45EF7614 - c26f4gp03.ppd.pok.ibm.com (not found) directly attached
t65nsd4b 0972364D45EF7614 /dev/hdisk26 c26f4gp04.ppd.pok.ibm.com
```

From this output we can tell that NSD **t65nsd4b** is:

- Known as **hdisk92** on node **c26f4gp01** and **c26f4gp02**.
- Known as **hdisk26** on node **c26f4gp04**
- Is not attached to node **c26f4gp03**

To display extended information about a node's view of its NSDs, the **mmlsnsd -X** command can be used:

```
mmlsnsd -X -d "gpfs10nsd;gpfs11nsd;sdbnsd;vsdn06"
```

Output is similar to this:

```
mmlsnsd -X -d "hd3n97;sdfnsd;hd5n98"
```

The system displays information similar to:

```
 Disk name    NSD volume ID      Device         Devtype  Node name                Remarks
-------------------------------------------------------------------------------------------------
 hd3n97       0972846145C8E927   /dev/hdisk3    hdisk    c5n97g.ppd.pok.ibm.com   server node,pr=no
 hd3n97       0972846145C8E927   /dev/hdisk3    hdisk    c5n98g.ppd.pok.ibm.com   server node,pr=no
 hd5n98       0972846245EB501C   /dev/hdisk5    hdisk    c5n97g.ppd.pok.ibm.com   server node,pr=no
 hd5n98       0972846245EB501C   /dev/hdisk5    hdisk    c5n98g.ppd.pok.ibm.com   server node,pr=no
 sdfnsd       0972845E45F02E81   /dev/sdf       generic  c5n94g.ppd.pok.ibm.com   server node
 sdfnsd       0972845E45F02E81   /dev/sdm       generic  c5n96g.ppd.pok.ibm.com   server node
```

From this output we can tell that:

- Disk **hd3n97** is an hdisk known as **/dev/hdisk3** on NSD server node **c5n97** and **c5n98**.
- Disk **sdfnsd** is a generic disk known as **/dev/sdf** and **/dev/sdm** on NSD server node **c5n94g** and **c5n96g**, respectively.
- In addition to the above information, the NSD volume ID is displayed for each disk.

Note: The **-m**, **-M** and **-X** options of the **mmlsnsd** command can be very time consuming, especially on large clusters. Use these options judiciously.

NSD creation fails with a message referring to an existing NSD

The following message:

6027-1662
Disk descriptor *descriptor* refers to an existing NSD *name*.

may appear when trying to create a Network Shared Disk on a disk that had previously been used as an NSD by GPFS. This situation arises if the **mmdelnsd** command failed to complete. These instructions describe how to delete an NSD that is not being used by GPFS anymore.

NSDs are deleted with the **mmdelnsd** command. Internally, this is a two-step process:

1. Remove the NSD definitions from the GPFS control information.
2. Zero-out GPFS-specific data structures on the disk.

If for some reason the second step fails, for example because the disk is damaged and cannot be written to, the **mmdelnsd** command issues a message describing the error and then another message stating the exact command to issue to complete the deletion of the NSD. If these instructions are not successfully completed, a subsequent **mmcrnsd** command can fail with

6027-1662
Disk descriptor *descriptor* refers to an existing NSD *name*.

This error message indicates that the disk is either an existing NSD, or that the disk was previously an NSD that had been removed from the GPFS cluster using the **mmdelnsd -p** command, and had not been marked as available.

If the GPFS data structures are not removed from the disk, it might be unusable for other purposes. For example, if you are trying to create an AIX volume group on the disk, the **mkvg** command might fail with messages similar to:

```
0516-1339 /usr/sbin/mkvg: Physical volume contains some 3rd party volume group.
0516-1397 /usr/sbin/mkvg: The physical volume hdisk5, will not be added to the volume group.
0516-862 /usr/sbin/mkvg: Unable to create volume group.
```

The easiest way to recover such a disk is to temporarily define it as an NSD again (using the **-v no** option) and then delete the just-created NSD. For example:

```
mmcrnsd -F filename -v no
mmdelnsd -F filename
```

GPFS has declared NSDs as down

There are several situations in which disks can appear to fail to GPFS. Almost all of these situations involve a failure of the underlying disk subsystem. The following information describes how GPFS reacts to these failures and how to find the cause.

GPFS will stop using a disk that is determined to have failed. This event is marked as `MMFS_DISKFAIL` in an error log entry (see “The operating system error log facility” on page 2). The state of a disk can be checked by issuing the **mmlsdisk** command.

The consequences of stopping disk usage depend on what is stored on the disk:

- Certain data blocks may be unavailable because the data residing on a stopped disk is not replicated.
- Certain data blocks may be unavailable because the controlling metadata resides on a stopped disk.
- In conjunction with other disks that have failed, all copies of critical data structures may be unavailable resulting in the unavailability of the entire file system.

The disk will remain unavailable until its status is explicitly changed through the **mmchdisk** command. After that command is issued, any replicas that exist on the failed disk are updated before the disk is used.

GPFS can declare disks **down** for a number of reasons:

- If the first NSD server goes down and additional NSD servers were not assigned, or all of the additional NSD servers are also down and no local device access is available on the node, the disks are marked as stopped.
- A failure of an underlying disk subsystem such as the IBM Virtual Shared Disk subsystem may result in a similar marking of disks as stopped. Locate the appropriate IBM Reliable Scalable Cluster Technology (RSCT) document at: publib.boulder.ibm.com/clresctr/windows/public/rsctbooks.html and search on *diagnosing IBM Virtual Shared Disk problems*. Follow the problem determination and repair actions specified. should this subsystem fail.
  1. Issue the **mmlsdisk** command to verify the status of the disks in the file system.
  2. Issue the **mmchdisk** command with the **-a** option to start all stopped disks.
- Disk failures should be accompanied by error log entries (see The operating system error log facility) for the failing disk. GPFS error log entries labelled **MMFS_DISKFAIL** will occur on the node detecting the error. This error log entry will contain the identifier of the failed disk. Follow the problem determination and repair actions specified in your disk vendor problem determination guide. After performing problem determination and repair issue the **mmchdisk** command to bring the disk back up.

Deleting IBM Virtual Shared Disks

IBM Virtual Shared Disks are never destroyed by GPFS. If you want to destroy any virtual shared disks, after removing them from GPFS using the **mmdelnsd** command, you must use the **removevsd** command. For detailed information on the **removevsd** command, locate the appropriate IBM Reliable Scalable Cluster Technology (RSCT) document at: publib.boulder.ibm.com/clresctr/windows/public/rsctbooks.html.

Unable to access disks

If you cannot open a disk, the specification of the disk may be incorrect. It is also possible that a configuration failure may have occurred during disk subsystem initialization. For example, on Linux you should consult **/var/log/messages** to determine if disk device configuration errors have occurred.

```
Feb 16 13:11:18 host123 kernel: SCSI device sdu: 35466240 512-byte hdwr sectors (18159 MB)
Feb 16 13:11:18 host123 kernel:  sdu: I/O error: dev 41:40, sector 0
Feb 16 13:11:18 host123 kernel:  unable to read partition table
```

On AIX, consult "The operating system error log facility" on page 2 for hardware configuration error log entries. For AIX IBM Virtual Shared disks, locate the appropriate IBM Reliable Scalable Cluster Technology (RSCT) document at: publib.boulder.ibm.com/clresctr/windows/public/rsctbooks.html and search on *diagnosing IBM Virtual Shared Disk problems*. Follow the problem determination and repair actions specified.

Accessible disk devices will generate error log entries similar to this example for a SSA device:

```
---------------------------------------------------------------------------
LABEL:          SSA_DEVICE_ERROR
IDENTIFIER:     FE9E9357

Date/Time:       Wed Sep  8 10:28:13 edt
Sequence Number: 54638
Machine Id:      000203334C00
Node Id:         c154n09
Class:           H
Type:            PERM
Resource Name:   pdisk23
Resource Class:  pdisk
Resource Type:   scsd
Location:        USSA4B33-D3
VPD:
        Manufacturer................IBM
        Machine Type and Model......DRVC18B
        Part Number.................09L1813
        ROS Level and ID............0022
        Serial Number...............6800D2A6HK
        EC Level....................E32032
        Device Specific.(Z2)........CUSHA022
        Device Specific.(Z3)........09L1813
        Device Specific.(Z4)........99168

Description
DISK OPERATION ERROR

Probable Causes
DASD DEVICE

Failure Causes
DISK DRIVE

        Recommended Actions
        PERFORM PROBLEM DETERMINATION PROCEDURES

Detail Data
ERROR CODE
2310 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000
---------------------------------------------------------------------------
```

or this one from GPFS:

```
---------------------------------------------------------------------------
LABEL:          MMFS_DISKFAIL
IDENTIFIER:     9C6C05FA

Date/Time:       Tue Aug  3 11:26:34 edt
Sequence Number: 55062
Machine Id:      000196364C00
Node Id:         c154n01
Class:           H
Type:            PERM
Resource Name:   mmfs
Resource Class:  NONE
Resource Type:   NONE
Location:

Description
DISK FAILURE

Probable Causes
STORAGE SUBSYSTEM
DISK

Failure Causes
STORAGE SUBSYSTEM
DISK

        Recommended Actions
        CHECK POWER
        RUN DIAGNOSTICS AGAINST THE FAILING DEVICE

Detail Data
EVENT CODE
     1027755
VOLUME
fs3
RETURN CODE
          19
PHYSICAL VOLUME
vp31n05
---------------------------------------------------------------------------
```

Guarding against disk failures

There are various ways to guard against the loss of data due to disk media failures. For example, the use of a RAID controller, which masks disk failures with parity disks, or a twin-tailed disk, could prevent the need for using these recovery steps.

GPFS offers a method of protection called *replication,* which overcomes disk failure at the expense of additional disk space. GPFS allows replication of data and metadata. This means that two instances of data, metadata, or both can be automatically created and maintained for any file in a GPFS file system. If one instance becomes unavailable due to disk failure, another instance is used instead. You can set different replication specifications for each file, or apply default settings specified at file system creation. Refer to *General Parallel File System: Concepts, Planning, and Installation Guide* and search for *file system recoverability parameters*.

Disk media failure

Regardless of whether you have chosen additional hardware or replication to protect your data against media failures, you first need to determine that the disk has completely failed. If the disk has completely failed and it is not the path to the disk which has failed, follow the procedures defined by your disk vendor. Otherwise:

1. Check on the states of the disks for the file system:

```
mmlsdisk fs1 -e
```

GPFS will mark disks **down** if there have been problems accessing the disk.

2. To prevent any I/O from going to the down disk, issue these commands *immediately*:

   ```
   mmchdisk fs1 suspend -d gpfs1nsd
   mmchdisk fs1 stop -d gpfs1nsd
   ```

   Note: If there are any GPFS file systems with pending I/O to the down disk, the I/O will timeout if the system administrator does not stop it.

   To see if there are any threads that have been waiting a long time for I/O to complete, on all nodes issue:

   ```
   mmfsadm dump waiters 10 | grep "I/O completion"
   ```

3. The next step is *irreversible*! Do not run this command unless data and metadata have been replicated. This step may take a while to scan for bad disk addresses.

   ```
   mmdeldisk fs1 gpfs1n12 -p
   ```

   Note: Use only if you have replicated metadata.

4. Invoke the **mmfileid** command with the operand **:BROKEN**:

   ```
   mmfileid :BROKEN
   ```

   For more information, see “The mmfileid command” on page 31.

5. After the disk is properly repaired and available for use, you can add it back to the file system.

Replicated metadata and data

If you have replicated metadata and data and only disks in a single failure group have failed, everything should still be running normally but with slightly degraded performance. You can determine the replication values set for the file system by issuing the **mmlsfs** command. Proceed with the appropriate course of action:

1. After the failed disk has been repaired, issue an **mmadddisk** command to add the disk to the file system:

   ```
   mmadddisk fs1 gpfs12nsd
   ```

   You can rebalance the file system at the same time by issuing:

   ```
   mmadddisk fs1 gpfs12nsd -r
   ```

   Note: Rebalancing of files is an I/O intensive and time consuming operation, and is important only for file systems with large files that are mostly invariant. In many cases, normal file update and creation will rebalance your file system over time, without the cost of the rebalancing.

2. To re-replicate data that only has single copy, issue:

   ```
   mmrestripefs fs1 -r
   ```

   Optionally, use the **-b** flag instead of the **-r** flag to rebalance across all disks.

   Note: Rebalancing of files is an I/O intensive and time consuming operation, and is important only for file systems with large files that are mostly invariant. In many cases, normal file update and creation will rebalance your file system over time, without the cost of the rebalancing.

3. Optionally, check the file system for metadata inconsistencies by issuing the offline version of **mmfsck**:

   ```
   mmfsck fs1
   ```

   If **mmfsck** succeeds, you may still have errors that occurred. Check to verify no files were lost. If files containing user data were lost, you will have to restore the files from the backup media.

   If **mmfsck** fails, sufficient metadata was lost and you need to recreate your file system and restore the data from backup media.

Replicated metadata only

If you have only replicated metadata, you should be able to recover some, but not all, of the user data. Recover any data to be kept using normal file operations or erase the file. If you read a file in blocksize chunks and get a failure return code and an **EIO** errno, that block of the file has been lost. The rest of the file may have useful data to recover, or it can be erased.

Strict replication

If you delete, replace, or suspend a disk with strict replication enforced, you may receive an **errno** of ENOSPC when you create or append data to an existing file. If you need to delete, replace, or suspend a disk, and you need to write new data while the disk is offline, you can disable strict replication by issuing the **mmchfs -K no** command before you perform the disk action. However, data written while replication is disabled will not be replicated properly. Therefore, after you perform the disk action, you must re-enable strict replication by issuing the **mmchfs** command with the original value of the **-K** option (**always** or **whenpossible**) and then run the **mmrestripe -r** command. To determine if a disk has strict replication enforced, issue the **mmlsfs -K** command.

No replication

When there is no replication, the system metadata has been lost and the file system is basically irrecoverable. You may be able to salvage some of the user data, but it will take work and time. A forced unmount of the file system will probably already have occurred. If not, it probably will very soon if you try to do any recovery work. You can manually force the unmount yourself:

1. Mount the file system in **read-only** mode (see "Read-only mode mount" on page 21). This will bypass recovery errors and let you read whatever you can find. Directories may be lost and give errors, and parts of files will be missing. Get what you can now, for all will soon be gone. On a single node, issue:

   ```
   mount -o ro /dev/fs1
   ```

2. If you read a file in blocksize chunks and get an **EIO** return code that block of the file has been lost. The rest of the file may have useful data to recover or it can be erased. To save the file system parameters for recreation of the file system, issue:

   ```
   mmlsfs fs1   > fs1.saveparms
   ```

   Note: This next step is *irreversible*!

   To delete the file system, issue:

   ```
   mmdelfs fs1
   ```

3. To repair the disks, see your disk vendor problem determination guide. Follow the problem determination and repair actions specified.
4. Delete the affected NSDs. Issue:

   ```
   mmdelnsd nsdname
   ```

   The system displays output similar to this:

   ```
   mmdelnsd: Processing disk nsdname
   mmdelnsd: 6027-1371 Propagating the changes to all affected nodes.
   This is an asynchronous process.
   ```

5. Create a disk descriptor file for the disks to be used. This will include recreating NSDs for the new file system.
6. Recreate the file system with either different parameters or the same as you used before. Use the disk descriptor file.
7. Restore lost data from backups.

GPFS error messages for disk media failures

Disk media failures can be associated with these GPFS message numbers:

6027-304
Remount failed for *device* after daemon restart.

6027-305
Perform **mmlsdisk** for any disk failures and remount.

6027-418
Inconsistent file system quorum. readQuorum=*value* writeQuorum=*value* quorumSize=*value*

6027-636
Disks marked as stopped or offline.

Error numbers specific to GPFS application calls when disk failure occurs

When a disk failure has occurred, GPFS may report these error numbers in the operating system error log, or return them to an application:

EOFFLINE = 208, Operation failed because a disk is offline
This error is most commonly returned when an attempt to open a disk fails. Since GPFS will attempt to continue operation with failed disks, this will be returned when the disk is first needed to complete a command or application request. If this return code occurs, check your disk for stopped states, and check to determine if the network path exists. In rare situations, this will be reported if virtual shared disk definitions are incorrect.

To repair the disks, see your disk vendor problem determination guide. Follow the problem determination and repair actions specified.

ENO_MGR = 212, The current file system manager failed and no new manager could be appointed.
This error usually occurs when a large number of disks are unavailable or when there has been a major network failure. Run the **mmlsdisk** command to determine whether disks have failed. If disks have failed, take corrective action by issuing the **mmchdisk** command. Check the operating system error log on all nodes for indications of RVSD errors.

To repair the disks, see your disk vendor problem determination guide. Follow the problem determination and repair actions specified.

Disk connectivity failure and recovery

If a disk is defined to have a local connection and to be connected to defined NSD servers, and the local connection fails, GPFS bypasses the broken local connection and uses the NSD servers to maintain disk access. The following error message appears in the **MMFS log**:

6027-361
Local access to *disk* failed with EIO, switching to access the disk remotely.

This is the default behavior, and can be changed with the **useNSDserver** file system mount option. See *General Parallel File System: Concepts, Planning, and Installation Guide* and search for *NSD server considerations*.

For a file system using the default mount option **useNSDserver=asneeded**, disk access fails over from local access to remote NSD access. Once local access is restored, GPFS detects this fact and switches back to local access. The detection and switch over are not instantaneous, but occur at approximately five minute intervals.

Partial disk failure

If the disk has only partially failed and you have chosen not to implement hardware protection against media failures, the steps to restore your data depends on whether you have used replication. If you have replicated neither your data nor metadata, you will need to issue the offline version of the **mmfsck** command, and then restore the lost information from the backup media. If it is just the data which was not replicated, you will need to restore the data from the backup media. There is no need to run the **mmfsck** command if the metadata is intact.

If both your data and metadata have been replicated, implement these recovery actions:

1. Unmount the file system:
   ```
   mmumount fs1 -a
   ```
2. Delete the disk from the file system:
   ```
   mmdeldisk fs1 gpfs10nsd -c
   ```
3. If you are replacing the disk, add the new disk to the file system:
   ```
   mmadddisk fs1 gpfs11nsd
   ```
4. Then restripe the file system:
   ```
   mmrestripefs fs1 -b
   ```

   Note: Ensure there is sufficient space elsewhere in your file system for the data to be stored by using the **mmdf** command.

GPFS has declared NSDs built on top of AIX logical volumes as down

Earlier releases of GPFS allowed AIX logical volumes to be used in GPFS file systems. Using AIX logical volumes in GPFS file systems is now discouraged as they are limited with regard to their clustering ability and cross platform support.

Existing file systems using AIX logical volumes are however still supported, and this information might be of use when working with those configurations.

Verify logical volumes are properly defined for GPFS use

To verify your logical volume configuration, you must first determine the mapping between the GPFS NSD and the underlying disk device. Issue the command:

```
mmlsnsd -m
```

which will display any underlying physical device present on this node which is backing the NSD. If the underlying device is a logical volume, perform a mapping from the logical volume to the volume group.

Issue the commands:

```
lsvg -o | lsvg -i -l
```

The output will be a list of logical volumes and corresponding volume groups. Now issue the **lsvg** command for the volume group containing the logical volume. For example:

```
lsvg gpfs1vg
```

The system displays information similar to:

```
VOLUME GROUP:   gpfs1vg                 VG IDENTIFIER:  000195600004c00000000ee60c66352
VG STATE:       active                  PP SIZE:        16 megabyte(s)
VG PERMISSION:  read/write              TOTAL PPs:      542 (8672 megabytes)
MAX LVs:        256                     FREE PPs:       0 (0 megabytes)
LVs:            1                       USED PPs:       542 (8672 megabytes)
OPEN LVs:       1                       QUORUM:         2
TOTAL PVs:      1                       VG DESCRIPTORS: 2
STALE PVs:      0                       STALE PPs:      0
ACTIVE PVs:     1                       AUTO ON:        no
MAX PPs per PV: 1016                    MAX PVs:        32
LTG size:       128 kilobyte(s)         AUTO SYNC:      no
HOT SPARE:      no
```

Check the volume group on each node

Make sure that all disks are properly defined to all nodes in the GPFS cluster:

1. Issue the AIX **lspv** command on all nodes in the GPFS cluster and save the output.
2. Compare the **pvid** and volume group fields for all GPFS volume groups.

Each volume group must have the same **pvid** and volume group name on each node. The **hdisk** name for these disks may vary.

For example, to verify the volume group **gpfs1vg** on the five nodes in the GPFS cluster, for each node in the cluster issue:

```
lspv | grep gpfs1vg
```

The system displays information similar to:

```
k145n01:  hdisk3         00001351566acb07    gpfs1vg    active
k145n02:  hdisk3         00001351566acb07    gpfs1vg    active
k145n03:  hdisk5         00001351566acb07    gpfs1vg    active
k145n04:  hdisk5         00001351566acb07    gpfs1vg    active
k145n05:  hdisk7         00001351566acb07    gpfs1vg    active
```

Here the output shows that on each of the five nodes the volume group **gpfs1vg** is the same physical disk (has the same **pvid**). The **hdisk** numbers vary, but the fact that they may be called different **hdisk** names on different nodes has been accounted for in the GPFS product. This is an example of a properly defined volume group.

If any of the **pvid**s were different for the same volume group, this would indicate that the same volume group name has been used when creating volume groups on different physical volumes. This will not work for GPFS. A volume group name can be used only for the same physical volume shared among nodes in a cluster. Refer to *AIX 5L™ System Management Guide: Operating System and Devices* at: publib.boulder.ibm.com/infocenter/pseries/index.jsp and search for *changing the attribute of a device*.

Volume group varyon problems

If an NSD backed by an underlying logical volume will not come online to a node, it may be due to **varyonvg** problems at the volume group layer. Issue the **varyoffvg** command for the volume group at all nodes and restart GPFS. On startup, GPFS will **varyon** any underlying volume groups in proper sequence.

Incorrect disk states after failures in mmcrvsd

The **mmcrvsd** command allows you to create IBM Virtual Shared Disks that can be used later as input to the **mmcrnsd** command. However, there are a number of factors that could cause **all** attempts to create a virtual shared disk to fail.

These factors include:

1. The state of the disk
2. The path to the disk
3. Incorrect or inconsistent input data to the command.

Note: Use this information only if your system uses IBM Virtual Shared Disks.

If the command fails, virtual shared disk that were successfully created remain in existence. Virtual shared disks are never destroyed by GPFS. If you want to destroy any virtual shared disks, you must issue the **removevsd** command.

The standard way of handling this error is to repair the specific disk error and reissue the **mmcrvsd -F** *DescFile* command, specifying the same disk descriptor file as the one specified in the first invocation of the command.

However, if you want to remove the successfully created virtual shared disks, issue:

```
removevsd -f -v vsdname
```

For detailed information on the **removevsd** command, locate the appropriate IBM Reliable Scalable Cluster Technology (RSCT) document at: publib.boulder.ibm.com/clresctr/windows/public/rsctbooks.html.

After removing a file system and deleting the NSDs with the **mmdelnsd** command, if you want to remove the virtual shared disks, follow the procedure in the appropriate IBM Reliable Scalable Cluster Technology (RSCT) document at: publib.boulder.ibm.com/clresctr/windows/public/rsctbooks.html.

Disk accessing commands fail to complete due to problems with some non-IBM disks

Certain disk commands, such as **mmcrfs**, **mmadddisk**, **mmrpldisk**, **mmmount** and the operating system's **mount**, might issue the **varyonvg -u** command if the NSD is backed by an AIX logical volume.

For some non-IBM disks, when many **varyonvg -u** commands are issued in parallel, some of the AIX **varyonvg -u** invocations do not complete, causing the disk command to hang.

This situation is recognized by the GPFS disk command not completing after a long period of time, and the persistence of the **varyonvg** processes as shown by the output of the **ps -ef** command on some of the nodes of the cluster. In these cases, **kill** the **varyonvg** processes that were issued by the GPFS disk command on the nodes of the cluster. This allows the GPFS disk command to complete. Before mounting the affected file system on any node where a **varyonvg** process was killed, issue the **varyonvg -u** command (**varyonvg -u** *vgname*) on the node to make the disk available to GPFS. Do this on each of the nodes in question, one by one, until all of the GPFS volume groups are varied online.

Long response time for virtual shared disk I/O operations

There are indications that lead you to the conclude that there are I/O problems. There are various courses of action that you can take if you are experiencing long response times.

Note: Use this procedure only if your system has IBM Virtual Shared Disks.

Some of those indications include:

- Applications are running abnormally slow.

  Note: Certain GPFS commands such as **mmrestripefs**, **mmdeldisk**, **mmrpldisk**, may take a long time to complete. This is normal.

If you are experiencing long response times, try these actions:

- Check for stopped or suspended disks by issuing the **lsvsd** command. See the appropriate IBM Reliable Scalable Cluster Technology (RSCT) document at: publib.boulder.ibm.com/clresctr/windows/public/rsctbooks.html for details. If there are stopped or suspended disks, locate the appropriate IBM Reliable Scalable Cluster Technology (RSCT) document at: publib.boulder.ibm.com/clresctr/windows/public/rsctbooks.html and search on *diagnosing IBM Virtual Shared Disk problems*. Follow the problem determination and repair actions specified.
- If you are experiencing nonresponsive virtual shared disks causing retries, run the **statvsd** command and notice if there are a number of retries on the application nodes. If the number is increasing, locate the appropriate IBM Reliable Scalable Cluster Technology (RSCT) document at: publib.boulder.ibm.com/clresctr/windows/public/rsctbooks.html and search on *diagnosing IBM Virtual Shared Disk problems*. Follow the problem determination and repair actions specified.
- Certain conditions can occur in the IBM Virtual Shared Disk which cause I/O operations to take an unusually long time (tens of minutes) to complete. These conditions often involve excessive loads on the server nodes. This happens when IBM Virtual Shared Disk retries operations and suspends specific disks that are waiting for failed components to recover or for memory resources on server nodes to become available.

In these situations, GPFS might present disk I/O operations to the IBM Virtual Shared Disk, and receive no response while retries and recovery proceed. If there is no error indication from GPFS, examine the state of the disks, the number of retries reported in the IBM Virtual Shared Disk statistics, and the number of retries reported at both the application node and the server in the communication subsystem. Locate the appropriate IBM Reliable Scalable Cluster Technology (RSCT) document at: publib.boulder.ibm.com/clresctr/windows/public/rsctbooks.html for additional information.

Persistent Reserve errors

You can use Persistent Reserve (PR) to provide faster failover times between disks that support this feature. PR allows the stripe group manager to ″fence″ disks during node failover by removing the reservation keys for that node. In contrast, non-PR disk failovers cause the system to wait until the disk lease expires.

GPFS allows file systems to have a mix of PR and non-PR disks. In this configuration, GPFS will fence PR disks for node failures and recovery and non-PR disk will use disk leasing. If all of the disks are PR disks, disk leasing is not used, so recovery times improve.

GPFS uses the **mmchconfig** command to enable PR. Issuing this command with the appropriate **usePersistentReserve** option configures disks automatically. If this command fails, the most likely cause is either a hardware or device driver problem. Other PR-related errors will probably be seen as file system unmounts that are related to disk reservation problems. This type of problem should be debugged with existing trace tools.

Understanding Persistent Reserve

Persistent Reserve (PR) refers to a set of Small Computer Systems Interface-3 (SCSI-3) standard commands and command options. These PR commands and command options give SCSI initiators the ability to establish, preempt, query, and reset a reservation policy with a specified target disk. The functions provided by PR commands are a superset of current reserve and release mechanisms. These functions are not compatible with legacy reserve and release mechanisms. Target disks can only support reservations from either the legacy mechanisms or the current mechanisms.

Note: Attempting to mix Persistent Reserve commands with legacy reserve and release commands will result in the target disk returning a reservation conflict error.

Persistent Reserve establishes an interface through a *reserve_policy* attribute for SCSI disks. You can optionally use this attribute to specify the type of reservation that the device driver will establish before accessing data on the disk. For devices that do not support the *reserve_policy* attribute, the drivers will use the value of the *reserve_lock* attribute to determine the type of reservation to use for the disk. GPFS supports four values for the *reserve_policy* attribute:

no_reserve::
Specifies that no reservations are used on the disk.

single_path::
Specifies that legacy reserve/release commands are used on the disk.

PR_exclusive::
Specifies that Persistent Reserve is used to establish exclusive host access to the disk.

PR_shared::
Specifies that Persistent Reserve is used to establish shared host access to the disk.

Persistent Reserve support affects both the parallel (scdisk) and SCSI-3 (scsidisk) disk device drivers and configuration methods. When a device is opened (for example, when the **varyonvg** command opens the underlying **hdisks**), the device driver checks the ODM for *reserve_policy* and *PR_key_value* and then

opens the device appropriately. For PR, each host attached to the shared disk must use unique registration key values for *reserve_policy* and *PR_key_value*. On AIX, you can display the values assigned to *reserve_policy* and *PR_key_value* by issuing:

```
lsattr -El hdiskx -a reserve_policy,PR_key_value
```

If needed, use the AIX **chdev** command to set *reserve_policy* and *PR_key_value*.

Note: GPFS manages *reserve_policy* and *PR_key_value* using `reserve_policy=PR_shared` when Persistent Reserve support is enabled and `reserve_policy=no_reserve` when Persistent Reserve is disabled.

Checking Persistent Reserve

For Persistent Reserve to function properly, you must have PR enabled on all of the disks that are PR-capable. To determine the PR status in the cluster:

1. Determine if PR is enabled on the cluster
   a. Issue **mmlsconfig**
   b. Check for `usePersistentReserve=yes`
2. Determine if PR is enabled for all disks on all nodes
   a. Make sure that GPFS has been started and mounted on all of the nodes
   b. Enable PR by issuing **mmchconfig**
   c. Issue the command **mmlsnsd -X** and look for `pr=yes` on all the hdisk lines

Notes:

1. To view the keys that are currently registered on a disk, issue the following command from a node that has access to the disk:

   ```
   /usr/lpp/mmfs/bin/tsprreadkeys hdiskx
   ```

2. To check the AIX ODM status of a single disk on a node, issue the following command from a node that has access to the disk:

   ```
   lsattr -El hdiskx -a reserve_policy,PR_key_value
   ```

Manually enabling or disabling Persistent Reserve

Attention: Manually enabling or disabling Persistent Reserve should only be done under the supervision of IBM Service with GPFS stopped on the node.

IBM Service will help you determine if the PR state is incorrect for a disk. If the PR state is incorrect, you may be directed to correct the situation by manually enabling or disabling PR on that disk.

Chapter 9. Other problem determination hints and tips

These hints and tips might be helpful when investigating problems encountered while using GPFS and the IBM eServer™ High Performance Switch, logical volumes, quorum nodes, or system performance.

See these topics for more information:

Disk processing when using the IBM eServer High Performance Switch

When using an IBM eServer High Performance Switch (HPS) in your configuration, it is suggested you process your disks in two steps.

1. Create virtual shared disks on each physical disk through the **mmcrvsd** command.
2. Using the rewritten disk descriptors from the **mmcrvsd** command, create NSDs through the **mmcrnsd** command.

Which virtual shared disk is associated with a volume group?

If an error report contains a reference to a volume group of an IBM Virtual Shared Disk pertaining to GPFS, you can use the **lsvg** and **vsdatalst** commands to determine which virtual shared disk corresponds to the volume group.

Note: Use this procedure only if the failure occurred on an AIX node and you are using IBM Virtual Shared Disks.

For example, if you wanted to know the virtual shared disk associated with volume group **gpfs2vg**, use the following steps:

1. Issue the **lsvg** command:

   ```
   lsvg -l gpfs2vg
   ```

   which returns output similar to:

   ```
   gpfs2vg:
   LV NAME             TYPE       LPs   PPs   PVs  LV STATE      MOUNT POINT
   gpfs2lv             jfs        537   537   1    open/syncd    N/A
   ```

   The name of the logical volume, `gpfs2lv`, appears in the third row first column.
2. Issue the **vsdatalst** command to cross reference the logical volume and IBM Virtual Shared Disk.

   ```
   vsdatalst -v | grep gpfs2lv
   ```

   The system displays output similar to this. Header lines have been added here for clarity, but do not appear in the output.

```
VSD Table
VSD name          logical volume  Global Volume Group     minor# size_in_MB
----------------- --------------- ----------------------- ------ ----------
gpfs2vsd          gpfs2lv         gpfs2gvg                     1 8672
```

The virtual shared disk associated with the logical volume is in the first column.

Which physical disk is associated with a virtual shared disk?

If there is a virtual shared disk timeout or a GPFS disk failure, the steps in this procedure will help you determine which physical disk is failing.

Note: Use this procedure only if the failure occurred on an AIX node and you are using IBM Virtual Shared Disks.

To determine which physical disk is failing, follow these steps:

1. Determine the state of the virtual shared disk by issuing:

   ```
   lsvsd -l gpfs87vsd
   ```

   The system displays output similar to:

   ```
   minor  state  server  lv_major  lv_minor  vsd_name    size (MB)  server_list
   83     STP      -1        0         0      gpfs87vsd      20         1
   ```

2. Determine the global volume group name by issuing:

   ```
   vsdatalst -v | grep gpfs87vsd
   ```

   The system displays output similar to this. Header lines have been added here for clarity, but do not appear in the output.

   ```
   VSD Table
   VSD name          logical volume  Global Volume Group     minor# size_in_MB
   ----------------- --------------- ----------------------- ------ ---------
   gpfs87vsd         gpfs87lv        gpfs87gvg               88     8672
   ```

   The global volume group name is in the third column.

3. Determine the node number for the virtual shared disk server by issuing:

   ```
   vsdatalst -g | grep gpfs87gvg
   ```

   The system displays output similar to this. Header lines have been added here for clarity, but do not appear in the output.

   ```
   VSD Global Volume Group Information
   Global Volume     Local     Server Node  Numbers:  eio_
   Group name        VG name   primary      backup    recovery  Recovery server_list
   ----------------  --------  -----------  --------  --------  -------- -----------
   gpfs87gvg         gpfs87vg     4          5          1         00         0
   ```

   The node number of the primary disk server is in the third column.

4. Determine the name of this node by issuing this command:

   ```
   vsdatalst -n
   ```

   The system produces output similar to this:

   ```
          VSD Node Information
                                              Initial Maximum     VSD      rw      Buddy Buffer
     node                 VSD      IP packet   cache   cache request request minimum maximum size: #
   number host_name       adapter    size    buffers buffers   count   count    size    size maxbufs
   ------ --------------- -------- --------- ------- ------- ------- ------- ------- ------- -------
        1 k164n01.kgn     css0         61440      64      64     256      64    4096  262144       6
        2 k164n02.kgn     css0         61440      64     256     256      48    4096  262144     128
        3 k164n03.kgn     css0         61440      12      64     256     256    4096  262144       4
   ```

```
     4 k164n04.kgn     css0         61440      64     256     256      48    4096  262144     128
     5 k164n05.kgn     css0         61440      64     256     256      48    4096  262144       1
     6 k164n06.kgn     css0         61440      64     256     256      48    4096  262144     128
```

5. On the primary virtual shared disk server (**k164n04**), determine the physical disk name by issuing:

   ```
   lspv | grep gpfs87vg
   ```

   The system displays output similar to this. Header lines have been added here for clarity, but do not appear in the **lspv** output.

   ```
   disk name      PVID            Global Volume Group
   -------------- --------------- ------------------
   hdisk3         00001351566acb07    gpfs87vg
   ```

   The physical disk name is in the first column.

6. Follow the problem determination and repair action in "NSD and underlying disk subsystem failures" on page 89.

Which physical disk is associated with a logical volume?

Earlier releases of GPFS allowed AIX logical volumes to be used in GPFS file systems. Their use is now discouraged because they are limited with regard to their clustering ability and cross platform support.

Existing file systems using AIX logical volumes are, however, still supported. This information might be of use when working with those configurations.

If an error report contains a reference to a logical volume pertaining to GPFS, you can use the **lslv -l** command to list the physical volume name. For example, if you want to find the physical disk associated with logical volume **gpfs7lv**, issue:

```
lslv -l gpfs44lv
```

Output is similar to this, with the physical volume name in column one.

```
  gpfs44lv:N/A
  PV                COPIES        IN BAND       DISTRIBUTION
  hdisk8            537:000:000   100%          108:107:107:107:108
```

Which nodes in my cluster are quorum nodes?

Use the **mmlscluster** command to determine which nodes in your cluster are quorum nodes.

Output is similar to this:

```
GPFS cluster information
========================
  GPFS cluster name:         cluster1.kgn.ibm.com
  GPFS cluster id:           680681562214606028
  GPFS UID domain:           cluster1.kgn.ibm.com
  Remote shell command:      /usr/bin/rsh
  Remote file copy command:  /usr/bin/rcp

GPFS cluster configuration servers:
-----------------------------------
  Primary server:    k164n06.kgn.ibm.com
  Secondary server:  k164n05.kgn.ibm.com

 Node  Daemon node name         IP address       Admin node name       Designation
--------------------------------------------------------------------------------
   1   k164n04.kgn.ibm.com     198.117.68.68   k164n04.kgn.ibm.com   quorum
   2   k164n05.kgn.ibm.com     198.117.68.71   k164n05.kgn.ibm.com   quorum
   3   k164n06.kgn.ibm.com     198.117.68.70   k164n06.kgn.ibm.com
```

In this example, **k164n04** and **k164n05** are quorum nodes and **k164n06** is a nonquorum node.

To change the quorum status of a node, use the **mmchnode** command. To change one quorum node to nonquorum, GPFS does not have to be stopped. If you are changing more than one node at the same time, GPFS needs to be down on all the affected nodes. GPFS does not have to be stopped when changing nonquorum nodes to quorum nodes, nor does it need to be stopped on nodes that are not affected.

For example, to make **k164n05** a nonquorum node, and **k164n06** a quorum node, issue these commands:

```
mmchnode --nonquorum -N k164n05
mmchnode --quorum -N k164n06
```

To set a node's quorum designation at the time that it is added to the cluster, see the **mmcrcluster** or **mmaddnode** commands.

Are there any special concerns regarding the use of mksysb images?

Using a **mksysb** image to install or reinstall a node that is running the IBM Virtual Shared Disk licensed product, presents some challenges to be aware of.

The RVSD subsystem expects that the underlying disks and volume groups are available to AIX. Therefore, if you are reinstalling an existing server node from a **mksysb** image, you may need to use the AIX **importvg** command to make AIX aware of the preexisting volume groups.

Note: All virtual shared disk volume groups *must* have the **auto varyon** attribute turned off. To accomplished this, after the **importvg**, issue:

```
chvg -a -n volume_group_name
```

If the **mksysb** image is not an exact copy of the node being reinstalled, you will need to remove the node from the RSCT peer domain and then add it back using the original RSCT peer domain node number. See the appropriate IBM Reliable Scalable Cluster Technology (RSCT) document at: publib.boulder.ibm.com/clresctr/windows/public/rsctbooks.html for more detail, but basically you can do this by issuing these commands:

```
rmrpnode  Node_name
preprpnode Node_name
addrpnode -f fileName
```

where *fileName* contains ″Node_name Node_number″.

What is stored in the /tmp/mmfs directory and why does it sometimes disappear?

When GPFS encounters an internal problem, certain state information is saved in the GPFS dump directory for later analysis by IBM service.

The default dump directory for GPFS is **/tmp/mmfs**. This directory might disappear on Linux if cron is set to run the **/etc/cron.daily/tmpwatch** script. The **tmpwatch** script removes files and directory in **/tmp** that have not been accessed recently. Administrators who want to use a different directory for GPFS dumps can change the directory by issuing this command:

```
mmchconfig dataStructureDump=/name_of_some_other_big_file_system
```

Why does my system load increase significantly during the night?

On some Linux distributions, cron runs the **/etc/cron.daily/slocate.cron** job every night. This will try to index all the files in GPFS. This will put a very large load on the GPFS token manager.

You can exclude all GPFS file systems by adding **gpfs** to the **excludeFileSytemType** list in this script, or exclude specific GPFS file systems in the **excludeFileSytemType** list.

```
/usr/bin/updatedb -f "excludeFileSystemType" -e "excludeFileSystem"
```

If indexing GPFS file systems is desired, only one node should run the **updatedb** command and build the database in a GPFS file system. If the database is built within a GPFS file system it will be visible on all nodes after one node finishes building it.

What do I do if I receive message 6027-648?

The **mmedquota** or **mmdefedquota** commands can fail with message **6027-648: EDITOR environment variable must be full path name**. This message occurs when the value of the EDITOR environment variable is not an absolute path name.

To resolve this error, do the following:

1. Change the value of the EDITOR environment variable to an absolute path name.
2. Check to see if the EDITOR variable is set in the **$HOME/.kshrc** file. If it is set, check to see if it is an absolute path name because the **mmedquota** or **mmdefedquota** command could retrieve the EDITOR environment variable from that file.

Why can't I see my newly mounted Windows file system?

On Windows, a newly mounted file system might not be visible to you if you are currently logged on to a system. This can happen if you have mapped a network share to the same drive letter as GPFS.

Once you start a new session (by logging out and logging back in), the use of the GPFS drive letter will supersede any of your settings for the same drive letter. This is standard behavior for all local file systems on Windows.

Chapter 10. Messages

6027-000 **Attention: A disk being removed reduces the number of failure groups to *nFailureGroups*, which is below the number required for replication: *nReplicas*.**

Explanation:
Replication cannot protect data against disk failures when there are insufficient failure groups.

User response:
Add more disks in new failure groups to the file system or accept the risk of data loss.

6027-300 **mmfsd ready**

Explanation:
The **mmfsd** server is up and running.

User response:
None. Informational message only.

6027-301 **Initialization file *fileName* could not be run.**

Explanation:
Failed to **execvp** the initialization shell script.

User response:
Check file existence and access permissions.

6027-302 **Could not execute *gpfsreadyScript*.**

Explanation:
The **verifyGpfsReady=yes** configuration attribute is set, but the **/var/mmfs/etc/gpfsready** script could not be executed.

User response:
Make sure that the **/var/mmfs/etc/gpfsready** script exists and is executable, or disable the **verifyGpfsReady** option by using **mmchconfig verifyGpfsReady=no**.

6027-303 ***gpfsreadyScript* killed by signal.**

Explanation:
The **verifyGpfsReady=yes** configuration attribute is set, but the **/var/mmfs/etc/gpfsready** script did not complete successfully.

User response:
Make sure that the **/var/mmfs/etc/gpfsready** script completes and returns a zero exit status, or disable the **verifyGpfsReady** option by using **mmchconfig verifyGpfsReady=no**.

6027-304 ***gpfsreadyScript* ended abnormally.**

Explanation:
The **verifyGpfsReady=yes** configuration attribute is set, but the **/var/mmfs/etc/gpfsready** script did not complete successfully.

User response:
Make sure that the **/var/mmfs/etc/gpfsready** script completes and returns a zero exit status, or disable the **verifyGpfsReady** option by using **mmchconfig verifyGpfsReady=no**.

6027-305 ***gpfsreadyScript* failed with exit code *exitCode*.**

Explanation:
The **verifyGpfsReady=yes** configuration attribute is set, but the **/var/mmfs/etc/gpfsready** script did not complete successfully.

User response:
Make sure that the **/var/mmfs/etc/gpfsready** script completes and returns a zero exit status, or disable the **verifyGpfsReady** option by using **mmchconfig verifyGpfsReady=no**.

6027-306 **Could not initialize internode communication.**

Explanation:
The GPFS daemon was unable to initialize the communications required to proceed.

User response:
User action depends on the return code shown in the accompanying message (**/usr/include/errno.h**). The communications failure that caused the failure must be corrected. One possibility is an **rc** value of 67, indicating that the required port is unavailable. This may mean that a previous version of the **mmfs** daemon is still running. Killing that daemon may resolve the problem.

6027-310 ***command* initializing. {Version *versionName*: Built *date time*}**

Explanation:
The **mmfsd** server has started execution.

User response:
None. Informational message only.

6027-311 ***programName* is shutting down.**

Explanation:
The stated program is about to terminate.

User response:
None. Informational message only.

6027-312 **Unknown trace class *traceClass*.**

Explanation:

The trace class is not recognized.

User response:
Specify a valid trace class.

6027-313 Cannot open configuration file *fileName*.

Explanation:
The configuration file could not be opened.

User response:
The configuration file is **/var/mmfs/gen/mmfs.cfg**. Verify that this file and **/var/mmfs/gen/mmsdrfs** exist in your system.

6027-314 *command* requires *SuperuserName* authority to execute.

Explanation:
The mmfsd server was started by a user without superuser authority.

User response:
Log on as a superuser and reissue the command.

6027-315 Bad config file entry in *fileName*, line *number*.

Explanation:
The configuration file has an incorrect entry.

User response:
Fix the syntax error in the configuration file. Verify that you are not using a configuration file that was created on a release of GPFS subsequent to the one that you are currently running.

6027-316 Unknown config parameter *parameter* in *fileName*, line *number*.

Explanation:
There is an unknown parameter in the configuration file.

User response:
Fix the syntax error in the configuration file. Verify that you are not using a configuration file that was created on a release of GPFS subsequent to the one you are currently running.

6027-317 Old server with PID *pid* still running.

Explanation:
An old copy of **mmfsd** is still running.

User response:
This message would occur only if the user bypasses the SRC. The normal message in this case would be an SRC message stating that multiple instances are not allowed. If it occurs, stop the previous instance and use the SRC commands to restart the daemon.

6027-318 Watchdog: Some process appears stuck; stopped the daemon process.

Explanation:
A high priority process got into a loop.

User response:
Stop the old instance of the **mmfs** server, then restart it.

6027-319 Could not create shared segment.

Explanation:
The shared segment could not be created.

User response:
This is an error from the AIX operating system. Check the accompanying error indications from AIX.

6027-320 Could not map shared segment.

Explanation:
The shared segment could not be attached.

User response:
This is an error from the AIX operating system. Check the accompanying error indications from AIX.

6027-321 Shared segment mapped at wrong address (is *value*, should be *value*).

Explanation:
The shared segment did not get mapped to the expected address.

User response:
Contact the IBM Support Center.

6027-322 Could not map shared segment in kernel extension.

Explanation:
The shared segment could not be mapped in the kernel.

User response:
If an **EINVAL** error message is displayed, the kernel extension could not use the shared segment because it did not have the correct GPFS version number. Unload the kernel extension and restart the GPFS daemon.

6027-323 Error unmapping shared segment

Explanation:
The shared segment could not be detached.

User response:
Check reason given by error message.

6027-324 Could not create message queue for main process

Explanation:
The message queue for the main process could not be created. This is probably an operating system error.

User response:
Contact the IBM Support Center.

6027-328 **Value** *value* **for** *parameter* **is out of range in** *fileName***. Valid values are** *value* **through** *value***.** *value* **used.**

Explanation:
An error was found in the **/var/mmfs/etc/mmfs.cfg** files.

User response:
Check the **mmfs.cfg** file.

6027-329 **Cannot pin the main shared segment:** *name*

Explanation:
Trying to pin the shared segment during initialization.

User response:
Check the **mmfs.cfg** file. The **pagepool** size may be too large. It cannot be more than 80% of real memory. If a previous **mmfsd** crashed, check for processes that begin with the name **mmfs** that may be holding on to an old pinned shared segment. Issue **mmchconfig** command to change the **pagepool** size.

6027-334 **Error initializing internal communications.**

Explanation:
The mailbox system used by the daemon for communication with the kernel cannot be initialized.

User response:
Increase the size of available memory using the **mmchconfig** command.

6027-335 **Configuration error: check** *fileName***.**

Explanation:
A configuration error is found.

User response:
Check the **mmfs.cfg** file and other error messages.

6027-336 **Value** *value* **for configuration parameter** *parameter* **is not valid. Check** *fileName***.**

Explanation:
A configuration error is found.

User response:
Check the **mmfs.cfg** file.

6027-337 **Waiting for resources to be reclaimed before exiting.**

Explanation:
The **mmfsd** daemon is attempting to terminate, but cannot because data structures in the daemon shared segment may still be referenced by kernel code. This message may be accompanied by other messages that show which disks still have I/O in progress.

User response:
None. Informational message only.

6027-338 **Waiting for** *number* **user(s) of shared segment to release it.**

Explanation:
The **mmfsd** daemon is attempting to terminate, but cannot because some process is holding the shared segment while in a system call. The message will repeat every 30 seconds until the count drops to zero.

User response:
Find the process that is not responding, and find a way to get it out of its system call.

6027-339 **Nonnumeric trace value** *value* **after class** *classname***.**

Explanation:
The specified trace value is not recognized.

User response:
Specify a valid trace integer value.

6027-340 **Could not retrieve the network interface configuration.**

Explanation:
The socket or **ioctl** call to retrieve the interface configuration fails.

User response:
Check reason for given error.

6027-341 **Node** *nodeName* **is incompatible because its maximum compatible version (***number***) is less than the version of this node (***number***). [***value***/***value***]**

Explanation:
The GPFS daemon tried to make a connection with another GPFS daemon. However, the other daemon is not compatible. Its maximum compatible version is less than the version of the daemon running on this node. The numbers in square brackets are for IBM Service use.

User response:
Verify your GPFS daemon version.

6027-342 **Node** *nodeName* **is incompatible because its minimum compatible version (***number***) is greater than the version of this node (***number***). [***value***/***value***]**

Explanation:

The GPFS daemon tried to make a connection with another GPFS daemon. However, the other daemon is not compatible. Its minimum compatible version is greater than the version of the daemon running on this node. The numbers in square brackets are for IBM Service use.

User response:
Verify your GPFS daemon version.

6027-343 **Node** *nodeName* **is incompatible because its version (***number***) is less than the minimum compatible version of this node (***number***). [***value***/***value***]**

Explanation:
The GPFS daemon tried to make a connection with another GPFS daemon. However, the other daemon is not compatible. Its version is less than the minimum compatible version of the daemon running on this node. The numbers in square brackets are for IBM Service use.

User response:
Verify your GPFS daemon version.

6027-344 **Node** *nodeName* **is incompatible because its version (***number***) is greater than the maximum compatible version of this node (***number***). [***value***/***value***]**

Explanation:
The GPFS daemon tried to make a connection with another GPFS daemon. However, the other daemon is not compatible. Its version is greater than the maximum compatible version of the daemon running on this node. The numbers in square brackets are for IBM Service use.

User response:
Verify your GPFS daemon version.

6027-345 **Network error on** *ipAddress***, check connectivity.**

Explanation:
A TCP error has caused GPFS to exit due to a bad return code from an error. Exiting allows recovery to proceed on another node and resources are not tied up on this node.

User response:
Follow network problem determination procedures.

6027-346 **Incompatible daemon version. My version =***number***, repl.my_version =** *number***.**

Explanation:
The GPFS daemon tried to make a connection with another GPFS daemon. However, the other GPFS daemon is not the same version and it sent a reply indicating its version number is incompatible.

User response:
Verify your GPFS daemon version.

6027-347 **Remote host** *ipAddress* **refused connection because IP address** *ipAddress* **was not in the node list file.**

Explanation:
The GPFS daemon tried to make a connection with another GPFS daemon. However, the other GPFS daemon sent a reply indicating it did not recognize the IP address of the connector.

User response:
Add the IP address of the local host to the node list file on the remote host.

6027-348 **Bad ″subnets″ configuration: invalid subnet** *ipAddress***.**

Explanation:
A subnet specified by the **subnets** configuration parameter could not be parsed.

User response:
Run the **mmlsconfig** command and check the value of the **subnets** parameter. Each subnet must be specified as a dotted-decimal IP address. Run the **mmchconfig subnets** command to correct the value.

6027-349 **Bad subnets configuration: invalid cluster name pattern** *clusterNamePattern***.**

Explanation:
A cluster name pattern specified by the **subnets** configuration parameter could not be parsed.

User response:
Run the **mmlsconfig** command and check the value of the **subnets** parameter. The optional cluster name pattern following subnet address must be a shell-style pattern allowing '*', '/' and '[...]' as wild cards. Run the **mmchconfig subnets** command to correct the value.

6027-350 **Bad subnets configuration: primary IP address** *ipAddress* **is on a private subnet. Use a public IP address instead.**

Explanation:
GPFS is configured to allow multiple IP addresses per node (**subnets** configuration parameter), but the primary IP address of the node (the one specified when the cluster was created or when the node was added to the cluster) was found to be on a private subnet. If multiple IP addresses are used, the primary address must be a public IP address.

User response:
Remove the node from the cluster; then add it back using a public IP address.

6027-358 Communication with mmspsecserver through socket *name* failed, err *value*: *errorString*, msgType *messageType*.

Explanation:
Communication failed between **spsecClient** (the daemon) and **spsecServer**.

User response:
Verify both the communication socket and the **mmspsecserver** process.

6027-359 The mmspsecserver process is shutting down. Reason: *explanation*.

Explanation:
The **mmspsecserver** process received a signal from the **mmfsd** daemon or encountered an error on execution.

User response:
Verify the reason for shutdown.

6027-360 Disk *name* must be removed from the /etc/filesystems stanza before it can be deleted. Another disk in the file system can be added in its place if needed.

Explanation:
A disk being deleted is found listed in the *disks=* list for a file system.

User response:
Remove the disk from list.

6027-361 Local access to *disk* failed with EIO, switching to access the disk remotely.

Explanation:
Local access to the disk failed. To avoid unmounting of the file system, the disk will now be accessed remotely.

User response:
Wait until work continuing on the local node completes. Then determine why local access to the disk failed, correct the problem and restart the daemon. This will cause GPFS to begin accessing the disk locally again.

6027-362 Attention: No disks were deleted, but some data was migrated. The file system may no longer be properly balanced.

Explanation:
The **mmdeldisk** command did not complete migrating data off the disks being deleted. The disks were restored to normal **ready**, status, but the migration has left the file system unbalanced. This may be caused by having too many disks unavailable or insufficient space to migrate all of the data to other disks.

User response:
Check disk availability and space requirements. Determine the reason that caused the command to end before successfully completing the migration and disk deletion. Reissue the **mmdeldisk** command.

6027-363 I/O error writing disk descriptor for disk *name*.

Explanation:
An I/O error occurred when the **mmadddisk** command was writing a disk descriptor on a disk. This could have been caused by either a configuration error or an error in the path to the disk.

User response:
Determine the reason the disk is inaccessible for writing and reissue the **mmadddisk** command.

6027-364 Error processing disks.

Explanation:
An error occurred when the **mmadddisk** command was reading disks in the file system.

User response:
Determine the reason why the disks are inaccessible for reading, then reissue the **mmadddisk** command.

6027-365 Rediscovered local access to *disk*.

Explanation:
Rediscovered local access to disk, which failed earlier with **EIO**. For good performance, the disk will now be accessed locally.

User response:
Wait until work continuing on the local node completes. This will cause GPFS to begin accessing the disk locally again.

6027-369 I/O error writing file system descriptor for disk *name*.

Explanation:
mmadddisk detected an I/O error while writing a file system descriptor on a disk.

User response:
Determine the reason the disk is inaccessible for writing and reissue the **mmadddisk** command.

6027-370 mmdeldisk completed.

Explanation:
The **mmdeldisk** command has completed.

User response:
None. Informational message only.

6027-371 Cannot delete all disks in the file system

Explanation:
An attempt was made to delete all the disks in a file system.

User response:
Either reduce the number of disks to be deleted or use the **mmdelfs** command to delete the file system.

6027-372 Replacement disk must be in the same failure group as the disk being replaced.

Explanation:
An improper failure group was specified for **mmrpldisk**.

User response:
Specify a failure group in the disk descriptor for the replacement disk that is the same as the failure group of the disk being replaced.

6027-373 Disk *diskName* is being replaced, so status of disk *diskName* must be replacement.

Explanation:
The **mmrpldisk** command failed when retrying a replace operation because the new disk does not have the correct status.

User response:
Issue the **mmlsdisk** command to display disk status. Then either issue the **mmchdisk** command to change the status of the disk to **replacement** or specify a new disk that has a status of **replacement**.

6027-374 Disk *name* may not be replaced.

Explanation:
A disk being replaced with **mmrpldisk** does not have a status of **ready** or **suspended**.

User response:
Use the **mmlsdisk** command to display disk status. Issue the **mmchdisk** command to change the status of the disk to be replaced to either **ready** or **suspended**.

6027-375 Disk name *diskName* already in file system.

Explanation:
The replacement disk name specified in the **mmrpldisk** command already exists in the file system.

User response:
Specify a different disk as the replacement disk.

6027-376 Previous replace command must be completed before starting a new one.

Explanation:
The **mmrpldisk** command failed because the status of other disks shows that a replace command did not complete.

User response:
Issue the **mmlsdisk** command to display disk status. Retry the failed **mmrpldisk** command or issue the **mmchdisk** command to change the status of the disks that have a status of **replacing** or **replacement**.

6027-377 Cannot replace a disk that is in use.

Explanation:
Attempting to replace a disk in place, but the disk specified in the **mmrpldisk** command is still available for use.

User response:
Use the **mmchdisk** command to stop GPFS's use of the disk.

6027-378 I/O still in progress near sector *number* on disk *name*.

Explanation:
The **mmfsd** daemon is attempting to terminate, but cannot because data structures in the daemon shared segment may still be referenced by kernel code. In particular, the daemon has started an I/O that has not yet completed. It is unsafe for the daemon to terminate until the I/O completes, because of asynchronous activity in the device driver that will access data structures belonging to the daemon.

User response:
Either wait for the I/O operation to time out, or issue a device-dependent command to terminate the I/O.

6027-379 Could not invalidate disk(s).

Explanation:
Trying to delete a disk and it could not be written to in order to invalidate its contents.

User response:
No action needed if removing that disk permanently. However, if the disk is ever to be used again, the **-v** flag must be specified with a value of **no** when using either the **mmcrfs** or **mmadddisk** command.

6027-380 Disk name missing from disk descriptor list entry *name*.

Explanation:
When parsing disk lists, no disks were named.

User response:
Check the argument list of the command.

6027-381 **Too many items in disk descriptor list entry** *name***.**

Explanation:
When parsing a disk descriptor, too many fields were specified for one disk.

User response:
Correct the disk descriptor to use the correct disk descriptor syntax.

6027-382 **Value** *value* **for the 'sector size' option for disk** *disk* **is not a multiple of** *value***.**

Explanation:
When parsing disk lists, the sector size given is not a multiple of the default sector size.

User response:
Specify a correct sector size.

6027-383 **Disk name** *name* **appears more than once.**

Explanation:
When parsing disk lists, a duplicate name is found.

User response:
Remove the duplicate name.

6027-384 **Disk name** *name* **already in file system.**

Explanation:
When parsing disk lists, a disk name already exists in the file system.

User response:
Rename or remove the duplicate disk.

6027-385 **Value** *value* **for the 'sector size' option for disk** *name* **is out of range. Valid values are** *number* **through** *number***.**

Explanation:
When parsing disk lists, the sector size given is not valid.

User response:
Specify a correct sector size.

6027-386 **Value** *value* **for the 'sector size' option for disk** *name* **is invalid.**

Explanation:
When parsing disk lists, the sector size given is not valid.

User response:
Specify a correct sector size.

6027-387 **Value** *value* **for the 'failure group' option for disk** *name* **is out of range. Valid values are** *number* **through** *number***.**

Explanation:
When parsing disk lists, the failure group given is not valid.

User response:
Specify a correct failure group.

6027-388 **Value** *value* **for the 'failure group' option for disk** *name* **is invalid.**

Explanation:
When parsing disk lists, the failure group given is not valid.

User response:
Specify a correct failure group.

6027-389 **Value** *value* **for the 'has metadata' option for disk** *name* **is out of range. Valid values are** *number* **through** *number***.**

Explanation:
When parsing disk lists, the 'has metadata' value given is not valid.

User response:
Specify a correct 'has metadata' value.

6027-390 **Value** *value* **for the 'has metadata' option for disk** *name* **is invalid.**

Explanation:
When parsing disk lists, the 'has metadata' value given is not valid.

User response:
Specify a correct 'has metadata' value.

6027-391 **Value** *value* **for the 'has data' option for disk** *name* **is out of range. Valid values are** *number* **through** *number***.**

Explanation:
When parsing disk lists, the 'has data' value given is not valid.

User response:
Specify a correct 'has data' value.

6027-392 **Value** *value* **for the 'has data' option for disk** *name* **is invalid.**

Explanation:
When parsing disk lists, the 'has data' value given is not valid.

User response:

Specify a correct 'has data' value.

6027-393 Either the 'has data' option or the 'has metadata' option must be '1' for disk *diskName*.

Explanation:
When parsing disk lists the 'has data' or 'has metadata' value given is not valid.

User response:
Specify a correct 'has data' or 'has metadata' value.

6027-394 Too many disks specified for file system. Maximum = *number*.

Explanation:
Too many disk names were passed in the disk descriptor list.

User response:
Check the disk descriptor list or the file containing the list.

6027-399 Not enough items in disk descriptor list entry, need *fields*.

Explanation:
When parsing a disk descriptor, not enough fields were specified for one disk.

User response:
Correct the disk descriptor to use the correct disk descriptor syntax.

6027-416 Incompatible file system descriptor version or not formatted.

Explanation:
Possible reasons for the error are:

1. A file system descriptor version that is not valid was encountered.
2. No file system descriptor can be found.
3. Disks are not correctly defined on all active nodes.
4. Disks, logical volumes, network shared disks, or virtual shared disks were incorrectly re-configured after creating a file system.

User response:
Verify:

1. The disks are correctly defined on all nodes.
2. The paths to the disks are correctly defined and operational.

6027-417 Bad file system descriptor.

Explanation:
A file system descriptor that is not valid was encountered.

User response:
Verify:

1. The disks are correctly defined on all nodes.
2. The paths to the disks are correctly defined and operational.

6027-418 Inconsistent file system quorum. readQuorum=*value* writeQuorum=*value* quorumSize=*value*.

Explanation:
A file system descriptor that is not valid was encountered.

User response:
Start any disks that have been stopped by the **mmchdisk** command or by hardware failures. If the problem persists, run offline **mmfsck**.

6027-419 Failed to read a file system descriptor.

Explanation:
Not enough valid replicas of the file system descriptor could be read from the file system.

User response:
Start any disks that have been stopped by the **mmchdisk** command or by hardware failures. Verify that paths to all disks are correctly defined and operational.

6027-420 Inode size must be greater than zero.

Explanation:
An internal consistency check has found a problem with file system parameters.

User response:
Record the above information. Contact the IBM Support Center.

6027-421 Inode size must be a multiple of logical sector size.

Explanation:
An internal consistency check has found a problem with file system parameters.

User response:
Record the above information. Contact the IBM Support Center.

6027-422 Inode size must be at least as large as the logical sector size.

Explanation:
An internal consistency check has found a problem with file system parameters.

User response:
Record the above information. Contact the IBM Support Center.

6027-423 Minimum fragment size must be a multiple of logical sector size.

Explanation:
An internal consistency check has found a problem with file system parameters.

User response:
Record the above information. Contact the IBM Support Center.

6027-424 Minimum fragment size must be greater than zero.

Explanation:
An internal consistency check has found a problem with file system parameters.

User response:
Record the above information. Contact the IBM Support Center.

6027-425 File system block size is larger than maxblocksize parameter.

Explanation:
The file system's block size is larger than the specified **maxblocksize** parameter.

User response:
Use the **mmchconfig maxblocksize=xxx** command to increase the maximum allowable block size.

6027-426 Warning: mount detected unavailable disks. Use mmlsdisk *fileSystem* **to see details.**

Explanation:
The **mount** command detected that some disks needed for the file system are unavailable.

User response:
Without file system replication enabled, the mount will fail. If it has replication, the mount may succeed depending on which disks are unavailable. Use **mmlsdisk** to see details of the disk status.

6027-427 Indirect block size must be at least as large as the minimum fragment size.

Explanation:
An internal consistency check has found a problem with file system parameters.

User response:
Record the above information. Contact the IBM Support Center.

6027-428 Indirect block size must be a multiple of the minimum fragment size.

Explanation:
An internal consistency check has found a problem with file system parameters.

User response:
Record the above information. Contact the IBM Support Center.

6027-429 Indirect block size must be less than full data block size.

Explanation:
An internal consistency check has found a problem with file system parameters.

User response:
Record the above information. Contact the IBM Support Center.

6027-430 Default metadata replicas must be less than or equal to default maximum number of metadata replicas.

Explanation:
An internal consistency check has found a problem with file system parameters.

User response:
Record the above information. Contact the IBM Support Center.

6027-431 Default data replicas must be less than or equal to default maximum number of data replicas.

Explanation:
An internal consistency check has found a problem with file system parameters.

User response:
Record the above information. Contact the IBM Support Center.

6027-432 Default maximum metadata replicas must be less than or equal to *value***.**

Explanation:
An internal consistency check has found a problem with file system parameters.

User response:
Record the above information. Contact the IBM Support Center.

6027-433 Default maximum data replicas must be less than or equal to *value***.**

Explanation:
An internal consistency check has found a problem with file system parameters.

User response:
Record the above information. Contact the IBM Support Center.

6027-434 Indirect blocks must be at least as big as inodes.

Explanation:
An internal consistency check has found a problem with file system parameters.

User response:
Record the above information. Contact the IBM Support Center.

6027-438 Duplicate disk name *name*.

Explanation:
An internal consistency check has found a problem with file system parameters.

User response:
Record the above information. Contact the IBM Support Center.

6027-439 Disk *name* sector size *value* does not match sector size *value* of other disk(s).

Explanation:
An internal consistency check has found a problem with file system parameters.

User response:
Record the above information. Contact the IBM Support Center.

6027-441 Unable to open disk *name*.

Explanation:
A disk name that is not valid was specified in **mmcrfs** command.

User response:
Correct the indicated command parameters.

6027-445 Value for option '-m' cannot exceed the number of metadata failure groups.

Explanation:
The current number of replicas of metadata cannot be larger than the number of failure groups that are enabled to hold metadata.

User response:
Use a smaller value for **-m** on the **mmchfs** command, or increase the number of failure groups by adding disks to the file system.

6027-446 Value for option '-r' cannot exceed the number of data failure groups.

Explanation:
The current number of replicas of data cannot be larger than the number of failure groups that are enabled to hold data.

User response:
Use a smaller value for **-r** on the **mmchfs** command, or increase the number of failure groups by adding disks to the file system.

6027-451 No disks= list found in mount options.

Explanation:
No 'disks=' clause found in the mount options list when opening a file system.

User response:
Check the operating system's file system database and local **mmsdrfs** file for this file system.

6027-452 No disks found in disks= list.

Explanation:
No disks listed when opening a file system.

User response:
Check the operating system's file system database and local **mmsdrfs** file for this file system.

6027-453 No disk name found in a clause of the list.

Explanation:
No disk name found in a clause of the**disks=** list.

User response:
Check the operating system's file system database and local **mmsdrfs** file for this file system.

6027-461 Unable to find *name* device.

Explanation:
Self explanatory.

User response:
There must be a **/dev/sg***name* special device defined. Check the error code. This could indicate a configuration error in the specification of disks, logical volumes, network shared disks, or virtual shared disks.

6027-462 *name* must be a char or block special device.

Explanation:
Opening a file system.

User response:
There must be a **/dev/sg***name* special device defined. This could indicate a configuration error in the

specification of disks, logical volumes, network shared disks, or virtual shared disks.

6027-463 SubblocksPerFullBlock was not 32.

Explanation:
The value of the SubblocksPerFullBlock variable was not 32. This situation should never exist, and indicates an internal error.

User response:
Record the above information and contact the IBM Support Center.

6027-465 The average file size must be at least as large as the minimum fragment size.

Explanation:
When parsing the command line of **tscrfs**, it was discovered that the average file size is smaller than the minimum fragment size.

User response:
Correct the indicated command parameters.

6027-468 Disk *name* **listed in** *fileName* **or local mmsdrfs file, not found in device** *name***. Run: mmcommon recoverfs** *name***.**

Explanation:
Tried to access a file system but the disks listed in the operating system's file system database or the local **mmsdrfs** file for the device do not exist in the file system.

User response:
Check the configuration and availability of disks. Run the **mmcommon recoverfs** *device* command. If this does not resolve the problem, configuration data in the SDR may be incorrect. If no user modifications have been made to the SDR, contact the IBM Support Center. If user modifications have been made, correct these modifications.

6027-469 File system *name* **does not match** *descriptor***.**

Explanation:
The file system name found in the descriptor on disk does not match the corresponding device name in **/etc/filesystems**.

User response:
Check the operating system's file system database.

6027-470 Disk *name* **may still belong to an active file system.**

Explanation:
The disk being added by the **mmcrfs**, **mmadddisk**, or **mmrpldisk** command appears to still belong to some file system.

User response:
Verify that the disks you are adding do not belong to an active file system, and use the **-v no** option to bypass this check. Use this option only if you are certain that no other file system has this disk configured because you can cause data corruption in both file systems if this is not the case.

6027-471 You have requested that the file system be upgraded to version *number***. This will enable new functionality but will prevent you from using the file system with earlier releases of GPFS. Do you want to continue?**

Explanation:
Verification request in response to the **mmchfs -V** command. This is a request to upgrade the file system and activate functions that are incompatible with the previous release of GPFS.

User response:
Enter **yes** if you want the conversion to take place.

6027-472 Failed to enable extended attribute support.

Explanation:
An error occurred while attempting to upgrade the file system to a version that supports extended attributes. The file to store extended attribute data could not be created. An additional GPFS message will be issued providing more specific information.

User response:
Follow the suggested action for the other message that occurs with this one.

6027-473 File System *fileSystem* **unmounted by the system with return code** *value* **and reason code** *value***.**

Explanation:
Console log entry caused by a forced unmount due to disk or communication failure.

User response:
Correct the underlying problem and remount the file system.

6027-474 Recovery Log I/O Failed, unmounting *fileSystem***.**

Explanation:
I/O to the recovery log failed.

User response:
Check the paths to all disks making up the file system. Run the **mmlsdisk** command to determine if GPFS has

declared any disks unavailable. Repair any paths to disks that have failed. Remount the file system.

6027-475 The option -F is not enabled. Use option -V to enable the most recent features.

Explanation:
mmchfs -F is not enabled under the current file system format version.

User response:
Issue **mmchfs -V** to change the file system format to the latest format supported.

Note: This is a request to upgrade the file system and activate functions that are incompatible with the previous release of GPFS.

6027-476 Restricted mount using only available file system descriptor.

Explanation:
Fewer than the necessary number of file system descriptors were successfully read. Using the best available descriptor to allow the restricted mount to continue.

User response:
Informational message only.

6027-477 The option -z is not enabled. Use the -V option to enable most recent features.

Explanation:
The file system format version does not support the **-z** option on the **mmchfs** command.

User response:
Change the file system format version by issuing **mmchfs -V**.

Note: This is a request to upgrade the file system and activate functions that are incompatible with the previous release of GPFS.

6027-478 The option -z could not be changed. *fileSystem* is still in use.

Explanation:
The file system is still mounted or another GPFS administration command (**mm...**) is running against the file system.

User response:
Unmount the file system if it is mounted, and wait for any command that is running to complete before reissuing the **mmchfs -z** command.

6027-479 This file system utilizes features that are not available in this version of GPFS.

Explanation:
The file system descriptor specifies a wide disk address, which is not available on the 32-bit version of GPFS.

User response:
Use the 64-bit version of GPFS with wide disk address support.

6027-480 Cannot enable DMAPI in a file system with existing snapshots.

Explanation:
The user is not allowed to enable DMAPI for a file system with existing snapshots.

User response:
Delete all existing snapshots in the file system and repeat the **mmchfs** command.

6027-481 Remount failed for mountid *id*: *errnoDescription*.

Explanation:
mmfsd restarted and tried to remount any file systems that the VFS layer thinks are still mounted.

User response:
Check the errors displayed and the **errno** description.

6027-482 Remount failed for device *id*: *errnoDescription*.

Explanation:
mmfsd restarted and tried to remount any file systems that the VFS layer thinks are still mounted.

User response:
Check the errors displayed and the **errno** description.

6027-483 Remounted *name*.

Explanation:
mmfsd restarted and remounted the specified file system because it was in the kernel's list of previously mounted file systems.

User response:
Informational message only.

6027-484 Remount failed for *device* after daemon restart.

Explanation:
A remount failed after daemon restart. This ordinarily occurs because one or more disks are unavailable. Other possibilities include loss of connectivity to one or more disks.

User response:

Issue the **mmlsdisk** command and check for **down** disks. Issue the **mmchdisk** command to start any **down** disks, then remount the file system. If there is another problem with the disks or the connections to the disks, take necessary corrective actions and remount the file system.

6027-485 Perform mmlsdisk for any disk failures and re-mount.

Explanation:
Occurs in conjunction with 6027-484.

User response:
Follow the *User Response* for 6027-484.

6027-486 No local device specified for *fileSystemName* **in** *clusterName***.**

Explanation:
While attempting to mount a remote file system from another cluster, GPFS was unable to determine the local device name for this file system.

User response:
There must be a **/dev/sgname** special device defined. Check the error code. This is probably a configuration error in the specification of a remote file system. Run **mmremotefs show** to check that the remote file system is properly configured.

6027-487 Failed to write the file system descriptor to disk *diskName***.**

Explanation:
An error occurred when **mmfsctl include** was writing a copy of the file system descriptor to one of the disks specified on the command line. This could have been caused by a failure of the corresponding disk device, or an error in the path to the disk.

User response:
Verify that the disks are correctly defined on all nodes. Verify that paths to all disks are correctly defined and operational.

6027-488 Error opening the exclusion disk file *fileName***.**

Explanation:
Unable to retrieve the list of excluded disks from an internal configuration file.

User response:
Ensure that GPFS executable files have been properly installed on all nodes. Perform required configuration steps prior to starting GPFS.

6027-489 Attention: The desired replication factor exceeds the number of available *dataOrMetadata* **failure groups in storage pool** *storagePool***. This is allowed, but files in this storage pool will not be replicated and will therefore be at risk.**

Explanation:
You have specified a number of replicas that exceeds the number of failure groups available.

User response:
Reissue the command with a smaller replication factor or increase the number of failure groups.

6027-490 The descriptor replica on disk *name* **has been excluded.**

Explanation:
The file system descriptor quorum has been overridden and, as a result, the specified disk was excluded from all operations on the file system descriptor quorum.

User response:
None. Informational message only.

6027-492 The file system is already at file system version *number*

Explanation:
The user tried to upgrade the file system format using **mmchfs -V --version=v**, but the specified version is smaller than the current version of the file system.

User response:
Specify a different value for the **--version** option.

6027-493 File system version *number* **is not supported on** *nodeName* **nodes in the cluster.**

Explanation:
The user tried to upgrade the file system format using **mmchfs -V**, but some nodes in the local cluster are still running an older GPFS release that does support the new format version.

User response:
Install a newer version of GPFS on those nodes.

6027-494 File system version *number* **is not supported on** *nodeName* **remote nodes mounting the file system.**

Explanation:
The user tried to upgrade the file system format using **mmchfs -V**, but the file system is still mounted on some nodes in remote clusters that do support the new format version.

User response:

Unmount the file system on the nodes that do not support the new format version.

6027-495 **You have requested that the file system be upgraded to version *number*. This will enable new functionality but will prevent you from using the file system with earlier releases of GPFS. Do you want to continue?**

Explanation:
Verification request in response to the **mmchfs -V full** command. This is a request to upgrade the file system and activate functions that are incompatible with a previous release of GPFS.

User response:
Enter **yes** if you want the conversion to take place.

6027-496 **You have requested that the file system version for local access be upgraded to version *number*. This will enable some new functionality but will prevent local nodes from using the file system with earlier releases of GPFS. Remote nodes are not affected by this change. Do you want to continue?**

Explanation:
Verification request in response to the **mmchfs -V** command. This is a request to upgrade the file system and activate functions that are incompatible with a previous release of GPFS.

User response:
Enter **yes** if you want the conversion to take place.

6027-497 **The file system has already been upgraded to *number* using -V full. It is not possible to revert back.**

Explanation:
The user tried to upgrade the file system format using **mmchfs -V compat**, but the file system has already been fully upgraded.

User response:
Informational message only.

6027-498 **Incompatible file system format. Only file systems formatted with GPFS 3.2.1.5 or later can be mounted on this platform.**

Explanation:
A user running GPFS on Microsoft Windows tried to mount a file system that was formatted with a version of GPFS that did not have Windows support.

User response:
Create a new file system using current GPFS code.

6027-500 ***name*** **loaded and configured.**

Explanation:
The kernel extension was loaded and configured.

User response:
None. Informational message only.

6027-501 ***name*****:module** ***moduleName*** **unloaded.**

Explanation:
The kernel extension was unloaded.

User response:
None. Informational message only.

6027-502 **Incorrect parameter:** ***name*****.**

Explanation:
mmfsmnthelp was called with an incorrect parameter.

User response:
Contact the IBM Support Center.

6027-504 **Not enough memory to allocate internal data structure.**

Explanation:
Self explanatory.

User response:
Increase ulimit or paging space

6027-505 **Internal error, aborting.**

Explanation:
Self explanatory.

User response:
Contact the IBM Support Center.

6027-506 **program:** ***loadFile*** **is already loaded at** ***address*****.**

Explanation:
The program was already loaded at the address displayed.

User response:
None. Informational message only.

6027-507 **program:** ***loadFile*** **is not loaded.**

Explanation:
The program could not be loaded.

User response:
None. Informational message only.

6027-510 **Cannot mount** *fileSystem* **on** *mountPoint*: *errorString*

Explanation:
There was an error mounting the GPFS file system.

User response:
Determine action indicated by the error messages and error log entries. Errors in the disk path often cause this problem.

6027-511 **Cannot unmount** *fileSystem*: *errorDescription*

Explanation:
There was an error unmounting the GPFS file system.

User response:
Take the action indicated by **errno** description.

6027-512 *name* **not listed in /etc/vfs**

Explanation:
Error occurred while installing the GPFS kernel extension, or when trying to mount a file system.

User response:
Check for the **mmfs** entry in **/etc/vfs**

6027-514 **Cannot mount** *fileSystem* **on** *mountPoint*: **Already mounted.**

Explanation:
An attempt has been made to mount a file system that is already mounted.

User response:
None. Informational message only.

6027-515 **Cannot mount** *fileSystem* **on** *mountPoint*

Explanation:
There was an error mounting the named GPFS file system. Errors in the disk path usually cause this problem.

User response:
Take the action indicated by other error messages and error log entries.

6027-516 **Cannot mount** *fileSystem*

Explanation:
There was an error mounting the named GPFS file system. Errors in the disk path usually cause this problem.

User response:
Take the action indicated by other error messages and error log entries.

6027-517 **Cannot mount** *fileSystem*: *errorString*

Explanation:
There was an error mounting the named GPFS file system. Errors in the disk path usually cause this problem.

User response:
Take the action indicated by other error messages and error log entries.

6027-518 **Cannot mount** *fileSystem*: **Already mounted.**

Explanation:
An attempt has been made to mount a file system that is already mounted.

User response:
None. Informational message only.

6027-519 **Cannot mount** *fileSystem* **on** *mountPoint*: **File system table full.**

Explanation:
An attempt has been made to mount a file system when the file system table is full.

User response:
None. Informational message only.

6027-520 **Cannot mount** *fileSystem*: **File system table full.**

Explanation:
An attempt has been made to mount a file system when the file system table is full.

User response:
None. Informational message only.

6027-530 **Mount of** *name* **failed: cannot mount restorable file system for read/write.**

Explanation:
A file system marked as **enabled** for restore cannot be mounted **read/write**.

User response:
None. Informational message only.

6027-531 **The following disks of** *name* **will be formatted on node** *nodeName*: *list*.

Explanation:
Output showing which disks will be formatted by the **mmcrfs** command.

User response:
None. Informational message only.

6027-532 The quota record *number* in file *fileName* is not valid.

Explanation:
A quota entry contained a checksum that is not valid.

User response:
Remount the file system with quotas disabled. Restore the quota file from back up, and run **mmcheckquota**.

6027-533 File system *fileSystem* is approaching the limit for the maximum number of inodes/files.

Explanation:
The number of files created is approaching the file system limit.

User response:
Increase the maximum number of files with the **mmchfs** command to avoid reaching the inode limit and possible performance degradation.

6027-534 Cannot create a snapshot in a DMAPI-enabled file system, rc=*returnCode*.

Explanation:
You cannot create a snapshot in a DMAPI-enabled file system.

User response:
Use the **mmchfs** command to disable DMAPI, and reissue the command.

6027-535 Disks up to size *size* can be added to storage pool '*pool*'.

Explanation:
Based on the parameters given to **mmcrfs** and the size and number of disks being formatted, GPFS has formatted its allocation maps to allow disks up the given size to be added to this storage pool by the **mmadddisk** command.

User response:
None. Informational message only. If the reported maximum disk size is smaller than necessary, delete the file system with **mmdelfs** and rerun **mmcrfs** with either larger disks or a larger value for the **-n** parameter.

6027-536 Insufficient system memory to run GPFS daemon. Reduce page pool memory size with the mmchconfig command or add additional RAM to system.

Explanation:
Insufficient memory for GPFS internal data structures with current system and GPFS configuration.

User response:
Reduce page pool usage with the **mmchconfig** command, or add additional RAM to system.

6027-537 Disks up to size *size* can be added to this file system.

Explanation:
Based on the parameters given to the **mmcrfs** command and the size and number of disks being formatted, GPFS has formatted its allocation maps to allow disks up the given size to be added to this file system by the **mmadddisk** command.

User response:
None, informational message only. If the reported maximum disk size is smaller than necessary, delete the file system with **mmdelfs** and reissue the **mmcrfs** command with larger disks or a larger value for the **-n** parameter.

6027-538 Error accessing disks.

Explanation:
The **mmcrfs** command encountered an error accessing one or more of the disks.

User response:
Verify that the disk descriptors are coded correctly and that all named disks exist and are online.

6027-539 Unable to clear descriptor areas for *fileSystem*.

Explanation:
The **mmdelfs** command encountered an error while invalidating the file system control structures on one or more disks in the file system being deleted.

User response:
If the problem persists, specify the **-p** option on the **mmdelfs** command.

6027-540 Formatting file system.

Explanation:
The **mmcrfs** command began to write file system data structures onto the new disks.

User response:
None. Informational message only.

6027-541 Error formatting file system.

Explanation:
mmcrfs command encountered an error while formatting a new file system. This is often an I/O error.

User response:
Check the subsystems in the path to the disk. Follow the instructions from other messages that appear with this one.

6027-542 Fileset *filesetName* (id *filesetId*) has been incompletely deleted.

Explanation:
A fileset delete operation was interrupted, leaving this fileset in an incomplete state.

User response:
Reissue the fileset delete command.

6027-543 Error writing file system descriptor for*fileSystem*.

Explanation:
The **mmcrfs** command could not successfully write the file system descriptor in a particular file system. Check the subsystems in the path to the disk. This is often an I/O error.

User response:
Check system error log, rerun **mmcrfs**.

6027-544 Could not invalidate *disk* of *fileSystem*.

Explanation:
A disk could not be written to invalidate its contents. Check the subsystems in the path to the disk. This is often an I/O error.

User response:
Ensure the indicated logical volume is writable.

6027-545 Error processing fileset metadata file.

Explanation:
There is no I/O path to critical metadata or metadata has been corrupted.

User response:
Verify that the I/O paths to all disks are valid and that all disks are either in the 'recovering' or 'up' availability states. If all disks are available and the problem persists, issue the **mmfsck** command to repair damaged metadata

6027-546 Error processing allocation map for storage pool *poolName*.

Explanation:
There is no I/O path to critical metadata, or metadata has been corrupted.

User response:
Verify that the I/O paths to all disks are valid, and that all disks are either in the 'recovering' or 'up' availability. Issue the **mmlsdisk** command.

6027-547 Fileset *filesetName* was unlinked.

Explanation:
Fileset was already unlinked.

User response:
None. Informational message only.

6027-548 Fileset '*filesetName*' unlinked from '*filesetName*'.

Explanation:
A fileset being deleted contains junctions to other filesets. The cited fileset were unlinked.

User response:
None. Informational message only.

6027-549 Failed to open *name*.

Explanation:
The **mount** command was unable to access a file system. Check the subsystems in the path to the disk. This is often an I/O error.

User response:
Follow the suggested actions for the other messages that occur with this one.

6027-550 Allocation manager for *fileSystem* failed to revoke ownership from node *nodeName*.

Explanation:
An irrecoverable error occurred trying to revoke ownership of an allocation region. The allocation manager has panicked the file system to prevent corruption of on-disk data.

User response:
Remount the file system.

6027-551 *fileSystem* is still in use.

Explanation:
The **mmdelfs** or **mmcrfs** command found that the named file system is still mounted or that another GPFS command is running against the file system.

User response:
Unmount the file system if it is mounted, or wait for GPFS commands in progress to terminate before retrying the command.

6027-552 Scan completed successfully.

Explanation:
The scan function has completed without error.

User response:
None. Informational message only.

6027-553 Scan failed on *number* user or system files.

Explanation:
Data may be lost as a result of pointers that are not valid or unavailable disks.

User response:
Some files may have to be restored from backup copies. Issue the **mmlsdisk** command to check the availability of all the disks that make up the file system.

6027-554 Scan failed on *number* out of *number* user or system files.

Explanation:
Data may be lost as a result of pointers that are not valid or unavailable disks.

User response:
Some files may have to be restored from backup copies. Issue the **mmlsdisk** command to check the availability of all the disks that make up the file system.

6027-555 The desired replication factor exceeds the number of available failure groups.

Explanation:
You have specified a number of replicas that exceeds the number of failure groups available.

User response:
Reissue the command with a smaller replication factor or increase the number of failure groups.

6027-556 Not enough space for the desired number of replicas.

Explanation:
In attempting to restore the correct replication, GPFS ran out of space in the file system. The operation can continue but some data is not fully replicated.

User response:
Make additional space available and reissue the command.

6027-557 Not enough space or available disks to properly balance the file.

Explanation:
In attempting to stripe data within the file system, data was placed on a disk other than the desired one. This is normally not a problem.

User response:
Run **mmrestripefs** to rebalance all files.

6027-558 Some data are unavailable.

Explanation:
An I/O error has occurred or some disks are in the stopped state.

User response:
Check the availability of all disks by issuing the **mmlsdisk** command and check the path to all disks. Reissue the command.

6027-559 Some data could not be read or written.

Explanation:
An I/O error has occurred or some disks are in the stopped state.

User response:
Check the availability of all disks and the path to all disks, and reissue the command.

6027-560 File system is already suspended.

Explanation:
The **tsfsctl** command was asked to suspend a suspended file system.

User response:
None. Informational message only.

6027-561 Error migrating log.

Explanation:
There are insufficient available disks to continue operation.

User response:
Restore the unavailable disks and reissue the command.

6027-562 Error processing inodes.

Explanation:
There is no I/O path to critical metadata or metadata has been corrupted.

User response:
Verify that the I/O paths to all disks are valid and that all disks are either in the `recovering` or `up` availability. Issue the **mmlsdisk** command.

6027-563 File system is already running.

Explanation:
The **tsfsctl** command was asked to resume a file system that is already running.

User response:
None. Informational message only.

6027-564 Error processing inode allocation map.

Explanation:
There is no I/O path to critical metadata or metadata has been corrupted.

User response:
Verify that the I/O paths to all disks are valid and that all disks are either in the `recovering` or `up` availability. Issue the **mmlsdisk** command.

6027-565 Scanning user file metadata.

Explanation:
Self explanatory.

User response:
None. Informational message only.

6027-566 Error processing user file metadata.

Explanation:
Self explanatory.

User response:
None. Informational message only.

6027-567 Waiting for pending file system scan to finish.

Explanation:
Self explanatory.

User response:
None. Informational message only.

6027-568 Waiting for *number* pending file system scans to finish.

Explanation:
Self explanatory.

User response:
None. Informational message only.

6027-569 Incompatible parameters. Unable to allocate space for file system metadata. Change one or more of the following as suggested and try again:

Explanation:
Incompatible file system parameters were detected.

User response:
Refer to the details given and correct the file system parameters.

6027-570 Incompatible parameters. Unable to create file system. Change one or more of the following as suggested and try again:

Explanation:
Incompatible file system parameters were detected.

User response:
Refer to the details given and correct the file system parameters.

6027-571 Logical sector size *value* must be the same as disk sector size.

Explanation:
This message is produced by the **mmcrfs** command if the sector size given by the **-l** option is not the same as the sector size given for disks in the **-d** option.

User response:
Correct the options and reissue the command.

6027-572 Completed creation of file system *fileSystem*.

Explanation:
The **mmcrfs** command has successfully completed.

User response:
None. Informational message only.

6027-573 All data on following disks of *fileSystem* will be destroyed:

Explanation:
Produced by the **mmdelfs** command to list the disks in the file system that is about to be destroyed. Data stored on the disks will be lost.

User response:
None. Informational message only.

6027-574 Completed deletion of file system *fileSystem*.

Explanation:
The **mmdelfs** command has successfully completed.

User response:
None. Informational message only.

6027-575 Unable to complete low level format for *fileSystem*.

Explanation:
The **mmcrfs** command was unable to create the low level file structures for the file system.

User response:
Check other error messages and the error log. This is usually an error accessing disks.

6027-576 Storage pools have not been enabled for file system *fileSystem*.

Explanation:
User invoked a command with a storage pool option (**-p** or **-P**) before storage pools were enabled.

User response:
Enable storage pools with the **mmchfs -V** command, or correct the command invocation and reissue the command.

6027-577 Attention: *number* user or system files are not properly replicated.

Explanation:
GPFS has detected files that are not replicated correctly due to a previous failure.

User response:
Issue the **mmrestripefs** command at the first opportunity.

6027-578 Attention: *number* out of *number* user or system files are not properly replicated.

Explanation:
GPFS has detected files that are not replicated correctly due to a previous failure.

User response:
Issue the **mmrestripefs** command at the first opportunity.

6027-579 Some unreplicated file system metadata has been lost. File system usable only in restricted mode.

Explanation:
A disk was deleted that contained vital file system metadata that was not replicated.

User response:
Mount the file system in restricted mode (**-o rs**) and copy any user data that may be left on the file system. Then delete the file system.

6027-580 Unable to access vital system metadata, too many disks are unavailable.

Explanation:
Metadata is unavailable because the disks on which the data reside are stopped, or an attempt was made to delete them.

User response:
Either start the stopped disks, try to delete the disks again, or recreate the file system.

6027-581 Unable to access vital system metadata, file system corrupted.

Explanation:
When trying to access the files system, the metadata was unavailable due to a disk being deleted.

User response:
Determine why a disk is unavailable.

6027-582 Some data has been lost.

Explanation:
An I/O error has occurred or some disks are in the stopped state.

User response:
Check the availability of all disks by issuing the **mmlsdisk** command and check the path to all disks. Reissue the command.

6027-583 Disk *name* cannot be added to storage pool *name*. Allocation map cannot accommodate disks larger than *size* MB.

Explanation:
Cannot add a disk that is so much larger than the disks initially used to create the file system that the allocation map cannot accommodate the new disk.

User response:
Break the disk into pieces smaller than the size indicated in the message by creating logical volumes, or recreate the file system.

6027-584 Incompatible parameters. Unable to allocate space for root directory. Change one or more of the following as suggested and try again:

Explanation:
Inconsistent parameters have been passed to the **mmcrfs** command, which would result in the creation of an inconsistent file system. Suggested parameter changes are given.

User response:
Reissue the **mmcrfs** command with the suggested parameter changes.

6027-585 Incompatible parameters. Unable to allocate space for ACL data. Change one or more of the following as suggested and try again:

Explanation:
Inconsistent parameters have been passed to the **mmcrfs** command, which would result in the creation of an inconsistent file system. The parameters entered require more space than is available. Suggested parameter changes are given.

User response:
Reissue the **mmcrfs** command with the suggested parameter changes.

6027-586 Quota server initialization failed.

Explanation:
Quota server initialization has failed. This message may appear as part of the detail data in the quota error log.

User response:
Check status and availability of the disks. If quota files have been corrupted, restore them from the last available backup. Finally, reissue the command.

6027-587 **Unable to initialize quota client because there is no quota server. Please check error log on the file system manager node. The mmcheckquota command must be run with the file system unmounted before retrying the command.**

Explanation:
`startQuotaClient` failed.

User response:
If the quota file could not be read (check error log on file system manager. Issue the **mmlsmgr** command to determine which node is the file system manager), then the **mmcheckquota** command must be run with the file system unmounted.

6027-588 **No more than** *number* **nodes can mount a file system.**

Explanation:
The limit of the number of nodes that can mount a file system was exceeded.

User response:
Observe the stated limit for how many nodes can mount a file system.

6027-589 **Scanning file system metadata, phase** *number*

Explanation:
Self-explanatory.

User response:
None. Informational message only.

6027-590 **GPFS is experiencing a shortage of pagepool. This message will not be repeated for at least one hour.**

Explanation:
Pool starvation occurs, buffers have to be continually stolen at high aggressiveness levels.

User response:
Issue the **mmchconfig** command to increase the size of **pagepool**.

6027-591 **Unable to allocate sufficient inodes for file system metadata. Increase the value for** *option* **and try again.**

Explanation:
Too few inodes have been specified on the **-N** option of the **mmcrfs** command.

User response:
Increase the value for the **-N** option and reissue the **mmcrfs** command

---

6027-592 **Mount of** *fileSystem* **is waiting for the mount disposition to be set by some data management application.**

Explanation:
Data management utilizing DMAPI is enabled for the file system, but no data management application has set a disposition for the mount event.

User response:
Start the data management application and verify that the application sets the mount disposition.

6027-593 **The root quota entry is not found in its assigned record.**

Explanation:
On mount, the root entry is not found in the first record of the quota file.

User response:
Issue the **mmcheckquota** command to verify that the use of root has not been lost.

6027-594 **Disk** *diskName* **cannot be added to storage pool** *poolName***. Allocation map cannot accommodate disks larger than** *number* **MB.**

Explanation:
Cannot add a disk that is so much larger than the disks initially used to create the file system that the allocation map cannot accommodate the new disk.

User response:
Break the disk into pieces smaller than the size indicated in the message by creating logical volumes, or recreate the file system.

6027-595 **While creating quota files, file** *file name***, with no valid quota information, was found in the root directory. Please remove files with reserved quota file names (for example, user.quota) without valid quota information from the root directory by: 1 mounting the file system without quotas. 2. removing the files. 3. remounting the file system with quotas to recreate new quota files. To use quota file names other than the reserved names, use the mmcheckquota command.**

Explanation:
While mounting a file system, the state of the file system descriptor indicates that quota files do not exist. However, files that do not contain quota information but

have one of the reserved names: **user.quota**, **group.quota**, or **fileset.quota** exist in the root directory.

User response:
To mount the file system so that new quota files will be created, perform these steps:

1. Mount the file system without quotas.
2. Verify that there are no files in the root directory with the reserved names: **user.quota**, **group.quota**, or **fileset.quota**.
3. Remount the file system with quotas. To mount the file system with other files used as quota files, issue the **mmcheckquota** command.

6027-596 **While creating quota files, file** *file name* **containing quota information was found in the root directory. This file will be used as** *quota type* **quota file.**

Explanation:
While mounting a file system, the state of the file system descriptor indicates that quota files do not exist. However, files that have one of the reserved names **user.quota**, **group.quota**, or **fileset.quota** and contain quota information, exist in the root directory. The file with the reserved name will be used as the quota file.

User response:
None. Informational message.

6027-597 **The quota command is requested to process quotas for a type (user, group, or fileset), which is not enabled**

Explanation:
A quota command is requested to process quotas for a **user**, **group**, or **fileset** quota type, which is not enabled.

User response:
Verify that the **user**, **group**, or **fileset** quota type is enabled and reissue the command.

6027-598 **The supplied file does not contain quota information.**

Explanation:
A file supplied as a quota file does not contain quota information.

User response:
Change the file so it contains valid quota information and reissue the command.

To mount the file system so that new quota files are created:

1. Mount the file system without quotas.
2. Verify there are no files in the root directory with the reserved **user.quota** or **group.quota** name.
3. Remount the file system with quotas.

6027-599 **File supplied to the command does not exist in the root directory.**

Explanation:
The user-supplied name of a new quota file has not been found.

User response:
Ensure that a file with the supplied name exists. Then reissue the command.

6027-600 **On node** *nodeName* **an earlier error may have caused some file system data to be inaccessible at this time. Check error log for additional information. After correcting the problem, the file system can be mounted again to restore normal data access.**

Explanation:
An earlier error may have caused some file system data to be inaccessible at this time.

User response:
Check the error log for additional information. After correcting the problem, the file system can be mounted again.

6027-601 **Error changing pool size.**

Explanation:
The **mmchconfig** command failed to change the pool size to the requested value.

User response:
Follow the suggested actions in the other messages that occur with this one.

6027-602 **ERROR: file system not mounted. Mount file system** *fileSystem* **and retry command.**

Explanation:
A GPFS command that requires the file system be mounted was issued.

User response:
Mount the file system and reissue the command.

6027-603 **Current pool size:** *value***K =** *value***M, max block size:** *value***K =** *value***M.**

Explanation:
Displays the current pool size.

User response:
None. Informational message only.

6027-604 Parameter incompatibility. File system block size is larger than maxblocksize parameter.

Explanation:
An attempt is being made to mount a file system whose block size is larger than the **maxblocksize** parameter as set by **mmchconfig**.

User response:
Use the **mmchconfig maxblocksize=xxx** command to increase the maximum allowable block size.

6027-605 File system has been renamed.

Explanation:
Self-explanatory.

User response:
None. Informational message only.

6027-606 The node number *nodeNumber* is not defined in the node list.

Explanation:
A node matching *nodeNumber* was not found in the GPFS configuration file.

User response:
Perform required configuration steps prior to starting GPFS on the node.

6027-607 mmcommon getEFOptions *fileSystem* failed. Return code *value*.

Explanation:
The **mmcommon getEFOptions** command failed while looking up the names of the disks in a file system. This error usually occurs during **mount** processing.

User response:
Check the preceding messages. A frequent cause for such errors is lack of space in **/var**.

6027-608 file system manager takeover failed.

Explanation:
An attempt to takeover as file system manager failed. The file system is unmounted to allow another node to try.

User response:
Check the return code. This is usually due to network or disk connectivity problems. Issue the **mmlsdisk** command to determine if the paths to the disk are unavailable, and issue the **mmchdisk** if necessary.

6027-609 File system *fileSystem* unmounted because it does not have a manager.

Explanation:
The file system had to be unmounted because a file system manager could not be assigned. An accompanying message tells which node was the last manager.

User response:
Examine error log on the last file system manager. Issue the **mmlsdisk** command to determine if a number of disks are down. Examine the other error logs for an indication of network, disk, or virtual shared disk problems. Repair the base problem and issue the **mmchdisk** command if required.

6027-610 Cannot mount file system *fileSystem* because it does not have a manager.

Explanation:
The file system had to be unmounted because a file system manager could not be assigned. An accompanying message tells which node was the last manager.

User response:
Examine error log on the last file system manager node. Issue the **mmlsdisk** command to determine if a number of disks are down. Examine the other error logs for an indication of disk or network shared disk problems. Repair the base problem and issue the **mmchdisk** command if required.

6027-611 Recovery: *file system*, delay *number* sec. for safe recovery.

Explanation:
Informational. When disk leasing is in use, wait for the existing lease to expire before performing log and token manager recovery.

User response:
None.

6027-612 Unable to run *command* while the file system is suspended.

Explanation:
A command that can alter data in a file system was issued while the file system was suspended.

User response:
Resume the file system and reissue the command.

6027-613 Expel *node* request from *node*. Expelling: *node*.

Explanation:
One node is asking to have another node expelled from the cluster, usually because they have communications problems between them. The cluster manager node will decide which one will be expelled.

User response:
Check that the communications paths are available between the two nodes.

6027-614 Value *value* for option *name* is out of range. Valid values are *number* through *number*.

Explanation:
The value for an option in the command line arguments is out of range.

User response:
Correct the command line and reissue the command.

6027-615 mmcommon getContactNodes *clusterName* failed. Return code *value*.

Explanation:
mmcommon getContactNodes failed while looking up contact nodes for a remote cluster, usually while attempting to mount a file system from a remote cluster.

User response:
Check the preceding messages, and consult the earlier chapters of this document. A frequent cause for such errors is lack of space in **/var**.

6027-616 Duplicate address *ipAddress* in node list.

Explanation:
The IP address appears more than once in the node list file.

User response:
Check the node list shown by the **mmlscluster** command.

6027-617 Recovered *number* nodes for cluster *clusterName*.

Explanation:
The asynchronous part (phase 2) of node failure recovery has completed.

User response:
None. Informational message only.

6027-618 Local host not found in node list (local ip interfaces: *interfaceList*).

Explanation:
The local host specified in the node list file could not be found.

User response:
Check the node list shown by the **mmlscluster** command.

6027-619 Negative grace times are not allowed.

Explanation:
The **mmedquota** command received a negative value for the **-t** option.

User response:
Reissue the **mmedquota** command with a nonnegative value for grace time.

6027-620 Hard quota limit must not be less than soft limit.

Explanation:
The hard quota limit must be greater than or equal to the soft quota limit.

User response:
Reissue the **mmedquota** command and enter valid values when editing the information.

6027-621 Negative quota limits are not allowed.

Explanation:
The quota value must be positive.

User response:
Reissue the **mmedquota** command and enter valid values when editing the information.

6027-622 Failed to join remote cluster *name*.

Explanation:
The node was not able to establish communication with another cluster, usually while attempting to mount a file system from a remote cluster.

User response:
Check other console messages for additional information. Verify that contact nodes for the remote cluster are set correctly. Run **mmremotefs show** and **mmremotecluster show** to display information about the remote cluster.

6027-623 All disks up and ready.

Explanation:
Self-explanatory.

User response:
None. Informational message only.

6027-624 No disks

Explanation:
Self-explanatory.

User response:
None. Informational message only.

6027-625 Migrate already pending.

Explanation:
A request to migrate the file system manager failed because a previous migrate request has not yet completed.

User response:
None. Informational message only.

6027-626 Migrate to node *nodeName* already pending.

Explanation:
A request to migrate the file system manager failed because a previous migrate request has not yet completed.

User response:
None. Informational message only.

6027-627 Node *nodeName* is already manager for *fileSystem*.

Explanation:
A request has been made to change the file system manager node to the node that is already the manager.

User response:
None. Informational message only.

6027-628 Sending migrate request to current manager node *nodeName*.

Explanation:
A request has been made to change the file system manager node.

User response:
None. Informational message only.

6027-629 Node *nodeName* resigned as manager for *fileSystem*.

Explanation:
Progress report produced by the **mmchmgr** command.

User response:
None. Informational message only.

6027-630 Node *nodeName* appointed as manager for *fileSystem*.

Explanation:
The **mmchmgr** command successfully changed the node designated as the file system manager.

User response:
None. Informational message only.

6027-631 Failed to appoint node *nodeName* as manager for *fileSystem*.

Explanation:
A request to change the file system manager node has failed.

User response:
Accompanying messages will describe the reason for the failure. Also, see the **mmfs.log** file on the target node.

6027-632 Failed to appoint a new manager for *fileSystem*.

Explanation:
An attempt to change the file system manager node has failed.

User response:
Accompanying messages will describe the reason for the failure. Also, see the **mmfs.log** file on the target node.

6027-633 Best choice node *nodeName* already manager for *fileSystem*.

Explanation:
Informational message about the progress and outcome of a migrate request.

User response:
None. Informational message only.

6027-634 Node name or number *node* is not valid.

Explanation:
A node number, IP address, or host name that is not valid has been entered in the configuration file or as input for a command.

User response:
Validate your configuration information and the condition of your network. This message may result from an inability to translate a node name.

6027-635 The current file system manager failed and no new manager will be appointed.

Explanation:
The file system manager node could not be replaced. This is usually caused by other system errors, such as disk or communication errors.

User response:
See accompanying messages for the base failure.

6027-636 Disks marked as stopped or offline.

Explanation:
A disk continues to be marked **down** due to a previous error and was not opened again.

User response:
Check the disk status by issuing the **mmlsdisk** command, then issue the **mmchdisk start** command to restart the disk.

6027-637 RVSD is not active.

Explanation:
The RVSD subsystem needs to be activated.

User response:

See the appropriate IBM Reliable Scalable Cluster Technology (RSCT) document at: publib.boulder.ibm.com/clresctr/windows/public/rsctbooks.html and search on *diagnosing IBM Virtual Shared Disk problems*.

6027-638 File system *fileSystem* unmounted by node *nodeName*.

Explanation:
Produced in the console log on a forced unmount of the file system caused by disk or communication failures.

User response:
Check the error log on the indicated node. Correct the underlying problem and remount the file system.

6027-639 File system cannot be mounted in restricted mode and ro or rw concurrently.

Explanation:
There has been an attempt to concurrently mount a file system on separate nodes in both a normal mode and in 'restricted' mode.

User response:
Decide which mount mode you want to use, and use that mount mode on both nodes.

6027-640 File system is mounted.

Explanation:
A command has been issued that requires that the file system be unmounted.

User response:
Unmount the file system and reissue the command.

6027-641 Unable to access vital system metadata. Too many disks are unavailable or the file system is corrupted.

Explanation:
An attempt has been made to access a file system, but the metadata is unavailable. This can be caused by:

1. The disks on which the metadata resides are either stopped or there was an unsuccessful attempt to delete them.
2. The file system is corrupted.

User response:
To access the file system:

1. If the disks are the problem either start the stopped disks or try to delete them.
2. If the file system has been corrupted, you will have to recreate it from backup medium.

6027-642 File system has been deleted.

Explanation:
Self-explanatory.

User response:
None. Informational message only.

6027-643 Node *nodeName* completed takeover for *fileSystem*.

Explanation:
The **mmchmgr** command completed successfully.

User response:
None. Informational message only.

6027-644 The previous error was detected on node *nodeName*.

Explanation:
An unacceptable error was detected. This usually occurs when attempting to retrieve file system information from the operating system's file system database or the cached GPFS system control data. The message identifies the node where the error was encountered.

User response:
See accompanying messages for the base failure. A common cause for such errors is lack of space in **/var**.

6027-645 Attention: mmcommon getEFOptions *fileSystem* failed. Checking *fileName*.

Explanation:
The names of the disks in a file system were not found in the cached GPFS system data, therefore an attempt will be made to get the information from the operating system's file system database.

User response:
If the command fails, see "File system will not mount" on page 61. A common cause for such errors is lack of space in **/var**.

6027-646 File system unmounted due to loss of cluster membership.

Explanation:
Quorum was lost, causing file systems to be unmounted.

User response:
Get enough nodes running the GPFS daemon to form a quorum.

6027-647 File *fileName* could not be run

Explanation:
The named shell script could not be processed. This message is followed by the error string returned by exec.

User response:
Check file existence and access permissions.

6027-648 EDITOR environment variable must be an absolute pathname.

Explanation:
The value of the EDITOR environment variable is not an absolute path name.

User response:
Change the value of the EDITOR environment variable to an absolute path name.

6027-649 Error reading the mmpmon command file.

Explanation:
An error occurred when reading the **mmpmon** command file.

User response:
Check file existence and access permissions.

6027-650 The mmfs daemon is shutting down abnormally.

Explanation:
The GPFS daemon is shutting down as a result of an irrecoverable condition, typically a resource shortage.

User response:
Review error log entries, correct a resource shortage condition, and restart the GPFS daemon.

6027-660 Error displaying message from mmfsd.

Explanation:
GPFS could not properly display an output string sent from the **mmfsd** daemon due to some error. A description of the error follows.

User response:
Check that GPFS is properly installed.

6027-661 mmfsd waiting for primary node *nodeName*.

Explanation:
The **mmfsd** server has to wait during start up because **mmfsd** on the primary node is not yet ready.

User response:
None. Informational message only.

6027-662 mmfsd timed out waiting for primary node *nodeName*.

Explanation:
The **mmfsd** server is about to terminate.

User response:
Ensure that the **mmfs.cfg** configuration file contains the correct host name or IP address of the primary node. Check **mmfsd** on the primary node.

6027-663 Lost connection to file system daemon.

Explanation:
The connection between a GPFS command and the **mmfsd** daemon has broken. The daemon has probably crashed.

User response:
Ensure that the **mmfsd** daemon is running. Check the error log.

6027-664 Unexpected message from file system daemon.

Explanation:
The version of the **mmfsd** daemon does not match the version of the GPFS command.

User response:
Ensure that all GPFS software components are at the same version.

6027-665 Failed to connect to file system daemon: *errorString*.

Explanation:
An error occurred while trying to create a session with **mmfsd**.

User response:
Ensure that the **mmfsd** daemon is running. Also, only root can run most GPFS commands. The mode bits of the commands must be **set-user-id** to **root**.

6027-666 Failed to determine file system manager.

Explanation:
While running a GPFS command in a multiple node configuration, the local file system daemon is unable to determine which node is managing the file system affected by the command.

User response:
Check internode communication configuration and ensure that enough GPFS nodes are up to form a quorum.

6027-667 Could not set up socket

Explanation:
One of the calls to create or bind the socket used for sending parameters and messages between the command and the daemon failed.

User response:
Check additional error messages.

6027-668 Could not send message to file system daemon

Explanation:
Attempt to send a message to the file system failed.

User response:
Check if the file system daemon is up and running.

6027-669 Could not connect to file system daemon.

Explanation:
The TCP connection between the command and the daemon could not be established.

User response:
Check additional error messages.

6027-670 Value for '*option*' is not valid. Valid values are *list*.

Explanation:
The specified value for the given command option was not valid. The remainder of the line will list the valid keywords.

User response:
Correct the command line.

6027-671 Keyword missing or incorrect.

Explanation:
A missing or incorrect keyword was encountered while parsing command line arguments

User response:
Correct the command line.

6027-672 Too few arguments specified.

Explanation:
Too few arguments were specified on the command line.

User response:
Correct the command line.

6027-673 Too many arguments specified.

Explanation:
Too many arguments were specified on the command line.

User response:
Correct the command line.

6027-674 Too many values specified for option *name*.

Explanation:
Too many values were specified for the given option on the command line.

User response:
Correct the command line.

6027-675 Required value for *option* is missing.

Explanation:
A required value was not specified for the given option on the command line.

User response:
Correct the command line.

6027-676 Option *option* specified more than once.

Explanation:
The named option was specified more than once on the command line.

User response:
Correct the command line.

6027-677 Option *option* is incorrect.

Explanation:
An incorrect option was specified on the command line.

User response:
Correct the command line.

6027-678 Misplaced or incorrect parameter *name*.

Explanation:
A misplaced or incorrect parameter was specified on the command line.

User response:
Correct the command line.

6027-679 Device *name* is not valid.

Explanation:
An incorrect device name was specified on the command line.

User response:
Correct the command line.

6027-681 Required option *name* was not specified.

Explanation:
A required option was not specified on the command line.

User response:
Correct the command line.

6027-682 Device argument is missing.

Explanation:
The device argument was not specified on the command line.

User response:
Correct the command line.

6027-683 Disk *name* is invalid.

Explanation:
An incorrect disk name was specified on the command line.

User response:
Correct the command line.

6027-684 Value *value* for *option* is incorrect.

Explanation:
An incorrect value was specified for the named option.

User response:
Correct the command line.

6027-685 Value *value* for option *option* is out of range. Valid values are *number* through *number*.

Explanation:
An out of range value was specified for the named option.

User response:
Correct the command line.

6027-686 *option* (*value*) exceeds option (*value*).

Explanation:
The value of the first option exceeds the value of the second option. This is not permitted.

User response:
Correct the command line.

6027-687 Disk *name* is specified more than once.

Explanation:
The named disk was specified more than once on the command line.

User response:
Correct the command line.

6027-688 Failed to read file system descriptor.

Explanation:
The disk block containing critical information about the file system could not be read from disk.

User response:
This is usually an error in the path to the disks. If there are associated messages indicating an I/O error such as **ENODEV** or **EIO**, correct that error and retry the operation. If there are no associated I/O errors, then run the **mmfsck** command with the file system unmounted.

6027-689 Failed to update file system descriptor.

Explanation:
The disk block containing critical information about the file system could not be written to disk.

User response:
This is a serious error, which may leave the file system in an unusable state. Correct any I/O errors, then run the **mmfsck** command with the file system unmounted to make repairs.

6027-690 Failed to allocate I/O buffer.

Explanation:
Could not obtain enough memory (RAM) to perform an operation.

User response:
Either retry the operation when the **mmfsd** daemon is less heavily loaded, or increase the size of one or more of the memory pool parameters by issuing the **mmchconfig** command.

6027-691 Failed to send message to node *nodeName*.

Explanation:
A message to another file system node could not be sent.

User response:
Check additional error message and the internode communication configuration.

6027-692 Value for *option* is not valid. Valid values are yes, no.

Explanation:
An option that is required to be **yes** or **no** is neither.

User response:
Correct the command line.

6027-693 Cannot open disk *name*.

Explanation:
Could not access the given disk.

User response:
Check the disk hardware and the path to the disk.

6027-694 Disk not started; disk *name* has a bad volume label.

Explanation:
The volume label on the disk does not match that expected by GPFS.

User response:
Check the disk hardware. For hot-pluggable drives, ensure that the proper drive has been plugged in.

6027-695 File system is read-only.

Explanation:
An operation was attempted that would require modifying the contents of a file system, but the file system is read-only.

User response:
Make the file system R/W before retrying the operation.

6027-696 Too many disks are unavailable.

Explanation:
A file system operation failed because all replicas of a data or metadata block are currently unavailable.

User response:
Issue the **mmlsdisk** command to check the availability of the disks in the file system; correct disk hardware problems, and then issue the **mmchdisk** command with the **start** option to inform the file system that the disk or disks are available again.

6027-697 No log available.

Explanation:
A file system operation failed because no space for logging metadata changes could be found.

User response:
Check additional error message. A likely reason for this error is that all disks with available log space are currently unavailable.

6027-698 Not enough memory to allocate internal data structure.

Explanation:
A file system operation failed because no memory is available for allocating internal data structures.

User response:
Stop other processes that may have main memory pinned for their use.

6027-699 Inconsistency in file system metadata.

Explanation:
File system metadata on disk has been corrupted.

User response:
This is an extremely serious error that may cause loss of data. Issue the **mmfsck** command with the file system unmounted to make repairs. There will be a **POSSIBLE FILE CORRUPTION** entry in the system error log that should be forwarded to the IBM Support Center.

6027-700 Log recovery failed.

Explanation:
An error was encountered while restoring file system metadata from the log.

User response:
Check additional error message. A likely reason for this error is that none of the replicas of the log could be accessed because too many disks are currently unavailable. If the problem persists, issue the **mmfsck** command with the file system unmounted.

6027-701 Some file system data are inaccessible at this time.

Explanation:
The file system has encountered an error that is serious enough to make some or all data inaccessible. This message indicates that an occurred that left the file system in an unusable state.

User response:
Possible reasons include too many unavailable disks or insufficient memory for file system control structures. Check other error messages as well as the error log for additional information. Unmount the file system and correct any I/O errors. Then remount the file system and try the operation again. If the problem persists, issue the **mmfsck** command with the file system unmounted to make repairs.

6027-702 Some file system data are inaccessible at this time. Check error log for additional information. After correcting the problem, the file system must be unmounted and then mounted to restore normal data access.

Explanation:
The file system has encountered an error that is serious enough to make some or all data inaccessible. This message indicates that an error occurred that left the file system in an unusable state.

User response:
Possible reasons include too many unavailable disks or insufficient memory for file system control structures. Check other error messages as well as the error log for additional information. Unmount the file system and correct any I/O errors. Then remount the file system and try the operation again. If the problem persists, issue the **mmfsck** command with the file system unmounted to make repairs.

6027-703 Some file system data are inaccessible at this time. Check error log for additional information.

Explanation:
The file system has encountered an error that is serious enough to make some or all data inaccessible. This

message indicates that an error occurred that left the file system in an unusable state.

User response:
Possible reasons include too many unavailable disks or insufficient memory for file system control structures. Check other error messages as well as the error log for additional information. Unmount the file system and correct any I/O errors. Then remount the file system and try the operation again. If the problem persists, issue the **mmfsck** command with the file system unmounted to make repairs.

6027-704 Attention: Due to an earlier error normal access to this file system has been disabled. Check error log for additional information. After correcting the problem, the file system must be unmounted and then mounted again to restore normal data access.

Explanation:
The file system has encountered an error that is serious enough to make some or all data inaccessible. This message indicates that an error occurred that left the file system in an unusable state.

User response:
Possible reasons include too many unavailable disks or insufficient memory for file system control structures. Check other error messages as well as the error log for additional information. Unmount the file system and correct any I/O errors. Then remount the file system and try the operation again. If the problem persists, issue the **mmfsck** command with the file system unmounted to make repairs.

6027-705 Error code *value*.

Explanation:
Provides additional information about an error.

User response:
See accompanying error messages.

6027-706 The device *name* has no corresponding entry in *fileName* or has an incomplete entry.

Explanation:
The command requires a device that has a file system associated with it.

User response:
Check the operating system's file system database (the given file) for a valid device entry.

6027-707 Unable to open file *fileName*.

Explanation:
The named file cannot be opened.

User response:
Check that the file exists and has the correct permissions.

6027-708 Keyword *name* is incorrect. Valid values are *list*.

Explanation:
An incorrect keyword was encountered.

User response:
Correct the command line.

6027-709 Incorrect response. Valid responses are yes, no, or noall.

Explanation:
A question was asked that requires a **yes** or **no** answer. The answer entered was neither **yes**, **no**, nor **noall**.

User response:
Enter a valid response.

6027-710 Attention:

Explanation:
Precedes an attention messages.

User response:
None. Informational message only.

6027-711 Specified entity, such as a disk or file system, does not exist.

Explanation:
A file system operation failed because the specified entity, such as a disk or file system, could not be found.

User response:
Specify existing disk, file system, etc.

6027-712 Error in communications between mmfsd daemon and client program.

Explanation:
A message sent between the **mmfsd** daemon and the client program had an incorrect format or content.

User response:
Verify that the **mmfsd** daemon is running.

6027-713 Unable to start because conflicting program *name* is running. Waiting until it completes.

Explanation:
A program detected that it cannot start because a conflicting program is running. The program will automatically start once the conflicting program has ended, as long as there are no other conflicting programs running at that time.

User response:
None. Informational message only.

6027-714 Terminating because conflicting program *name* is running.

Explanation:
A program detected that it must terminate because a conflicting program is running.

User response:
Reissue the command once the conflicting program has ended.

6027-715 *command* is finished waiting. Starting execution now.

Explanation:
A program detected that it can now begin running because a conflicting program has ended.

User response:
None. Information message only.

6027-716 Some file system data or metadata has been lost.

Explanation:
Unable to access some piece of file system data that has been lost due to the deletion of disks beyond the replication factor.

User response:
If the function did not complete, try to mount the file system in **restricted** mode.

6027-717 Must execute mmfsck before mount.

Explanation:
An attempt has been made to mount a file system on which an incomplete **mmfsck** command was run.

User response:
Reissue the **mmfsck** command to the repair file system, then reissue the **mount** command.

6027-718 The mmfsd daemon is not ready to handle commands yet.

Explanation:
The **mmfsd** daemon is not accepting messages because it is restarting or stopping.

User response:
None. Informational message only.

6027-719 Device type not supported.

Explanation:
A disk being added to a file system with the **mmadddisk** or **mmcrfs** command is not a character mode special file, or has characteristics not recognized by GPFS.

User response:
Check the characteristics of the disk being added to the file system.

6027-720 Actual sector size does not match given sector size.

Explanation:
A disk being added to a file system with the **mmadddisk** or **mmcrfs** command has a physical sector size that differs from that given in the disk description list.

User response:
Check the physical sector size of the disk being added to the file system.

6027-721 Host *name* in *fileName* is not valid.

Explanation:
A host name or IP address that is not valid was found in a configuration file.

User response:
Check the configuration file specified in the error message.

6027-722 Attention: Due to an earlier error normal access to this file system has been disabled. Check error log for additional information. The file system must be mounted again to restore normal data access.

Explanation:
The file system has encountered an error that is serious enough to make some or all data inaccessible. This message indicates that an error occurred that left the file system in an unusable state. Possible reasons include too many unavailable disks or insufficient memory for file system control structures.

User response:
Check other error messages as well as the error log for additional information. Correct any I/O errors. Then, remount the file system and try the operation again. If the problem persists, issue the **mmfsck** command with the file system unmounted to make repairs.

6027-723 Attention: Due to an earlier error normal access to this file system has been disabled. Check error log for additional information. After correcting the problem, the file system must be mounted again to restore normal data access.

Explanation:
The file system has encountered an error that is serious enough to make some or all data inaccessible. This message indicates that an error occurred that left the file system in an unusable state. Possible reasons

include too many unavailable disks or insufficient memory for file system control structures.

User response:
Check other error messages as well as the error log for additional information. Correct any I/O errors. Then, remount the file system and try the operation again. If the problem persists, issue the **mmfsck** command with the file system unmounted to make repairs.

6027-724 Incompatible file system format.

Explanation:
An attempt was made to access a file system that was formatted with an older version of the product that is no longer compatible with the version currently running.

User response:
To change the file system format version to the current version, issue the **-V** option on the **mmchfs** command.

6027-725 The mmfsd daemon is not ready to handle commands yet. Waiting for quorum.

Explanation:
The GPFS **mmfsd** daemon is not accepting messages because it is waiting for quorum.

User response:
Determine why insufficient nodes have joined the group to achieve quorum and rectify the problem.

6027-726 Quota initialization/startup failed.

Explanation:
Quota manager initialization was unsuccessful. The file system manager finished without quotas. Subsequent client mount requests will fail.

User response:
Check the error log and correct I/O errors. It may be necessary to issue the **mmcheckquota** command with the file system unmounted.

6027-727 Specified driver type *type* does not match disk *name* driver type *type*.

Explanation:
The driver type specified on the **mmchdisk** command does not match the current driver type of the disk.

User response:
Verify the driver type and reissue the command.

6027-728 Specified sector size *value* does not match disk *name* sector size *value*.

Explanation:
The sector size specified on the **mmchdisk** command does not match the current sector size of the disk.

User response:
Verify the sector size and reissue the command.

6027-729 Attention: No changes for disk *name* were specified.

Explanation:
The disk descriptor in the **mmchdisk** command does not specify that any changes are to be made to the disk.

User response:
Check the disk descriptor to determine if changes are needed.

6027-730 *command* on *fileSystem*.

Explanation:
Quota was activated or deactivated as stated as a result of the **mmquotaon**, **mmquotaoff**, **mmdefquotaon**, or **mmdefquotaoff** commands.

User response:
None, informational only. This message is enabled with the **-v** option on the **mmquotaon**, **mmquotaoff**, **mmdefquotaon**, or **mmdefquotaoff** commands.

6027-731 Error *number* while performing *command* for *name* quota on *fileSystem*

Explanation:
An error occurred when switching quotas of a certain type on or off. If errors were returned for multiple file systems, only the error code is shown.

User response:
Check the error code shown by the message to determine the reason.

6027-732 Error while performing *command* on *fileSystem*.

Explanation:
An error occurred while performing the stated command when listing or reporting quotas.

User response:
None. Informational message only.

6027-733 Edit quota: Incorrect format!

Explanation:
The format of one or more edited quota limit entries was not correct.

User response:
Reissue the **mmedquota** command. Change only the values for the limits and follow the instructions given.

6027-734 Quota check for *fileSystem* ended prematurely.

Explanation:
The user interrupted and terminated the command.

User response:
If ending the command was not intended, reissue the **mmcheckquota** command.

6027-735 Error editing string from mmfsd.

Explanation:
An internal error occurred in the **mmfsd** when editing a string.

User response:
None. Informational message only.

6027-736 Attention: Due to an earlier error normal access to this file system has been disabled. Check error log for additional information. The file system must be unmounted and then mounted again to restore normal data access.

Explanation:
The file system has encountered an error that is serious enough to make some or all data inaccessible. This message indicates that an error occurred that left the file system in an unusable state. Possible reasons include too many unavailable disks or insufficient memory for file system control structures.

User response:
Check other error messages as well as the error log for additional information. Unmount the file system and correct any I/O errors. Then, remount the file system and try the operation again. If the problem persists, issue the **mmfsck** command with the file system unmounted to make repairs.

6027-737 Attention: No metadata disks remain.

Explanation:
The **mmchdisk** command has been issued, but no metadata disks remain.

User response:
None. Informational message only.

6027-738 Attention: No data disks remain.

Explanation:
The **mmchdisk** command has been issued, but no data disks remain.

User response:
None. Informational message only.

6027-739 Attention: Due to an earlier configuration change the file system is no longer properly balanced.

Explanation:
The **mmlsdisk** command found that the file system is not properly balanced.

User response:
Issue the **mmrestripefs -b** command at your convenience.

6027-740 Attention: Due to an earlier configuration change the file system is no longer properly replicated.

Explanation:
The **mmlsdisk** command found that the file system is not properly replicated.

User response:
Issue the **mmrestripefs -r** command at your convenience

6027-741 Attention: Due to an earlier configuration change the file system may contain data that is at risk of being lost.

Explanation:
The **mmlsdisk** command found that critical data resides on disks that are suspended or being deleted.

User response:
Issue the **mmrestripefs -m** command as soon as possible.

6027-742 Error occurred while executing a command for *fileSystem*.

Explanation:
A quota command encountered a problem on a file system. Processing continues with the next file system.

User response:
None. Informational message only.

6027-743 Initial disk state was updated successfully, but another error may have changed the state again.

Explanation:
The **mmchdisk** command encountered an error after the disk status or availability change was already recorded in the file system configuration. The most likely reason for this problem is that too many disks have become unavailable or are still unavailable after the disk state change.

User response:
Issue an **mmchdisk start** command when more disks are available.

6027-744 Unable to run *command* while the file system is mounted in restricted mode.

Explanation:
A command that can alter the data in a file system was issued while the file system was mounted in restricted mode.

User response:
Mount the file system in read-only or read-write mode or unmount the file system and then reissue the command.

6027-745 *fileSystem*: no *quotaType* quota management enabled.

Explanation:
A quota command of the cited type was issued for the cited file system when no quota management was enabled.

User response:
Enable quota management and reissue the command.

6027-746 Editing quota limits for this user or group not permitted.

Explanation:
The **root** user or **system** group was specified for quota limit editing in the **mmedquota** command.

User response:
Specify a valid user or group in the **mmedquota** command. Editing quota limits for the **root** user or **system** group is prohibited.

6027-747 Too many nodes in cluster (max *number*) or file system (max *number*).

Explanation:
The operation cannot succeed because too many nodes are involved.

User response:
Reduce the number of nodes to the applicable stated limit.

6027-748 *fileSystem*: no quota management enabled

Explanation:
A quota command was issued for the cited file system when no quota management was enabled.

User response:
Enable quota management and reissue the command.

6027-749 Pool size changed to *number* K = *number* M.

Explanation:
Pool size successfully changed.

User response:
None. Informational message only.

6027-750 The node address *ipAddress* is not defined in the node list.

Explanation:
An address does not exist in the GPFS configuration file.

User response:
Perform required configuration steps prior to starting GPFS on the node.

6027-751 Error code *value*.

Explanation:
Provides additional information about an error.

User response:
See accompanying error messages.

6027-752 Lost membership in cluster *clusterName*. Unmounting file systems.

Explanation:
This node has lost membership in the cluster. Either GPFS is no longer available on enough nodes to maintain quorum, or this node could not communicate with other members of the quorum. This could be caused by a communications failure between nodes, or multiple GPFS failures.

User response:
See associated error logs on the failed nodes for additional problem determination information.

6027-753 Could not run command *command*.

Explanation:
The GPFS daemon failed to run the specified command.

User response:
Verify correct installation.

6027-754 Error reading string for mmfsd.

Explanation:
GPFS could not properly read an input string.

User response:
Check that GPFS is properly installed.

6027-755 Waiting for challenge to be responded during disk election.

Explanation:
The node has challenged another node, which won the previous election and is waiting for the challenger to respond.

User response:
None. Informational message only.

6027-756 Configuration invalid or inconsistent between different nodes.

Explanation:
Self-explanatory.

User response:
Check cluster and file system configuration.

6027-757 *name* is not an excluded disk.

Explanation:
Some of the disks passed to the **mmfsctl include** command are not marked as **excluded** in the **mmsdrfs** file.

User response:
Verify the list of disks supplied to this command.

6027-758 Disk(s) not started; disk *name* has a bad volume label.

Explanation:
The volume label on the disk does not match that expected by GPFS.

User response:
Check the disk hardware. For hot-pluggable drives, make sure the proper drive has been plugged in.

6027-759 *fileSystem* is still in use.

Explanation:
The **mmfsctl include** command found that the named file system is still mounted, or another GPFS command is running against the file system.

User response:
Unmount the file system if it is mounted, or wait for GPFS commands in progress to terminate before retrying the command.

6027-762 No quota enabled file system found.

Explanation:
There is no quota-enabled file system in this cluster.

User response:
None. Informational message only.

6027-763 uidInvalidate: Incorrect option *option*.

Explanation:
An incorrect option passed to the **uidinvalidate** command.

User response:
Correct the command invocation.

6027-764 Error invalidating UID remapping cache for *domain*.

Explanation:
An incorrect domain name passed to the **uidinvalidate** command.

User response:
Correct the command invocation.

6027-765 Tick value hasn't changed for nearly *number* seconds.

Explanation:
Clock ticks incremented by AIX have not been incremented.

User response:
Check the error log for hardware or device driver problems that might cause timer interrupts to be lost.

6027-766 This node will be expelled from cluster *cluster* due to expel msg from *node*.

Explanation:
This node is being expelled from the cluster.

User response:
Check the network connection between this node and the node specified above.

6027-767 Request sent to *node* to expel *node* from cluster *cluster*.

Explanation:
This node sent an expel request to the cluster manager node to expel another node.

User response:
Check network connection between this node and the node specified above.

6027-768 Wrong number of operands for mmpmon command '*command*'.

Explanation:
The command read from the input file has the wrong number of operands.

User response:
Correct the command invocation and reissue the command.

6027-769 Malformed mmpmon command '*command*'.

Explanation:
The command read from the input file is malformed, perhaps with an unknown keyword.

User response:
Correct the command invocation and reissue the command.

6027-770 Error writing user.quota file.

Explanation:
An error occurred while writing the cited quota file.

User response:
Check the status and availability of the disks and reissue the command.

6027-771 Error writing group.quota file.

Explanation:
An error occurred while writing the cited quota file.

User response:
Check the status and availability of the disks and reissue the command.

6027-772 Error writing fileset.quota file.

Explanation:
An error occurred while writing the cited quota file.

User response:
Check the status and availability of the disks and reissue the command.

6027-774 *fileSystem*: **quota management is not enabled, or one or more quota clients are not available.**

Explanation:
An attempt was made to perform quotas commands without quota management enabled, or one or more quota clients failed during quota check.

User response:
Correct the cause of the problem, and then reissue the quota command.

6027-775 During mmcheckquota processing, *number* **node(s) failed. It is recommended that mmcheckquota be repeated.**

Explanation:
Nodes failed while an online quota check was running.

User response:
Reissue the quota check command.

6027-776 *fileSystem*: **There was not enough space for the report. Please repeat quota check!**

Explanation:
The **vflag** is set in the **tscheckquota** command, but either no space or not enough space could be allocated for the differences to be printed.

User response:
Correct the space problem and reissue the quota check.

6027-777 Recovering nodes: *nodeList*.

Explanation:
Recovery for one or more nodes has begun.

User response:
No response is needed if this message is followed by 'recovered nodes' entries specifying the nodes. If this message is not followed by such a message, determine why recovery did not complete.

6027-778 Recovering nodes in cluster *cluster*: *nodeList*.

Explanation:
Recovery for one or more nodes in the cited cluster has begun.

User response:
No response is need if this message is followed by a 'recovered nodes' entries specifying the nodes. If this message is not followed by such a message, determine why recovery did not complete.

6027-779 Incorrect fileset name *filesetName*.

Explanation:
The fileset name provided on the command line is incorrect.

User response:
Correct the fileset name and reissue the command.

6027-780 Incorrect path to fileset junction *junctionName*.

Explanation:
The path to the fileset junction is incorrect.

User response:
Correct the junction path and reissue the command.

6027-781 Storage pools have not been enabled for file system *fileSystem*.

Explanation:
The user invoked a command with a storage pool option (**-p** or **-P**) before storage pools were enabled.

User response:
Enable storage pools with the **mmchfs -V** command, or correct the command invocation and reissue the command.

6027-784 Device not ready.

Explanation:
A device is not ready for operation.

User response:
Check previous messages for further information.

6027-785 Cannot establish connection.

Explanation:
This node cannot establish a connection to another node.

User response:
Check previous messages for further information.

6027-786 Message failed because the destination node refused the connection.

Explanation:
This node sent a message to a node that refuses to establish a connection.

User response:
Check previous messages for further information.

6027-787 Security configuration data is inconsistent or unavailable.

Explanation:
There was an error configuring security on this node.

User response:
Check previous messages for further information.

6027-788 Failed to load or initialize security library.

Explanation:
There was an error loading or initializing the security library on this node.

User response:
Check previous messages for further information.

6027-789 Unable to read offsets *offset* **to** *offset* **for inode** *inode* **snap** *snap***, from disk** *diskName***, sector** *sector***.**

Explanation:
The **mmdeldisk -c** command found that the cited addresses on the cited disk represent data that is no longer readable.

User response:
Save this output for later use in cleaning up failing disks.

6027-790 Specified storage pool *poolName* **does not match disk** *diskName* **storage pool** *poolName***. Use mmdeldisk and mmadddisk to change a disk's storage pool.**

Explanation:
An attempt was made to change a disk's storage pool assignment using the **mmchdisk** command. This can only be done by deleting the disk from its current storage pool and then adding it to the new pool.

User response:
Delete the disk from its current storage pool and then add it to the new pool.

6027-792 Policies have not been enabled for file system *fileSystem***.**

Explanation:
The cited file system must be upgraded to use policies.

User response:
Upgrade the file system via the **mmchfs -V** command.

6027-793 No policy file was installed for file system *fileSystem***.**

Explanation:
No policy file was installed for this file system.

User response:
Install a policy file.

6027-794 Failed to read policy file for file system *fileSystem***.**

Explanation:
Failed to read the policy file for the requested file system.

User response:
Reinstall the policy file.

6027-795 Failed to open *fileName***:** *errorCode***.**

Explanation:
An incorrect file name was specified to **tschpolicy**.

User response:
Correct the command invocation and reissue the command.

6027-796 Failed to read *fileName***:** *errorCode***.**

Explanation:
An incorrect file name was specified to **tschpolicy**.

User response:
Correct the command invocation and reissue the command.

6027-797 Failed to stat *fileName***:** *errorCode***.**

Explanation:
An incorrect file name was specified to **tschpolicy**.

User response:
Correct the command invocation and reissue the command.

6027-798 Policy files are limited to *number* bytes.

Explanation:
A user-specified policy file exceeded the maximum-allowed length.

User response:
Install a smaller policy file.

6027-799 Policy *policyName* installed and broadcast to all nodes.

Explanation:
Self-explanatory.

User response:
None. Informational message only.

6027-850 Unable to issue this command from a non-root user.

Explanation:
tsiostat requires root privileges to run.

User response:
Get the system administrator to change the executable to set the UID to 0.

6027-851 Unable to process interrupt received.

Explanation:
An interrupt occurred that **tsiostat** cannot process.

User response:
Contact the IBM Support Center

6027-852 interval and count must be positive integers.

Explanation:
Incorrect values were supplied for **tsiostat** parameters.

User response:
Correct the command invocation and reissue the command.

6027-853 interval must be less than 1024.

Explanation:
An incorrect value was supplied for the interval parameter.

User response:
Correct the command invocation and reissue the command.

6027-854 count must be less than 1024.

Explanation:
An incorrect value was supplied for the count parameter.

User response:
Correct the command invocation and reissue the command.

6027-855 Unable to connect to server, mmfsd is not started.

Explanation:
The **tsiostat** command was issued but the file system is not started.

User response:
Contact your system administrator.

6027-856 No information to report.

Explanation:
The **tsiostat** command was issued but no file systems are mounted.

User response:
Contact your system administrator.

6027-857 Error retrieving values.

Explanation:
The **tsiostat** command was issued and an internal error occurred.

User response:
Contact the IBM Support Center.

6027-858 File system not mounted.

Explanation:
The requested file system is not mounted.

User response:
Mount the file system and reattempt the failing operation.

6027-859 Set DIRECTIO failed

Explanation:
The **tsfattr** call failed.

User response:
Check for additional error messages. Resolve the problems before reattempting the failing operation.

6027-860 -d is not appropriate for an NFSv4 ACL

Explanation:
Produced by the **mmgetacl** or **mmputacl** commands when the **-d** option was specified, but the object has an NFS Version 4 ACL (does not have a default).

User response:
None. Informational message only.

6027-862 Incorrect storage pool name *poolName*.

Explanation:
An incorrect storage pool name was provided.

User response:
Determine the correct storage pool name and reissue the command.

6027-863 File cannot be assigned to storage pool '*poolName*'.

Explanation:
The file cannot be assigned to the specified pool.

User response:
Determine the correct storage pool name and reissue the command.

6027-864 Set storage pool failed.

Explanation:
An incorrect storage pool name was provided.

User response:
Determine the correct storage pool name and reissue the command.

6027-865 Restripe file data failed.

Explanation:
An error occurred while restriping the file data.

User response:
Check the error code and reissue the command.

6027-866 Storage pools have not been enabled for this file system.

Explanation:
The user invoked a command with a storage pool option (**-p** or **-P**) before storage pools were enabled.

User response:
Enable storage pools via **mmchfs -V**, or correct the command invocation and reissue the command.

6027-867 Change storage pool is not permitted.

Explanation:
The user tried to change a file's assigned storage pool but was not root or superuser.

User response:
Reissue the command as root or superuser.

6027-868 mmchattr failed.

Explanation:
An error occurred while changing a file's attributes.

User response:
Check the error code and reissue the command.

6027-869 File replication exceeds number of failure groups in destination storage pool.

Explanation:
The **tschattr** command received incorrect command line arguments.

User response:
Correct the command invocation and reissue the command.

6027-870 Error on getcwd(): *errorString*. Try an absolute path instead of just *pathName*

Explanation:
The **getcwd** system call failed.

User response:
Specify an absolute path starting with '/' on the command invocation, so that the command will not need to invoke **getcwd**.

6027-871 Error on gpfs_get_pathname_from_fssnaphandle (*pathName*): *errorString*.

Explanation:
An error occurred during a **gpfs_get_pathname_from_fssnaphandle** operation.

User response:
Verify the invocation parameters and make sure the command is running under a user ID with sufficient authority (**root** or administrator privileges). Specify a GPFS file system device name or a GPFS directory path name as the first argument. Correct the command invocation and reissue the command.

6027-872 Is *pathName* a GPFS file system name or path?

Explanation:
An error occurred while attempting to access the named GPFS file system or path.

User response:
Verify the invocation parameters and make sure the command is running under a user ID with sufficient authority (**root** or administrator privileges). Correct the command invocation and reissue the command.

6027-874 Error: incorrect Date@Time (YYYY-MM-DD@HH:MM:SS) specification: *specification*.

Explanation:
The *Date@Time* command invocation argument could not be parsed.

User response:
Correct the command invocation and try again. The syntax should look similar to: **2005-12-25@07:30:00**.

6027-875 **Error on gpfs_stat(***pathName***:** *errorString***.**

Explanation:
An error occurred while attempting to **stat()** the cited path name.

User response:
Determine whether the cited path name exists and is accessible. Correct the command arguments as necessary and reissue the command.

6027-876 **Error starting directory scan(***pathName***:** *errorString***.**

Explanation:
The specified path name is not a directory.

User response:
Determine whether the specified path name exists and is an accessible directory. Correct the command arguments as necessary and reissue the command.

6027-877 **Error opening** *pathName errorString***.**

Explanation:
An error occurred while attempting to open the named file. Its pool and replication attributes remain unchanged.

User response:
Investigate the file and possibly reissue the command. The file may have been removed or locked by another application.

6027-878 **Error on gpfs_fcntl(***pathName***):** *errorString* **(offset=***offset***.**

Explanation:
An error occurred while attempting **fcntl** on the named file. Its pool or replication attributes may not have been adjusted.

User response:
Investigate the file and possibly reissue the command. Use the **mmlsattr** and **mmchattr** commands to examine and change the pool and replication attributes of the named file.

6027-879 **Error deleting** *pathName***:** *errorString***.**

Explanation:
An error occurred while attempting to delete the named file.

User response:
Investigate the file and possibly reissue the command. The file may have been removed or locked by another application.

6027-880 **Error on gpfs_seek_inode([***pathName*/*fileName***],** *inodeNumber***):** *errorString***.**

Explanation:
An error occurred during a **gpfs_seek_inode()** operation.

User response:
Reissue the command. If the problem persists, contact the IBM Support Center.

6027-881 **Error on gpfs_iopen([***rootPath*/*pathName***],** *inodeNumber***):** *errorString***.**

Explanation:
An error occurred during a **gpfs_iopen** operation.

User response:
Reissue the command. If the problem persists, contact the IBM Support Center.

6027-882 **Error on gpfs_ireaddir([***rootPath*/*pathName***],** *inodeNumber***):** *errorString***.**

Explanation:
An error occurred during a **gpfs_ireaddir()** operation.

User response:
Reissue the command. If the problem persists, contact the IBM Support Center.

6027-883 **Error on gpfs_stat_inode([***pathName*/*fileName***],** *inodeNumber***.***genNumber***):** *errorString***.**

Explanation:
An error occurred during a **gpfs_stat_inode()** operation.

User response:
Reissue the command. If the problem persists, contact the IBM Support Center.

6027-884 **Error during directory scan (***returnCode***).**

Explanation:
A terminal error occurred during the directory scan phase of the command.

User response:
Verify the command arguments. Reissue the command. If the problem persists, contact the IBM Support Center.

6027-885 **Error during inode scan:** *errorString***(***returnCode***).**

Explanation:
A terminal error occurred during the inode scan phase of the command.

User response:
Verify the command arguments. Reissue the command. If the problem persists, contact the IBM Support Center.

6027-886 **Error during policy decisions scan (*returnCode*).**

Explanation:
A terminal error occurred during the policy decisions phase of the command.

User response:
Verify the command arguments. Reissue the command. If the problem persists, contact the IBM Support Center.

6027-887 **Error on gpfs_igetstoragepool(*datapoolId*: *errorString*).**

Explanation:
An error occurred during a **gpfs_igetstoragepool** operation.

User response:
Reissue the command. If the problem persists, contact the IBM Support Center.

6027-888 **Error on gpfs_igetfilesetname(*filesetId*: *errorString*).**

Explanation:
An error occurred during a **gpfs_igetfilesetname** operation.

User response:
Reissue the command. If the problem persists, contact the IBM Support Center.

6027-889 **Error on gpfs_get_fssnaphandle(*rootPath*: *errorString*).**

Explanation:
An error occurred during a **gpfs_get_fssnaphandle** operation.

User response:
Reissue the command. If the problem persists, contact the IBM Support Center.

6027-890 **Error on gpfs_open_inodescan(*rootPath*: *errorString*.**

Explanation:
An error occurred during a **gpfs_open_inodescan()** operation.

User response:
Reissue the command. If the problem persists, contact the IBM Support Center.

6027-891 **WEIGHT(*thresholdValue*) UNKNOWN *pathName*.**

Explanation:
The named file was assigned the indicated weight, but the rule type is **UNKNOWN**.

User response:
Contact the IBM Support Center.

6027-892 **Error on pthread_create:** *where* **#***threadNumber_or_portNumber_or_socketNumber*: *errorString*.

Explanation:
An error occurred while creating the thread during a **pthread_create** operation.

User response:
Consider some of the command parameters that might affect memory usage. For further assistance, contact the IBM Support Center.

6027-893 **Error on pthread_mutex_init:** *errorString*.

Explanation:
An error occurred during a **pthread_mutex_init** operation.

User response:
Contact the IBM Support Center.

6027-894 **Error on pthread_mutex_lock:** *errorString*.

Explanation:
An error occurred during a **pthread_mutex_lock** operation.

User response:
Contact the IBM Support Center.

6027-895 **Error on pthread_mutex_unlock:** *errorString*.

Explanation:
An error occurred during a **pthread_mutex_unlock** operation.

User response:
Contact the IBM Support Center.

6027-896 **Error on pthread_cond_init:** *errorString*.

Explanation:
An error occurred during a **pthread_cond_init** operation.

User response:
Contact the IBM Support Center.

6027-897 **Error on pthread_cond_signal:** *errorString*.

Explanation:
An error occurred during a **pthread_cond_signal** operation.

User response:
Contact the IBM Support Center.

6027-898 **Error on pthread_cond_broadcast:** *errorString*.

Explanation:
An error occurred during a **pthread_cond_broadcast** operation.

User response:
Contact the IBM Support Center.

6027-899 **Error on pthread_cond_wait:** *errorString*.

Explanation:
An error occurred during a **pthread_cond_wait** operation.

User response:
Contact the IBM Support Center.

6027-900 **Error opening work file** *fileName*: *errorString*.

Explanation:
An error occurred while attempting to open the named work file.

User response:
Investigate the file and possibly reissue the command. Check that the path name is defined and accessible.

6027-901 **Error writing to work file** *fileName*: *errorString*.

Explanation:
An error occurred while attempting to write to the named work file.

User response:
Investigate the file and possibly reissue the command. Check that there is sufficient free space in the file system.

6027-902 **Error parsing work file** *fileName* **code** *number*.

Explanation:
An error occurred while attempting to read the named work file.

User response:
Investigate the file and possibly reissue the command. Check that there is sufficient free space in the file system. If the error persists, contact the IBM Support Center.

6027-903 **Error while loading policy rules: code** *number*.

Explanation:
An error occurred while attempting to read or parse the policy file, which may contain syntax errors. There should be subsequent messages with more information about the error.

User response:
Read all of the related error messages. Try to correct the problem.

6027-904 **Error** *returnCode* **from PD writer for inode=***inodeNumber* **pathname=***pathName*.

Explanation:
An error occurred while writing the policy decision for the candidate file with the indicated inode number and path name to a work file. There probably will be related error messages.

User response:
Read all the related error messages. Attempt to correct the problems.

6027-905 **Error: Out of memory, Code:** *number*.

Explanation:
The command has exhausted virtual memory.

User response:
Consider some of the command parameters that might affect memory usage. For further assistance, contact the IBM Support Center.

6027-906 **Error on system (***command***): rc=***number*.

Explanation:
An error occurred during the system call with the specified argument string.

User response:
Read and investigate related error messages.

6027-907 **Error code** *number* **from sort_file (***inodeListname*, *sortCommand*, *sortInodeOptions*, *tempDir*).

Explanation:
An error occurred while sorting the named work file using the named **sort** command with the given options and working directory.

User response:
Check these:

- The **sort** command is installed on your system.
- The **sort** command supports the given options.
- The working directory is accessible.
- The file system has sufficient free space.

6027-908 Attention: In RULE *ruleName* **(***ruleNumber***), the pool named by** *poolName* '*poolType*' **is not defined in the file system.**

Explanation:
The cited pool is not defined in the file system.

User response:
Correct the rule and reissue the command.

This is not an irrecoverable error; the command will continue to run. Of course it will not find any files in an incorrect **FROM POOL** and it will not be able to migrate any files to an incorrect **TO POOL**.

6027-909 Error on pthread_join: *where* **#***threadNumber***:** *errorString***.**

Explanation:
An error occurred while reaping the thread during a **pthread_join** operation.

User response:
Contact the IBM Support Center.

6027-910 Error during policy execution (*returnCode***).**

Explanation:
A terminating error occurred during the policy execution phase of the command.

User response:
Verify the command arguments and reissue the command. If the problem persists, contact the IBM Support Center.

6027-911 Error on *changeSpecification* **change for** *pathName***.** *errorString***.**

Explanation:
This provides more details on a **gpfs_fcntl()** error.

User response:
Use the **mmlsattr** and **mmchattr** commands to examine the file and then reissue the change command.

6027-912 Error on restriping of *pathName***.** *errorString***.**

Explanation:
This provides more details on a **gpfs_fcntl()** error.

User response:
Use the **mmlsattr** and **mmchattr** commands to examine the file and then reissue the restriping command.

6027-913 Desired replication exceeds number of failure groups.

Explanation:
While restriping a file, the **tschattr** or **tsrestripefile** command found that the desired replication exceeded the number of failure groups.

User response:
Reissue the command after adding or restarting file system disks.

6027-914 Insufficient space in one of the replica failure groups.

Explanation:
While restriping a file, the **tschattr** or **tsrestripefile** command found there was insufficient space in one of the replica failure groups.

User response:
Reissue the command after adding or restarting file system disks.

6027-915 Insufficient space to properly balance file.

Explanation:
While restriping a file, the **tschattr** or **tsrestripefile** command found that there was insufficient space to properly balance the file.

User response:
Reissue the command after adding or restarting file system disks.

6027-916 Too many disks unavailable to properly balance file.

Explanation:
While restriping a file, the **tschattr** or **tsrestripefile** command found that there were too many disks unavailable to properly balance the file.

User response:
Reissue the command after adding or restarting file system disks.

6027-917 All replicas of a data block were previously deleted.

Explanation:
While restriping a file, the **tschattr** or **tsrestripefile** command found that all replicas of a data block were previously deleted.

User response:
Reissue the command after adding or restarting file system disks.

6027-918 Cannot make this change to a nonzero length file.

Explanation:
GPFS does not support the requested change to the replication attributes.

User response:
You may want to create a new file with the desired attributes and then copy your data to that file and rename it appropriately. Be sure that there are sufficient disks assigned to the pool with different failure groups to support the desired replication attributes.

6027-919 Replication parameter range error (*value*, *value*).

Explanation:
Similar to message **6027-918**. The (a,b) numbers are the allowable range of the replication attributes.

User response:
You may want to create a new file with the desired attributes and then copy your data to that file and rename it appropriately. Be sure that there are sufficient disks assigned to the pool with different failure groups to support the desired replication attributes.

6027-920 Error on pthread_detach(self): *where*: *errorString*.

Explanation:
An error occurred during a **pthread_detach** operation.

User response:
Contact the IBM Support Center.

6027-921 Error on socket *socketName* (*hostName*): *errorString*.

Explanation:
An error occurred during a socket operation.

User response:
Verify any command arguments related to interprocessor communication and then reissue the command. If the problem persists, contact the IBM Support Center.

6027-922 Error in Mtconx - p_accepts should not be empty.

Explanation:
The program discovered an inconsistency or logic error within itself.

User response:
Contact the IBM Support Center.

6027-923 Error - command client is an incompatible version: *hostName protocolVersion*

Explanation:
While operating in master/client mode, the command discovered that the client is running an incompatible version.

User response:
Ensure the same version of the command software is installed on all nodes in the clusters and then reissue the command.

6027-924 Error - unrecognized client response from *hostName clientResponse*.

Explanation:
Similar to message **6027-923**, except this may be an internal logic error.

User response:
Ensure the latest, same version software is installed on all nodes in the clusters and then reissue the command. If the problem persists, contact the IBM Support Center.

6027-925 Directory cannot be assigned to storage pool '*poolName*'.

Explanation:
The file cannot be assigned to the specified pool.

User response:
Determine the correct storage pool name and reissue the command.

6027-926 Symbolic link cannot be assigned to storage pool '*poolName*'.

Explanation:
The file cannot be assigned to the specified pool.

User response:
Determine the correct storage pool name and reissue the command.

6027-927 System file cannot be assigned to storage pool '*poolName*'.

Explanation:
The file cannot be assigned to the specified pool.

User response:
Determine the correct storage pool name and reissue the command.

6027-928 Error: filesystem/device *fileSystem* has no snapshot with name *snapshotName*.

Explanation:
The specified file system does not have a snapshot with the specified snapshot name.

User response:
Use the **mmlssnapshot** command to list the snapshot names for the file system.

6027-929 **Attention: In RULE '*ruleName*' (*ruleNumber*), both pools '*poolName*' and '*poolType*' are EXTERNAL. This is not a supported migration.**

Explanation:
The command does not support migration between two **EXTERNAL** pools.

User response:
Correct the rule and reissue the command.

Note: This is not an unrecoverable error. The command will continue to run.

6027-930 **Attention: In RULE '*ruleName*' LIST name '*listName*' appears, but there is no corresponding EXTERNAL LIST '*listName*' EXEC ... OPTS ... rule to specify a program to process the matching files.**

Explanation:
There should be an **EXTERNAL LIST** rule for every list named by your **LIST** rules.

User response:
Add an "**EXTERNAL LIST** *listName* **EXEC** *scriptName* **OPTS** *opts*" rule.

Note: This is not an unrecoverable error. For execution with **-I defer**, file lists are generated and saved, so **EXTERNAL LIST** rules are not strictly necessary for correct execution.

6027-931 **Error - The policy evaluation phase did not complete.**

Explanation:
One or more errors prevented the policy evaluation phase from examining all of the files.

User response:
Consider other messages emitted by the command. Take appropriate action and then reissue the command.

6027-932 **Error - The policy execution phase did not complete.**

Explanation:
One or more errors prevented the policy execution phase from operating on each chosen file.

User response:
Consider other messages emitted by the command. Take appropriate action and then reissue the command.

6027-933 **EXEC '*wouldbeScriptPathname*' of EXTERNAL POOL or LIST '*PoolOrListName*' fails TEST with code *scriptReturnCode* on this node.**

Explanation:
Each EXEC defined in an **EXTERNAL POOL** or **LIST** rule is run in **TEST** mode on each node. Each invocation that fails with a nonzero return code is reported. Command execution is terminated on any node that fails any of these tests.

User response:
Correct the **EXTERNAL POOL** or **LIST** rule, the EXEC script, or do nothing because this is not necessarily an error. The administrator may suppress execution of the **mmapplypolicy** command on some nodes by deliberately having one or more EXECs return nonzero codes.

6027-934 **Attention: Specified snapshot: '*SnapshotName*' will be ignored because the path specified: '*PathName*' is not within that snapshot.**

Explanation:
The command line specified both a path name to be scanned and a snapshot name, but the snapshot name was not consistent with the path name.

User response:
If you wanted the entire snapshot, just specify the GPFS file system name or device name. If you wanted a directory within a snapshot, specify a path name within that snapshot (for example, **/gpfs/FileSystemName/.snapshots/SnapShotName/Directory**).

6027-935 **Attention: In RULE '*ruleName*' (*ruleNumber*) LIMIT or REPLICATE clauses are ignored; not supported for migration to EXTERNAL pool '*storagePoolName*'.**

Explanation:
GPFS does not support the **LIMIT** or **REPLICATE** clauses during migration to external pools.

User response:
Correct the policy rule to avoid this warning message.

6027-936 **Error - command master is an incompatible version.**

Explanation:
While operating in master/client mode, the command discovered that the master is running an incompatible version.

User response:
Upgrade the command software on all nodes and reissue the command.

6027-937 Error creating shared temporary subdirectory *subDirName*: *subDirPath*

Explanation:
The **mkdir** command failed on the named subdirectory path.

User response:
Specify an existing writable shared directory as the shared temporary directory argument to the policy command. The policy command will create a subdirectory within that.

6027-938 Error closing work file *fileName*: *errorString*

Explanation:
An error occurred while attempting to close the named work file or socket.

User response:
Record the above information. Contact the IBM Support Center.

6027-940 Open failed.

Explanation:
The **open()** system call was not successful.

User response:
Check additional error messages.

6027-941 Set replication failed.

Explanation:
The **open()** system call was not successful.

User response:
Check additional error messages.

6027-943 -M and -R are only valid for zero length files.

Explanation:
The **mmchattr** command received command line arguments that were not valid.

User response:
Correct command line and reissue the command.

6027-944 -m *value* **exceeds number of failure groups for metadata.**

Explanation:
The **mmchattr** command received command line arguments that were not valid.

User response:
Correct command line and reissue the command.

6027-945 -r *value* **exceeds number of failure groups for data.**

Explanation:
The **mmchattr** command received command line arguments that were not valid.

User response:
Correct command line and reissue the command.

6027-961 Cannot execute *command*.

Explanation:
The **mmeditacl** command cannot invoke the **mmgetacl** or **mmputacl** command.

User response:
Contact your system administrator.

6027-963 EDITOR environment variable not set

Explanation:
Self-explanatory.

User response:
Set the EDITOR environment variable and reissue the command.

6027-964 EDITOR environment variable must be an absolute path name

Explanation:
Self-explanatory.

User response:
Set the EDITOR environment variable correctly and reissue the command.

6027-965 Cannot create temporary file

Explanation:
Self-explanatory.

User response:
Contact your system administrator.

6027-966 Cannot access *fileName*

Explanation:
Self-explanatory.

User response:
Verify file permissions.

6027-967 Should the modified ACL be applied? yes or no

Explanation:
Self-explanatory.

User response:
Respond **yes** if you want to commit the changes, **no** otherwise.

6027-971 Cannot find *fileName*

Explanation:
Self-explanatory.

User response:
Verify the file name and permissions.

6027-972 *name* **is not a directory (-d not valid).**

Explanation:
Self-explanatory.

User response:
None, only directories are allowed to have default ACLs.

6027-973 Cannot allocate *number* **byte buffer for ACL.**

Explanation:
There was not enough available memory to process the request.

User response:
Contact your system administrator.

6027-974 Failure reading ACL (rc=*number***).**

Explanation:
An unexpected error was encountered by **mmgetacl** or **mmeditacl**.

User response:
Examine the return code, contact the IBM Support Center if necessary.

6027-976 Failure writing ACL (rc=*number***).**

Explanation:
An unexpected error encountered by **mmputacl** or **mmeditacl**.

User response:
Examine the return code, Contact the IBM Support Center if necessary.

6027-977 Authorization failure

Explanation:
An attempt was made to create or modify the ACL for a file that you do not own.

User response:
Only the owner of a file or the root user can create or change the access control list for a file.

6027-978 Incorrect, duplicate, or missing access control entry detected.

Explanation:
An access control entry in the ACL that was created had incorrect syntax, one of the required access control entries is missing, or the ACL contains duplicate access control entries.

User response:
Correct the problem and reissue the command.

6027-979 Incorrect ACL entry: *entry***.**

Explanation:
Self-explanatory.

User response:
Correct the problem and reissue the command.

6027-980 *name* **is not a valid user name.**

Explanation:
Self-explanatory.

User response:
Specify a valid user name and reissue the command.

6027-981 *name* **is not a valid group name.**

Explanation:
Self-explanatory.

User response:
Specify a valid group name and reissue the command.

6027-982 *name* **is not a valid ACL entry type.**

Explanation:
Specify a valid ACL entry type and reissue the command.

User response:
Correct the problem and reissue the command.

6027-983 *name* **is not a valid permission set.**

Explanation:
Specify a valid permission set and reissue the command.

User response:
Correct the problem and reissue the command.

6027-985 An error was encountered while deleting the ACL (rc=*value***).**

Explanation:
An unexpected error was encountered by **tsdelacl**.

User response:
Examine the return code and contact the IBM Support Center, if necessary.

6027-986 Cannot open *fileName*.

Explanation:
Self-explanatory.

User response:
Verify the file name and permissions.

6027-987 *name* is not a valid special name.

Explanation:
Produced by the **mmputacl** command when the NFS V4 'special' identifier is followed by an unknown special id string. *name* is one of the following: 'owner@', 'group@', 'everyone@'.

User response:
Specify a valid NFS V4 special name and reissue the command.

6027-988 *type* is not a valid NFS V4 type.

Explanation:
Produced by the **mmputacl** command when the type field in an ACL entry is not one of the supported NFS Version 4 type values. *type* is one of the following: 'allow' or 'deny'.

User response:
Specify a valid NFS V4 type and reissue the command.

6027-989 *name* is not a valid NFS V4 flag.

Explanation:
A flag specified in an ACL entry is not one of the supported values, or is not valid for the type of object (inherit flags are valid for directories only). Valid values are **FileInherit**, **DirInherit**, and **InheritOnly**.

User response:
Specify a valid NFS V4 option and reissue the command.

6027-990 Missing permissions (*value* found, *value* are required).

Explanation:
The permissions listed are less than the number required.

User response:
Add the missing permissions and reissue the command.

6027-991 Combining FileInherit and DirInherit makes the mask ambiguous.

Explanation:
Produced by the **mmputacl** command when WRITE/CREATE is specified without **MKDIR** (or the other way around), and both the **FILE_INHERIT** and **DIR_INHERIT** flags are specified.

User response:
Make separate **FileInherit** and **DirInherit** entries and reissue the command.

6027-992 Subdirectory *name* already exists. Unable to create snapshot.

Explanation:
tsbackup was unable to create a snapshot because the snapshot subdirectory already exists. This condition sometimes is caused by issuing a Tivoli restore operation without specifying a different subdirectory as the target of the restore.

User response:
Remove or rename the existing subdirectory and then retry the command.

6027-993 Keyword *aclType* is incorrect. Valid values are: 'posix', 'nfs4', 'native'.

Explanation:
One of the **mm*acl** commands specified an incorrect value with the **-k** option.

User response:
Correct the *aclType* value and reissue the command.

6027-994 ACL permissions cannot be denied to the file owner.

Explanation:
The **mmputacl** command found that the **READ_ACL**, **WRITE_ACL**, **READ_ATTR**, or **WRITE_ATTR** permissions are explicitly being denied to the file owner. This is not permitted, in order to prevent the file being left with an ACL that cannot be modified.

User response:
Do not select the **READ_ACL**, **WRITE_ACL**, **READ_ATTR**, or **WRITE_ATTR** permissions on **deny** ACL entries for the **OWNER**.

6027-995 This command will run on a remote node, *nodeName*.

Explanation:
The **mmputacl** command was invoked for a file that resides on a file system in a remote cluster, and UID remapping is enabled. To parse the user and group names from the ACL file correctly, the command will be run transparently on a node in the remote cluster.

User response:
None. Informational message only.

6027-996 Error (*returnCode*) reading policy text from: *fileName*.

Explanation:
An error occurred while attempting to open or read the specified policy file. The policy file may be missing or inaccessible.

User response:
Read all of the related error messages and try to correct the problem.

6027-997 **Attention: RULE '***ruleName***' attempts to redefine EXTERNAL POOL '***poolName***', ignored.**

Explanation:
Execution continues as if the specified rule was not present.

User response:
Correct or remove the policy rule.

6027-998 **Error in FLR/PDR serving for client** *clientHostNameAndPortNumber***: FLRs=***numberOfFileListRecords* **PDRs=***numberOfPolicyDecisionResponses* **pdrs=***numberOfPolicyDecisionResponseRecords*

Explanation:
A protocol error has been detected among cooperating **mmapplypolicy** processes.

User response:
Reissue the command. If the problem persists, contact the IBM Support Center.

6027-999 **Authentication failed:** *myNumericNetworkAddress* **with** *partnersNumericNetworkAddress* **(code=** *codeIndicatingProtocolStepSequence* **rc=***errnoStyleErrorCode***)**

Explanation:
Two processes at the specified network addresses failed to authenticate. The cooperating processes should be on the same network; they should not be separated by a firewall.

User response:
Correct the configuration and try the operation again. If the problem persists, contact the IBM Support Center.

6027-1004 **Incorrect [***nodelist***] format in file:** *nodeListLine*

Explanation:
A **[***nodelist***]** line in the input stream is not a comma-separated list of nodes.

User response:
Fix the format of the **[***nodelist***]** line in the **mmfs.cfg** input file. This is usually the *NodeFile* specified on the **mmchconfig** command.

If no user-specified **[***nodelist***]** lines are in error, contact the IBM Support Center.

If user-specified **[***nodelist***]** lines are in error, correct these lines.

6027-1005 **Common is not sole item on [] line** *number***.**

Explanation:
A **[***nodelist***]** line in the input stream contains common plus any other names.

User response:
Fix the format of the **[***nodelist***]** line in the **mmfs.cfg** input file. This is usually the *NodeFile* specified on the **mmchconfig** command.

If no user-specified **[***nodelist***]** lines are in error, contact the IBM Support Center.

If user-specified **[***nodelist***]** lines are in error, correct these lines.

6027-1006 **Incorrect custom [ ] line** *number***.**

Explanation:
A **[***nodelist***]** line in the input stream is not of the format: **[***nodelist***]**. This covers syntax errors not covered by messages 6027-1004 and 6027-1005.

User response:
Fix the format of the list of nodes in the **mmfs.cfg** input file. This is usually the *NodeFile* specified on the **mmchconfig** command.

If no user-specified lines are in error, contact the IBM Support Center.

If user-specified lines are in error, correct these lines.

6027-1007 **attribute found in common multiple times:** *attribute***.**

Explanation:
The attribute specified on the command line is in the main input stream multiple times. This is occasionally legal, such as with the trace attribute. These attributes, however, are not meant to be repaired by **mmfixcfg**.

User response:
Fix the configuration file (**mmfs.cfg** or **mmfscfg1** in the SDR). All attributes modified by GPFS configuration commands may appear only once in common sections of the configuration file.

6027-1008 **Attribute found in custom multiple times:** *attribute***.**

Explanation:
The attribute specified on the command line is in a custom section multiple times. This is occasionally legal. These attributes are not meant to be repaired by **mmfixcfg**.

User response:
Fix the configuration file (**mmfs.cfg** or **mmfscfg1** in the SDR). All attributes modified by GPFS configuration commands may appear only once in custom sections of the configuration file.

6027-1022 Missing mandatory arguments on command line.

Explanation:
Some, but not enough, arguments were specified to the **mmcrfsc** command.

User response:
Specify all arguments as per the usage statement that follows.

6027-1023 File system size arguments must be an integer: *value***.**

Explanation:
The first two arguments specified to the **mmcrfsc** command are not integers.

User response:
File system size is an internal argument. **mmcrfs** should never call the **mmcrfsc** command without a valid **fssize** argument. Contact the IBM Support Center.

6027-1028 Incorrect value for *name* **option.**

Explanation:
An incorrect argument was specified with an option requiring one of a limited number of legal values. (**-s** or any of the **yes | no** options).

User response:
Use one of the legal values for the indicated flag.

6027-1029 Incorrect characters in integer field for *option***.**

Explanation:
An incorrect character was specified with the indicated option.

User response:
Use a valid integer for the indicated option.

6027-1030 Value below minimum for *name* **option.**

Explanation:
The value specified with an option was below the minimum.

User response:
Use an integer in the legal range for the indicated option.

6027-1031 Value above maximum for option *name***.**

Explanation:
The value specified with an option was above the maximum.

User response:
Use an integer in the legal range for the indicated option.

6027-1032 Incorrect option *optionName***.**

Explanation:
An unknown option was specified.

User response:
Use only the options shown in the syntax.

6027-1033 Option *optionName* **specified twice.**

Explanation:
An option was specified more than once on the command line.

User response:
Use options only once.

6027-1034 Missing argument after *optionName* **option.**

Explanation:
An option was not followed by an argument.

User response:
All options need an argument. Specify one.

6027-1035 Option -*optionName* **is mandatory.**

Explanation:
A mandatory input option was not specified.

User response:
Specify all mandatory options.

6027-1036 Option expected at *string***.**

Explanation:
Something other than an expected option was encountered on the latter portion of the command line.

User response:
Follow the syntax shown. Options may not have multiple values. Extra arguments are not allowed.

6027-1038 IndirectSize must be <= BlockSize and must be a multiple of LogicalSectorSize (512).

Explanation:
The IndirectSize specified was not a multiple of 512 or the IndirectSize specified was larger than BlockSize.

User response:
Use valid values for IndirectSize and BlockSize.

6027-1039 InodeSize must be a multiple of LocalSectorSize (512).

Explanation:
The specified InodeSize was not a multiple of 512.

User response:
Use a valid value for InodeSize.

6027-1040 InodeSize must be less than or equal to Blocksize.

Explanation:
The specified InodeSize was not less than or equal to Blocksize.

User response:
Use a valid value for InodeSize.

6027-1042 DefaultMetadataReplicas must be less than or equal to MaxMetadataReplicas.

Explanation:
The specified DefaultMetadataReplicas was greater than MaxMetadataReplicas.

User response:
Specify a valid value for DefaultMetadataReplicas.

6027-1043 DefaultDataReplicas must be less than or equal MaxDataReplicas.

Explanation:
The specified DefaultDataReplicas was greater than MaxDataReplicas.

User response:
Specify a valid value for DefaultDataReplicas.

6027-1055 LogicalSectorSize must be a multiple of 512

Explanation:
The specified *LogicalSectorSize* was not a multiple of 512.

User response:
Specify a valid *LogicalSectorSize*.

6027-1056 *Blocksize* must be a multiple of *LogicalSectorSize* × 32

Explanation:
The specified *Blocksize* was not a multiple of *LogicalSectorSize* × 32.

User response:
Specify a valid value for *Blocksize*.

6027-1057 InodeSize must be less than or equal to Blocksize.

Explanation:
The specified *InodeSize* was not less than or equal to *Blocksize*.

User response:
Specify a valid value for *InodeSize*.

6027-1059 Mode must be M or S: *mode*

Explanation:
The first argument provided in the **mmcrfsc** command was not M or S.

User response:
The **mmcrfsc** command should not be called by a user. If any other command produces this error, contact the IBM Support Center.

6027-1060 BlockSize must be 16K, 64K, 256K, 512K, 1M, 2M or 4M.

Explanation:
The specified value for *BlockSize* was not valid.

User response:
Specify a valid *BlockSize* value.

6027-1084 The specified block size (*value*K) exceeds the maximum allowed block size currently in effect (*value*K). Either specify a smaller value for the -B parameter, or increase the maximum block size by issuing: mmchconfig maxblocksize=*value*K and restarting the GPFS daemon.

Explanation:
The specified value for block size was greater than the value of the **maxblocksize** configuration parameter.

User response:
Specify a valid value, or increase the value of the allowed block size by specifying a larger value on the **maxblocksize** parameter of the **mmchconfig** command.

6027-1113 Incorrect option: *option*.

Explanation:
The specified command option is not valid.

User response:
Specify a valid option and reissue the command.

6027-1114 Get adapter failed for IPA: *ipAddress*, **node** *nodeName* **while** *condition*.

Explanation:
A node being processed (from the command line or from another source) was recognized by the **host**

command and yet its IP address does not have a valid adapter in the SDR adapter table.

User response:
If the IP address is not valid, something is wrong with the indicated node. If the IP address is valid, and should have an adapter, check the integrity of the SDR. If the node name is **localhost**, the IP address will be valid, but will have no adapter. **localhost** should not be entered from the command line.

6027-1119 Obsolete option: *option***.**

Explanation:
A command received an option that is not valid any more.

User response:
Correct the command line and reissue the command.

6027-1120 Interrupt received: No changes made.

Explanation:
A GPFS administration command (**mm...**) received an interrupt before committing any changes.

User response:
None. Informational message only.

6027-1123 Disk name must be specified in disk descriptor.

Explanation:
The disk name positional parameter (the first field) in a disk descriptor was empty. The bad disk descriptor is displayed following this message.

User response:
Correct the input and rerun the command.

6027-1124 Disk usage must be dataOnly, metadataOnly, descOnly, or dataAndMetadata.

Explanation:
The disk usage positional parameter (field 4) in a disk descriptor has a value that is not valid. The incorrect disk descriptor is displayed following this message.

User response:
Correct the input and reissue the command.

6027-1130 User-specified failure group must be in the range from -1 to 4000.

Explanation:
A disk descriptor specified a failure group out of range. The bad disk descriptor is displayed following this message.

User response:
Correct the input and reissue the command.

6027-1132 Interrupt received: changes not propagated.

Explanation:
An interrupt was received after changes were committed but before the changes could be propagated to all the nodes.

User response:
All changes will eventually propagate as nodes recycle or other GPFS administration commands are issued. Changes can be activated now by manually restarting the GPFS daemons.

6027-1133 Interrupt received. Only -A, -Q, and -T were changed.

Explanation:
An interrupt was received in the **mmchfs** command. Any changes to the **-A**, **-Q**, and **-T** attributes were successfully completed. No other file system attributes were changed.

User response:
Reissue the command if you want to change additional attributes of the file system. The changes can be undone by running **mmchfs** with the original values for the affected attributes.

6027-1135 Restriping may not have finished.

Explanation:
An interrupt occurred during restriping.

User response:
Restart the restripe. Verify that the file system was not damaged by running the **mmfsck** command.

6027-1136 *option* **option specified twice.**

Explanation:
An option was specified multiple times on a command line.

User response:
Correct the error on the command line and reissue the command.

6027-1137 *option* **value must be yes or no.**

Explanation:
A yes or no option was used with something other than **yes** or **no**.

User response:
Correct the error on the command line and reissue the command.

6027-1138 Incorrect extra argument: *argument*

Explanation:
Non-option arguments followed the mandatory arguments.

User response:
Unlike most POSIX commands, the main arguments come first, followed by the optional arguments. Correct the error and reissue the command.

6027-1140 Incorrect integer for *option*: *number*.

Explanation:
An option requiring an integer argument was followed by something that cannot be parsed as an integer.

User response:
Specify an integer with the indicated option.

6027-1141 No disk descriptor file specified.

Explanation:
An **-F** flag was not followed by the path name of a disk descriptor file.

User response:
Specify a valid disk descriptor file.

6027-1142 File *fileName* **already exists.**

Explanation:
The specified file already exists.

User response:
Rename the file or specify a different file name and reissue the command.

6027-1143 Cannot open *fileName*.

Explanation:
A file could not be opened.

User response:
Verify that the specified file exists and that you have the proper authorizations.

6027-1144 Incompatible cluster types. You cannot move file systems that were created by GPFS cluster type *sourceCluster* **into GPFS cluster type** *targetCluster*.

Explanation:
The source and target cluster types are incompatible.

User response:
Contact the IBM Support Center for assistance.

6027-1145 *parameter* **must be greater than 0:** *value*

Explanation:
A negative value had been specified for the named parameter, which requires a positive value.

User response:
Correct the input and reissue the command.

6027-1147 Error converting *diskName* **into an NSD.**

Explanation:
Error encountered while converting a disk into an NSD.

User response:
Check the preceding messages for more information.

6027-1148 File system *fileSystem* **already exists in the cluster. Use mmchfs -W to assign a new device name for the existing file system.**

Explanation:
You are trying to import a file system into the cluster but there is already a file system with the same name in the cluster.

User response:
Remove or rename the file system with the conflicting name.

6027-1149 *fileSystem* **is defined to have mount point** *mountpoint*. **There is already such a mount point in the cluster. Use mmchfs -T to assign a new mount point to the existing file system.**

Explanation:
The cluster into which the file system is being imported already contains a file system with the same mount point as the mount point of the file system being imported.

User response:
Use the **-T** option of the **mmchfs** command to change the mount point of the file system that is already in the cluster and then rerun the **mmimportfs** command.

6027-1150 Error encountered while importing disk *diskName*.

Explanation:
The **mmimportfs** command encountered problems while processing the disk.

User response:
Check the preceding messages for more information.

6027-1151 Disk *diskName* already exists in the cluster.

Explanation:
You are trying to import a file system that has a disk with the same name as some disk from a file system that is already in the cluster.

User response:
Remove or replace the disk with the conflicting name.

6027-1153 At least one node in the cluster must be defined as a quorum node.

Explanation:
All nodes were explicitly designated or allowed to default to be nonquorum.

User response:
Specify which of the nodes should be considered quorum nodes and reissue the command.

6027-1154 Incorrect node *node* specified for command.

Explanation:
The user specified a node that is not valid.

User response:
Specify a valid node.

6027-1155 The NSD servers for the following disks from file system *fileSystem* were reset or not defined: *diskList*

Explanation:
Either the **mmimportfs** command encountered disks with no NSD servers, or was forced to reset the NSD server information for one or more disks.

User response:
After the **mmimportfs** command finishes, use the **mmchnsd** command to assign NSD server nodes to the disks as needed.

6027-1156 The NSD servers for the following free disks were reset or not defined: *diskList*

Explanation:
Either the **mmimportfs** command encountered disks with no NSD servers, or was forced to reset the NSD server information for one or more disks.

User response:
After the **mmimportfs** command finishes, use the **mmchnsd** command to assign NSD server nodes to the disks as needed.

6027-1157 Use the mmchnsd command to assign NSD servers as needed.

Explanation:
Either the **mmimportfs** command encountered disks with no NSD servers, or was forced to reset the NSD server information for one or more disks. See the preceding messages for detailed information.

User response:
After the **mmimportfs** command finishes, use the **mmchnsd** command to assign NSD server nodes to the disks as needed.

6027-1159 The following file systems were not imported: *fileSystemList*

Explanation:
The **mmimportfs** command was not able to import the specified file systems. See the preceding messages for error information.

User response:
Correct the problems and reissue the **mmimportfs** command.

6027-1160 The drive letters for the following file systems have been reset: *fileSystemList*.

Explanation:
The drive letters associated with the specified file systems are already in use by existing file systems and have been reset.

User response:
After the **mmimportfs** command finishes, use the **-t** option of the **mmchfs** command to assign new drive letters as needed.

6027-1161 Use the dash character (-) to separate multiple node designations.

Explanation:
A command detected an incorrect character used as a separator in a list of node designations.

User response:
Correct the command line and reissue the command.

6027-1162 Use the semicolon character (;) to separate the disk names.

Explanation:
A command detected an incorrect character used as a separator in a list of disk names.

User response:
Correct the command line and reissue the command.

6027-1163 GPFS is still active on *nodeName***.**

Explanation:
The GPFS daemon was discovered to be active on the specified node during an operation that requires the daemon to be stopped.

User response:
Stop the daemon on the specified node and rerun the command.

6027-1164 Use mmchfs -t to assign drive letters as needed.

Explanation:
The **mmimportfs** command was forced to reset the drive letters associated with one or more file systems. See the preceding messages for detailed information.

User response:
After the **mmimportfs** command finishes, use the **-t** option of the **mmchfs** command to assign new drive letters as needed.

6027-1165 The PR attributes for the following disks from file system *fileSystem* **were reset or not yet established:** *diskList*

Explanation:
The **mmimportfs** command disabled the Persistent Reserve attribute for one or more disks.

User response:
After the **mmimportfs** command finishes, use the **mmchconfig** command to enable Persistent Reserve in the cluster as needed.

6027-1166 The PR attributes for the following free disks were reset or not yet established: *diskList*

Explanation:
The **mmimportfs** command disabled the Persistent Reserve attribute for one or more disks.

User response:
After the **mmimportfs** command finishes, use the **mmchconfig** command to enable Persistent Reserve in the cluster as needed.

6027-1167 Use mmchconfig to enable Persistent Reserve in the cluster as needed.

Explanation:
The **mmimportfs** command disabled the Persistent Reserve attribute for one or more disks.

User response:
After the **mmimportfs** command finishes, use the **mmchconfig** command to enable Persistent Reserve in the cluster as needed.

6027-1178 *parameter* **must be from** *value* **to** *value***:** *valueSpecified*

Explanation:
A parameter value specified was out of range.

User response:
Keep the specified value within the range shown.

6027-1188 Duplicate disk specified: *disk*

Explanation:
A disk was specified more than once on the command line.

User response:
Specify each disk only once.

6027-1189 You cannot delete all the disks.

Explanation:
The number of disks to delete is greater than or equal to the number of disks in the file system.

User response:
Delete only some of the disks. If you want to delete them all, use the **mmdelfs** command.

6027-1197 *parameter* **must be greater than** *value***:** *value***.**

Explanation:
An incorrect value was specified for the named parameter.

User response:
Correct the input and reissue the command.

6027-1200 tscrfs failed. Cannot create *device*

Explanation:
tscrfs failed.

User response:
See the GPFS error message from **tscrfs**.

6027-1201 Disk *diskName* **does not belong to file system** *fileSystem***.**

Explanation:
The specified disk was not found to be part of the cited file system.

User response:
If the disk and file system were specified as part of a GPFS command, reissue the command with a disk that belongs to the specified file system.

6027-1203 Attention: File system *fileSystem* may have some disks that are in a non-ready state. Issue the command: mmcommon recoverfs *fileSystem*

Explanation:
The specified file system may have some disks that are in a non-ready state.

User response:
Run **mmcommon recoverfs** *fileSystem* to ensure that the GPFS configuration data for the file system is current, and then display the states of the disks in the file system using the **mmlsdisk** command.

If any disks are in a non-ready state, steps should be taken to bring these disks into the ready state, or to remove them from the file system. This can be done by mounting the file system, or by using the **mmchdisk** command for a mounted or unmounted file system. When maintenance is complete or the failure has been repaired, use the **mmchdisk** command with the **start** option. If the failure cannot be repaired without loss of data, you can use the **mmdeldisk** command to delete the disks.

6027-1204 *command* failed.

Explanation:
An internal command failed. This is usually a call to the GPFS daemon.

User response:
See the error message from the command that failed.

6027-1205 Failed to connect to remote cluster *clusterName*.

Explanation:
Attempt to establish a connection to the specified cluster was not successful. This can be caused by a number of reasons: GPFS is down on all of the contact nodes, the contact node list is obsolete, the owner of the remote cluster revoked authorization, and so forth.

User response:
If the error persists, contact the administrator of the remote cluster and verify that the contact node information is current and that the authorization key files are current as well.

6027-1206 File system *fileSystem* belongs to cluster *clusterName*. Command is not allowed for remote file systems.

Explanation:
The specified file system is not local to the cluster, but belongs to the cited remote cluster.

User response:
Choose a local file system, or issue the command on a node in the remote cluster.

6027-1207 There is already an existing file system using *value*.

Explanation:
The mount point or device name specified matches that of an existing file system. The device name and mount point must be unique within a GPFS cluster.

User response:
Choose an unused name or path.

6027-1208 File system *fileSystem* not found in cluster *clusterName*.

Explanation:
The specified file system does not belong to the cited remote cluster. The local information about the file system is not current. The file system may have been deleted, renamed, or moved to a different cluster.

User response:
Contact the administrator of the remote cluster that owns the file system and verify the accuracy of the local information. Use the **mmremotefs show** command to display the local information about the file system. Use the **mmremotefs update** command to make the necessary changes.

6027-1209 GPFS is down on this node.

Explanation:
GPFS is not running on this node.

User response:
Ensure that GPFS is running and reissue the command.

6027-1210 GPFS is not ready to handle commands yet.

Explanation:
GPFS is in the process of initializing or waiting for quorum to be reached.

User response:
Reissue the command.

6027-1211 *fileSystem* refers to file system *fileSystem* in cluster *clusterName*.

Explanation:
Informational message.

User response:
None.

6027-1212 File system *fileSystem* does not belong to cluster *clusterName*.

Explanation:
The specified file system refers to a file system that is remote to the cited cluster. Indirect remote file system access is not allowed.

User response:
Contact the administrator of the remote cluster that owns the file system and verify the accuracy of the local information. Use the **mmremotefs show** command to display the local information about the file system. Use the **mmremotefs update** command to make the necessary changes.

6027-1213 *command* **failed. Error code** *errorCode***.**

Explanation:
An internal command failed. This is usually a call to the GPFS daemon.

User response:
Examine the error code and other messages to determine the reason for the failure. Correct the problem and reissue the command.

6027-1214 **Unable to enable Persistent Reserve on the following disks:** *diskList*

Explanation:
The command was unable to set up all of the disks to use Persistent Reserve.

User response:
Examine the disks and the additional error information to determine if the disks should have supported Persistent Reserve. Correct the problem and reissue the command.

6027-1215 **Unable to reset the Persistent Reserve attributes on one or more disks on the following nodes:** *nodeList*

Explanation:
The command could not reset Persistent Reserve on at least one disk on the specified nodes.

User response:
Examine the additional error information to determine whether nodes were down or if there was a disk error. Correct the problems and reissue the command.

6027-1216 **File** *fileName* **contains additional error information.**

Explanation:
The command generated a file containing additional error information.

User response:
Examine the additional error information.

6027-1217 **A disk descriptor contains an incorrect separator character.**

Explanation:
A command detected an incorrect character used as a separator in a disk descriptor.

User response:
Correct the disk descriptor and reissue the command.

6027-1219 **NSD discovery on node** *nodeName* **failed with return code** *value***.**

Explanation:
The NSD discovery process on the specified node failed with the specified return code.

User response:
Determine why the node cannot access the specified NSDs. Correct the problem and reissue the command.

6027-1220 **Node** *nodeName* **cannot be used as an NSD server for Persistent Reserve disk** *diskName* **because it is not an AIX node.**

Explanation:
The node shown was specified as an NSD server for *diskName*, but the node does not support Persistent Reserve.

User response:
Specify a node that supports Persistent Reserve as an NSD server.

6027-1221 **The number of NSD servers exceeds the maximum (***value***) allowed.**

Explanation:
The number of NSD servers in the disk descriptor exceeds the maximum allowed.

User response:
Change the disk descriptor to specify no more NSD servers than the maximum allowed.

6027-1222 **Cannot assign a minor number for file system** *fileSystem* **(major number** *deviceMajorNumber***).**

Explanation:
The command was not able to allocate a minor number for the new file system.

User response:
Delete unneeded **/dev** entries for the specified major number and reissue the command.

6027-1223 *ipAddress* **cannot be used for NFS serving; it is used by the GPFS daemon.**

Explanation:
The IP address shown has been specified for use by the GPFS daemon. The same IP address cannot be used for NFS serving because it cannot be failed over.

User response:
Specify a different IP address for NFS use and reissue the command.

6027-1224 There is no file system with drive letter *driveLetter*.

Explanation:
No file system in the GPFS cluster has the specified drive letter.

User response:
Reissue the command with a valid file system.

6027-1225 Explicit drive letters are supported only in a Windows environment. Specify a mount point or allow the default settings to take effect.

Explanation:
An explicit drive letter was specified on the **mmmount** command but the target node does not run the Windows operating system.

User response:
Specify a mount point or allow the default settings for the file system to take effect.

6027-1226 Explicit mount points are not supported in a Windows environment. Specify a drive letter or allow the default settings to take effect.

Explanation:
An explicit mount point was specified on the **mmmount** command but the target node runs the Windows operating system.

User response:
Specify a drive letter or allow the default settings for the file system to take effect.

6027-1227 The main GPFS cluster configuration file is locked. Retrying...

Explanation:
Another GPFS administration command has locked the cluster configuration file. The current process will try to obtain the lock a few times before giving up.

User response:
None. Informational message only.

6027-1228 Lock creation successful.

Explanation:
The holder of the lock has released it and the current process was able to obtain it.

User response:
None. Informational message only. The command will now continue.

6027-1229 Timed out waiting for lock. Try again later.

Explanation:
Another GPFS administration command kept the main GPFS cluster configuration file locked for over a minute.

User response:
Try again later. If no other GPFS administration command is presently running, see "GPFS cluster configuration data files are locked" on page 42.

6027-1230 *diskName* is a tiebreaker disk and cannot be deleted.

Explanation:
A request was made to GPFS to delete a node quorum tiebreaker disk.

User response:
Specify a different disk for deletion.

6027-1231 GPFS detected more than eight quorum nodes while node quorum with tiebreaker disks is in use.

Explanation:
A GPFS command detected more than eight quorum nodes, but this is not allowed while node quorum with tiebreaker disks is in use.

User response:
Reduce the number of quorum nodes to a maximum of eight, or use the normal node quorum algorithm.

6027-1232 GPFS failed to initialize the tiebreaker disks.

Explanation:
A GPFS command unsuccessfully attempted to initialize the node quorum tiebreaker disks.

User response:
Examine prior messages to determine why GPFS was unable to initialize the tiebreaker disks and correct the problem. After that, reissue the command.

6027-1233 Incorrect keyword: *value*.

Explanation:
A command received a keyword that is not valid.

User response:
Correct the command line and reissue the command.

6027-1234 Adding node *node* to the cluster will exceed the quorum node limit.

Explanation:
An attempt to add the cited node to the cluster resulted in the quorum node limit being exceeded.

User response:

Change the command invocation to not exceed the node quorum limit, and reissue the command.

6027-1235 The *fileName* kernel extension does not exist.

Explanation:
The cited kernel extension does not exist.

User response:
Create the needed kernel extension by compiling a custom **mmfslinux** module for your kernel (see steps in **/usr/lpp/mmfs/src/README**), or copy the binaries from another node with the identical environment.

6027-1236 Unable to verify kernel/module configuration.

Explanation:
The **mmfslinux** kernel extension does not exist.

User response:
Create the needed kernel extension by compiling a custom **mmfslinux** module for your kernel (see steps in **/usr/lpp/mmfs/src/README**), or copy the binaries from another node with the identical environment.

6027-1237 The GPFS daemon is still running; use the mmshutdown command.

Explanation:
An attempt was made to unload the GPFS kernel extensions while the GPFS daemon was still running.

User response:
Use the **mmshutdown** command to shut down the daemon.

6027-1238 Module *fileName* is still in use. Unmount all GPFS file systems and issue the command: mmfsadm cleanup

Explanation:
An attempt was made to unload the cited module while it was still in use.

User response:
Unmount all GPFS file systems and issue the command **mmfsadm cleanup**. If this does not solve the problem, reboot the machine.

6027-1239 Error unloading module *moduleName*.

Explanation:
GPFS was unable to unload the cited module.

User response:
Unmount all GPFS file systems and issue the command **mmfsadm cleanup**. If this does not solve the problem, reboot the machine.

6027-1240 Module *fileName* is already loaded.

Explanation:
An attempt was made to load the cited module, but it was already loaded.

User response:
None. Informational message only.

6027-1241 *diskName* was not found in /proc/partitions.

Explanation:
The cited disk was not found in **/proc/partitions**.

User response:
Take steps to cause the disk to appear in **/proc/partitions**, and then reissue the command.

6027-1242 GPFS is waiting for *requiredCondition*.

Explanation:
GPFS is unable to come up immediately due to the stated required condition not being satisfied yet.

User response:
This is an informational message. As long as the required condition is not satisfied, this message will repeat every five minutes. You may want to stop the GPFS daemon after a while, if it will be a long time before the required condition will be met.

6027-1243 *command*: Processing user configuration file: *fileName*

Explanation:
Progress information for the **mmcrcluster** command.

User response:
None. Informational message only.

6027-1244 *configParameter* is set by the mmcrcluster processing. Line in error: *configLine*. The line will be ignored; processing continues.

Explanation:
The specified parameter is set by the **mmcrcluster** command and cannot be overridden by the user.

User response:
None. Informational message only.

6027-1245 *configParameter* must be set with the *command* command. Line in error: *configLine*. The line is ignored; processing continues.

Explanation:
The specified parameter has additional dependencies and cannot be specified prior to the completion of the **mmcrcluster** command.

User response:
After the cluster is created, use the specified command to establish the desired configuration parameter.

6027-1246 *configParameter* **is an obsolete parameter. Line in error:** *configLine***. The line is ignored; processing continues.**

Explanation:
The specified parameter is not used by GPFS anymore.

User response:
None. Informational message only.

6027-1247 *configParameter* **cannot appear in a node-override section. Line in error:** *configLine***. The line is ignored; processing continues.**

Explanation:
The specified parameter must have the same value across all nodes in the cluster.

User response:
None. Informational message only.

6027-1248 **Mount point can not be a relative path name:** *path*

Explanation:
The mount point does not begin with /.

User response:
Specify the absolute path name for the mount point.

6027-1249 *operand* **can not be a relative path name:** *path***.**

Explanation:
The specified path name does not begin with '/'.

User response:
Specify the absolute path name.

6027-1250 **Key file is not valid.**

Explanation:
While attempting to establish a connection to another node, GPFS detected that the format of the public key file is not valid.

User response:
Use the **mmremotecluster** command to specify the correct public key.

6027-1251 **Key file mismatch.**

Explanation:
While attempting to establish a connection to another node, GPFS detected that the public key file does not match the public key file of the cluster to which the file system belongs.

User response:
Use the **mmremotecluster** command to specify the correct public key.

6027-1252 **Node** *nodeName* **already belongs to the GPFS cluster.**

Explanation:
A GPFS command found that a node to be added to a GPFS cluster already belongs to the cluster.

User response:
Specify a node that does not already belong to the GPFS cluster.

6027-1253 **Incorrect value for** *option* **option.**

Explanation:
The provided value for the specified option is not valid.

User response:
Correct the error and reissue the command.

6027-1255 **There is nothing to commit. You must first run:** *command***.**

Explanation:
You are attempting to commit an SSL private key but such a key has not been generated yet.

User response:
Run the specified command to generate the public/private key pair.

6027-1256 **The current authentication files are already committed.**

Explanation:
You are attempting to commit public/private key files that were previously generated with the **mmauth** command. The files have already been committed.

User response:
None. Informational message.

6027-1257 **There are uncommitted authentication files. You must first run:** *command***.**

Explanation:
You are attempting to generate new public/private key files but previously generated files have not been committed yet.

User response:
Run the specified command to commit the current public/private key pair.

6027-1258 **You must establish a cipher list first. Run:** *command***.**

Explanation:
You are attempting to commit an SSL private key but a cipher list has not been established yet.

User response:
Run the specified command to specify a cipher list.

6027-1259 *command* **not found. Ensure the OpenSSL code is properly installed.**

Explanation:
The specified command was not found.

User response:
Ensure the OpenSSL code is properly installed and reissue the command.

6027-1260 **For the logical volume specification -l** *lvName* **to be valid** *lvName* **must be the only logical volume in the volume group. However, volume group** *vgName* **contains the following logical volumes:**

Explanation:
The specified logical volume is not the only logical volume in the volume group. The others are listed.

User response:
Check your logical volume and volume group names. Reissue the command.

6027-1261 *logicalVolume* **is not a valid logical volume.**

Explanation:
The specified logical volume is not a valid logical volume, with a corresponding volume group.

User response:
Reissue the command using a valid logical volume.

6027-1262 *vgName* **is not a valid volume group name.**

Explanation:
The specified volume group name is not correct.

User response:
Reissue the command using a valid volume group name.

6027-1263 **For the hdisk specification -h** *physicalDiskName* **to be valid,** *physicalDiskName* **must be the only disk in the volume group. However, volume group** *vgName* **contains the following disks:**

Explanation:
The specified volume group was found to contain the disks listed.

User response:
Check your configuration. Ensure that you are referring to the correct volume group and hdisks. Remove extraneous disks from the volume group and reissue the command.

6027-1264 *physicalDiskName* **is not a valid physical volume name.**

Explanation:
The specified physical disk name is not correct.

User response:
Reissue the command using a valid physical disk name.

6027-1265 *pvid* **is not a valid physical volume id.**

Explanation:
The specified physical volume ID is not correct.

User response:
Reissue the command using a valid physical volume ID.

6027-1268 **Missing arguments.**

Explanation:
A GPFS administration command received an insufficient number of arguments.

User response:
Correct the command line and reissue the command.

6027-1269 **The device name** *device* **starts with a slash, but not /dev/.**

Explanation:
The device name does not start with **/dev/**.

User response:
Correct the device name.

6027-1270 **The device name** *device* **contains a slash, but not as its first character.**

Explanation:
The specified device name contains a slash, but the first character is not a slash.

User response:
The device name must be an unqualified device name or an absolute device path name, for example: **fs0** or **/dev/fs0**.

6027-1271 **Unexpected error from** *command***. Return code:** *value***.**

Explanation:
A GPFS administration command (**mm...**) received an unexpected error code from an internally called command.

User response:
Perform problem determination. See “GPFS commands are unsuccessful” on page 56.

6027-1272 Unknown user name *userName*.

Explanation:
The specified value cannot be resolved to a valid user ID (UID).

User response:
Reissue the command with a valid user name.

6027-1273 Unknown group name *groupName*.

Explanation:
The specified value cannot be resolved to a valid group ID (GID).

User response:
Reissue the command with a valid group name.

6027-1274 The administrative adapter port name was not defined for nodes: *nodeList*. Correct the problems and run the mmchcluster -N command again.

Explanation:
For reasons cited in prior error messages, the adapter port names specified for the listed nodes are not known.

User response:
Examine the prior error messages. Correct the problems and issue the command again.

6027-1275 Daemon node adapter *Node* was not found on admin node *Node*.

Explanation:
An input node descriptor was found to be incorrect. The node adapter specified for GPFS daemon communications was not found to exist on the cited GPFS administrative node.

User response:
Correct the input node descriptor and reissue the command.

6027-1276 Command failed for disks: *diskList*.

Explanation:
A GPFS command was unable to complete successfully on the listed disks.

User response:
Correct the problems and reissue the command.

6027-1277 No contact nodes were provided for cluster *clusterName*.

Explanation:
A GPFS command found that no contact nodes have been specified for the cited cluster.

User response:
Use the **mmremotecluster** command to specify some contact nodes for the cited cluster.

6027-1278 None of the contact nodes in cluster *clusterName* can be reached.

Explanation:
A GPFS command was unable to reach any of the contact nodes for the cited cluster.

User response:
Determine why the contact nodes for the cited cluster cannot be reached and correct the problem, or use the **mmremotecluster** command to specify some additional contact nodes that can be reached.

6027-1287 Node *nodeName* returned ENODEV for disk *diskName*.

Explanation:
The specified node returned **ENODEV** for the specified disk.

User response:
Determine the cause of the **ENODEV** error for the specified disk and rectify it. The **ENODEV** may be due to disk fencing or the removal of a device that previously was present.

6027-1288 Remote cluster *clusterName* was not found.

Explanation:
A GPFS command found that the cited cluster has not yet been identified to GPFS as a remote cluster.

User response:
Specify a remote cluster known to GPFS, or use the **mmremotecluster** command to make the cited cluster known to GPFS.

6027-1289 Name *name* is not allowed. It contains the following invalid special character: *char*

Explanation:
The cited name is not allowed because it contains the cited invalid special character.

User response:
Specify a name that does not contain an invalid special character, and reissue the command.

6027-1290 GPFS configuration data for file system *fileSystem* may not be in agreement with the on-disk data for the file system. Issue the command: mmcommon recoverfs *fileSystem*

Explanation:
GPFS detected that the GPFS configuration database data for the specified file system may not be in agreement with the on-disk data for the file system. This may be caused by a GPFS disk command that did not complete normally.

User response:
Issue the specified command to bring the GPFS configuration database into agreement with the on-disk data.

6027-1291 Options *name* and *name* cannot be specified at the same time.

Explanation:
Incompatible options were specified on the command line.

User response:
Select one of the options and reissue the command.

6027-1292 *nodeList* cannot be used with attribute name.

Explanation:
The specified configuration attribute cannot be changed on only a subset of nodes. This attribute must be the same on all nodes in the cluster.

User response:
Certain attributes, such as **autoload**, may not be customized from node to node. Change the attribute for the entire cluster.

6027-1293 There are no remote file systems.

Explanation:
A value of **all** was specified for the remote file system operand of a GPFS command, but no remote file systems are defined.

User response:
None. There are no remote file systems on which to operate.

6027-1294 Remote file system *fileSystem* is not defined.

Explanation:
The specified file system was used for the remote file system operand of a GPFS command, but the file system is not known to GPFS.

User response:
Specify a remote file system known to GPFS.

6027-1295 The GPFS configuration information is incorrect or not available.

Explanation:
A problem has been encountered while verifying the configuration information and the execution environment.

User response:
Check the preceding messages for more information. Correct the problem and restart GPFS.

6027-1296 Device name cannot be 'all'.

Explanation:
A device name of **all** was specified on a GPFS command.

User response:
Reissue the command with a valid device name.

6027-1297 Each device specifies metadataOnly for disk usage. This file system could not store data.

Explanation:
All disk descriptors specify **metadataOnly** for disk usage.

User response:
Change at least one disk descriptor in the file system to indicate the usage of **dataOnly** or **dataAndMetadata**.

6027-1298 Each device specifies dataOnly for disk usage. This file system could not store metadata.

Explanation:
All disk descriptors specify **dataOnly** for disk usage.

User response:
Change at least one disk descriptor in the file system to indicate a usage of **metadataOnly** or **dataAndMetadata**.

6027-1299 Incorrect integer *number* specified for failure group.

Explanation:
A disk descriptor's specified failure group is not a valid integer.

User response:
Change the disk descriptor to indicate a valid failure group number.

6027-1300 No file systems were found.

Explanation:
A GPFS command searched for file systems, but none were found.

User response:
Create a GPFS file system before reissuing the command.

6027-1301 The NSD servers specified in the disk descriptor do not match the NSD servers currently in effect.

Explanation:
The set of NSD servers specified in the disk descriptor does not match the set that is currently in effect.

User response:

Specify the same set of NSD servers in the disk descriptor as is currently in effect or omit it from the disk descriptor and then reissue the command. Use the **mmchnsd** command to change the NSD servers as needed.

6027-1302 *clusterName* **is the name of the local cluster.**

Explanation:
The cited cluster name was specified as the name of a remote cluster, but it is already being used as the name of the local cluster.

User response:
Use the **mmchcluster** command to change the name of the local cluster, and then reissue the command that failed.

6027-1303 **The concurrent virtual shared disk flag cannot be applied to previously-created virtual shared disk** *diskName***. The flag will be ignored.**

Explanation:
The concurrent virtual shared disk flag was specified, but it cannot be applied to the cited virtual shared disk since it has already been created. The flag will be ignored.

User response:
None. Informational message only.

6027-1304 **Missing argument after** *option* **option.**

Explanation:
The specified command option requires a value.

User response:
Specify a value and reissue the command.

6027-1305 **Cannot execute (**_command_**): return code** *value*

Explanation:
A command was not successfully invoked.

User response:
Determine why the command is not accessible.

6027-1306 **Command** *command* **failed with return code** *value***.**

Explanation:
A command was not successfully processed.

User response:
Correct the failure specified by the command and reissue the command.

6027-1307 **Disk** *disk* **on node** *nodeName* **already has a volume group** *vgName* **that does not appear to have been created by this program in a prior invocation. Correct the descriptor file or remove the volume group and retry.**

Explanation:
The specified disk already belongs to a volume group.

User response:
Either remove the volume group or remove the disk descriptor and retry.

6027-1308 **Disk** *disk* **on node** *nodeName* **already has a logical volume** *vgName* **that doesn't appear to have been created by this program in a prior invocation. Correct the descriptor file or remove the logical volume and retry.**

Explanation:
The specified disk already has a logical volume.

User response:
Either remove the logical volume or remove the disk descriptor and retry.

6027-1309 **Disk** *disk* **on node** *nodeName* **already has multiple logical volumes that don't appear to have been created by this program in a prior invocation. Correct the descriptor file or remove the logical volumes and retry.**

Explanation:
The specified disk already had multiple logical volumes.

User response:
Either remove the logical volumes or remove the disk descriptor and retry.

6027-1311 **The global volume group** *vgName* **we're attempting to define for node** *nodeName* **and disk** *disk* **is already defined with different server nodes or volume group. Correct the descriptor file or remove the offending global volume group and retry.**

Explanation:
The global volume group was defined with different parameters.

User response:
Either remove the global volume group or remove the disk descriptor and retry.

6027-1312 **The virtual shared disk** *diskName* **we're attempting to define for global volume group** *vgName*, **node** *nodeName*, **and logical volume** *lvName* **is already defined with different parameters. Correct the descriptor file or remove the offending virtual shared disk and retry.**

Explanation:
The virtual shared disk was already defined.

User response:
Either remove the virtual shared disk or remove the disk descriptor and retry.

6027-1323 **A prior invocation of this command has recorded a partial completion in the file** *name*. **Should we restart at prior failing step (***number***)? [y/n]**

Explanation:
The **mmcrvsd** command was restarted on a prior failing descriptor file.

User response:
Either allow the command to continue at the failing step or remove the failing step specification from the descriptor file.

6027-1324 **Unable to rewrite new descriptor to the file** *fileName*. **Ensure the file system has enough space and retry.**

Explanation:
There was either not enough space or the file system was write protected.

User response:
Either expand the file system or remove the write protection.

6027-1325 **A descriptor was encountered whose volume (***volumeName***) is not a previously-defined virtual shared disk and for which no primary server was specified.**

Explanation:
A disk descriptor specified a volume that was not a previously-defined virtual shared disk and for which no primary server was specified.

User response:
Either specify an existing virtual shared disk as the volume, or specify a primary server within the offending disk descriptor.

6027-1327 **Cannot open (***fileName***): return code** *value*.

Explanation:
The file could not be opened.

User response:
Either correct the permission bits or specify the correct file name.

6027-1328 **Cannot determine node number for host** *nodeName*. **Either the node is not known to the IBM Virtual Shared Disk subsystem, or it is not a GPFS administration node name.**

Explanation:
A lookup in the cluster configuration data could not determine the node number. This can be caused by specifying a node interface that is not known to the IBM Virtual Shared Disk subsystem, or by specifying a node interface other than a GPFS administration node name.

User response:
Correct the node or host definition in the SDR.

6027-1330 *diskName* **is already in volume group** *vgName* **and cannot be added to** *vgName*.

Explanation:
The specified disk already belongs to a volume group.

User response:
Either remove the volume group or specify a different disk.

6027-1332 **Cannot find** *disk* **with** *command*.

Explanation:
The specified disk cannot be found.

User response:
Specify a correct disk name.

6027-1333 **The following nodes could not be restored:** *nodeList*. **Correct the problems and use the mmsdrrestore command to recover these nodes.**

Explanation:
The **mmsdrrestore** command was unable to restore the configuration information for the listed nodes.

User response:
Correct the problems and reissue the **mmsdrrestore** command for these nodes.

6027-1334 Incorrect value for option *option*. Valid values are: *validValues*.

Explanation:
An incorrect argument was specified with an option requiring one of a limited number of legal options.

User response:
Use one of the legal values for the indicated option.

6027-1335 Command completed: Not all required changes were made.

Explanation:
Some, but not all, of the required changes were made.

User response:
Examine the preceding messages, correct the problems, and reissue the command.

6027-1336 Volume group *vgName* cannot be imported on node *nodeName* because the disk with physical volume ID *value* cannot be found.

Explanation:
The specified volume group cannot be imported because a disk with the specified pvid cannot be found. This problem may be caused by another node having a reserve on the disk, thus preventing access by the node trying to import the disk.

User response:
Ensure the disk with the specified physical volume ID is known on the specified node. If a node in the cluster has a reserve on the disk, release the reserve. When the disk with the specified physical volume ID is known on the specified node, reissue the command.

6027-1337 Failed to obtain DCE credentials; dsrvtgt *name* command rc= *value*. Continuing.

Explanation:
An attempt to obtain **spbgroot** DCE credentials has failed. Processing continues, but there may be a authentication failure later on.

User response:
Go to the *Parallel System Support Programs for AIX: Diagnosis Guide* and search on *diagnosing per node key management (PNKM) problems*. Follow the problem determination and repair actions specified.

6027-1338 Command is not allowed for remote file systems.

Explanation:
A command for which a remote file system is not allowed was issued against a remote file system.

User response:
Choose a local file system, or issue the command on a node in the cluster that owns the file system.

6027-1339 Disk usage *value* is incompatible with storage pool *name*.

Explanation:
A disk descriptor specified a disk usage involving metadata and a storage pool other than **system**.

User response:
Change the descriptor's disk usage field to **dataOnly**, or do not specify a storage pool name.

6027-1340 File *fileName* not found. Recover the file or run mmauth genkey.

Explanation:
The cited file was not found.

User response:
Recover the file or run the **mmauth genkey** command to recreate it.

6027-1341 Starting force unmount of GPFS file systems.

Explanation:
Progress information for the **mmshutdown** command.

User response:
None. Informational message only.

6027-1342 Unmount not finished after *value* seconds. Waiting *value* more seconds.

Explanation:
Progress information for the **mmshutdown** command.

User response:
None. Informational message only.

6027-1343 Unmount not finished after *value* seconds.

Explanation:
Progress information for the **mmshutdown** command.

User response:
None. Informational message only.

6027-1344 Shutting down GPFS daemons.

Explanation:
Progress information for the **mmshutdown** command.

User response:
None. Informational message only.

6027-1345 Finished.

Explanation:
Progress information for the **mmshutdown** command.

User response:
None. Informational message only.

6027-1347 Disk with NSD volume id *NSD volume id* **no longer exists in the GPFS cluster configuration data but the NSD volume id was not erased from the disk. To remove the NSD volume id, issue: mmdelnsd -p** *NSD volume id*

Explanation:
A GPFS administration command (**mm...**) successfully removed the disk with the specified NSD volume id from the GPFS cluster configuration data but was unable to erase the NSD volume id from the disk.

User response:
Issue the specified command to remove the NSD volume id from the disk.

6027-1348 Disk with NSD volume id *NSD volume id* **no longer exists in the GPFS cluster configuration data but the NSD volume id was not erased from the disk. To remove the NSD volume id, issue: mmdelnsd -p** *NSD volume id* **-N** *nodeNameList*

Explanation:
A GPFS administration command (**mm...**) successfully removed the disk with the specified NSD volume id from the GPFS cluster configuration data but was unable to erase the NSD volume id from the disk.

User response:
Issue the specified command to remove the NSD volume id from the disk.

6027-1351 Could not load switch partition table.

Explanation:
Failed to load the switch partition table.

User response:
Contact the IBM Support Center.

6027-1352 *fileSystem* **is not a remote file system known to GPFS.**

Explanation:
The cited file system is not the name of a remote file system known to GPFS.

User response:
Use the **mmremotefs** command to identify the cited file system to GPFS as a remote file system, and then reissue the command that failed.

6027-1354 A prior invocation of this command has recorded a partial completion in the file (*fileName***). We will restart at prior failing step (***value***).**

Explanation:
The **mmcrvsd** command was restarted on a prior failing descriptor file.

User response:
None. Informational message only.

6027-1355 The nodes (*nodeName***) and (***nodeName***) specified for disk (***disk***) must be defined within the same concurrent IBM Virtual Shared Disk cluster.**

Explanation:
An attempt had been made to create concurrent virtual shared disks and the specified nodes are either not defined in the concurrent IBM Virtual Shared Disk cluster, or the specified nodes are not in the same cluster.

User response:
Specify nodes in the same concurrent IBM Virtual Shared Disk cluster using the **vsdnode** command.

6027-1356 The same primary and backup server node (*nodeName***) cannot be specified for disk (***disk***).**

Explanation:
The primary and backup server nodes specified for an virtual shared disk are the same node.

User response:
Specify different nodes for the primary and backup servers.

6027-1357 An internode connection between GPFS nodes was disrupted.

Explanation:
An internode connection between GPFS nodes was disrupted, preventing its successful completion.

User response:
Reissue the command. If the problem recurs, determine and resolve the cause of the disruption. If the problem persists, contact the IBM Support Center.

6027-1358 No clusters are authorized to access this cluster.

Explanation:
Self-explanatory.

User response:
This is an informational message.

6027-1359 Cluster *clusterName* is not authorized to access this cluster.

Explanation:
Self-explanatory.

User response:
This is an informational message.

6027-1361 Attention: There are no available valid VFS type values for mmfs in /etc/vfs.

Explanation:
An out of range number was used as the vfs number for GPFS.

User response:
The valid range is 8 through 32. Check **/etc/vfs** and remove unneeded entries.

6027-1362 There are no remote cluster definitions.

Explanation:
A value of **all** was specified for the remote cluster operand of a GPFS command, but no remote clusters are defined.

User response:
None. There are no remote clusters on which to operate.

6027-1363 Remote cluster *clusterName* is not defined.

Explanation:
The specified cluster was specified for the remote cluster operand of a GPFS command, but the cluster is not known to GPFS.

User response:
Specify a remote cluster known to GPFS.

6027-1364 No disks specified

Explanation:
There were no disks in the descriptor list or file.

User response:
Specify at least one disk.

6027-1365 Disk *diskName* already belongs to file system *fileSystem*.

Explanation:
The specified disk name is already assigned to a GPFS file system. This may be because the disk was specified more than once as input to the command, or because the disk was assigned to a GPFS file system in the past.

User response:
Specify the disk only once as input to the command, or specify a disk that does not belong to a file system.

6027-1366 File system *fileSystem* has some disks that are in a non-ready state.

Explanation:
The specified file system has some disks that are in a non-ready state.

User response:
Run **mmcommon recoverfs** *fileSystem* to ensure that the GPFS configuration data for the file system is current. If some disks are still in a non-ready state, display the states of the disks in the file system using the **mmlsdisk** command. Any disks in an undesired non-ready state should be brought into the ready state by using the **mmchdisk** command or by mounting the file system. If these steps do not bring the disks into the ready state, use the **mmdeldisk** command to delete the disks from the file system.

6027-1367 Attention: Not all disks were marked as available.

Explanation:
The process of marking the disks as available could not be completed.

User response:
Before adding these disks to a GPFS file system, you should either reformat them, or use the **-v no** option on the **mmcrfs** or **mmadddisk** command.

6027-1368 This GPFS cluster contains declarations for remote file systems and clusters. You cannot delete the last node.

Explanation:
An attempt has been made to delete a GPFS cluster that still has declarations for remote file systems and clusters.

User response:
Before deleting the last node of a GPFS cluster, delete all remote cluster and file system information. Use the **delete** option of the **mmremotecluster** and **mmremotefs** commands.

6027-1369 Virtual shared disk *diskName* that was specified as input does not have a value defined for global volume group or logical volume name. The defined global volume group is *vgName* and the defined logical volume name is *lvName*.

Explanation:
The specified input virtual shared disk does not have values defined for its global volume group or its logical volume name.

User response:
Recreate the offending virtual shared disk and try again.

6027-1370 The following nodes could not be reached:

Explanation:
A GPFS command was unable to communicate with one or more nodes in the cluster. A list of the nodes that could not be reached follows.

User response:
Determine why the reported nodes could not be reached and resolve the problem.

6027-1371 Propagating the cluster configuration data to all affected nodes. This is an asynchronous process.

Explanation:
A process is initiated to distribute the cluster configuration data to other nodes in the cluster.

User response:
This is an informational message. The command does not wait for the distribution to finish.

6027-1373 There is no file system information in input file *fileName*.

Explanation:
The cited input file passed to the **mmimportfs** command contains no file system information. No file system can be imported.

User response:
Reissue the **mmimportfs** command while specifying a valid input file.

6027-1374 File system *fileSystem* was not found in input file *fileName*.

Explanation:
The specified file system was not found in the input file passed to the **mmimportfs** command. The file system cannot be imported.

User response:
Reissue the **mmimportfs** command while specifying a file system that exists in the input file.

6027-1375 The following file systems were not imported: *fileSystem*.

Explanation:
The **mmimportfs** command was unable to import one or more of the file systems in the input file. A list of the file systems that could not be imported follows.

User response:
Examine the preceding messages, rectify the problems that prevented the importation of the file systems, and reissue the **mmimportfs** command.

6027-1376 Disk *nsdName* with pvid *pvid* already exists in the cluster configuration file. If this NSD is built on a virtual shared disk that no longer exists, remove it from the cluster configuration file by issuing: mmdelnsd *nsdName*.

Explanation:
The command could not create the desired disk because there is already a disk with the same **pvid** recorded in the cluster configuration file **/var/mmfs/gen/mmsdrfs**.

User response:
If *nsdName* is built on a virtual shared disk that no longer exists, issue the specified command to remove the NSD from the cluster configuration file.

6027-1377 Attention: Unknown attribute specified: *name*. Press the ENTER key to continue.

Explanation:
The **mmchconfig** command received an unknown attribute.

User response:
Unless directed otherwise by the IBM Support Center, press any key to bypass this attribute.

6027-1378 Incorrect record found in the mmsdrfs file (code *value*):

Explanation:
A line that is not valid was detected in the main GPFS cluster configuration file **/var/mmfs/gen/mmsdrfs**.

User response:
The data in the cluster configuration file is incorrect. If no user modifications have been made to this file, contact the IBM Support Center. If user modifications have been made, correct these modifications.

6027-1379 There is no file system with mount point *mountpoint*.

Explanation:
No file system in the GPFS cluster has the specified mount point.

User response:
Reissue the command with a valid file system.

6027-1380 File system *fileSystem* is already mounted at *mountpoint*.

Explanation:
The specified file system is mounted at a mount point different than the one requested on the **mmmount** command.

User response:
Unmount the file system and reissue the command.

6027-1381 Mount point cannot be specified when mounting all file systems.

Explanation:
A device name of **all** and a mount point were specified on the **mmmount** command.

User response:
Reissue the command with a device name for a single file system or do not specify a mount point.

6027-1382 This node does not belong to a GPFS cluster.

Explanation:
The specified node does not appear to belong to a GPFS cluster, or the GPFS configuration information on the node has been lost.

User response:
Informational message. If you suspect that there is corruption of the GPFS configuration information, recover the data following the procedures outlined in "Recovery from loss of GPFS cluster configuration data file" on page 43.

6027-1383 There is no record for this node in file *fileName*. Either the node is not part of the cluster, the file is for a different cluster, or not all of the node's adapter interfaces have been activated yet.

Explanation:
The **mmsdrrestore** command cannot find a record for this node in the specified cluster configuration file. The search of the file is based on the currently active IP addresses of the node as reported by the **ifconfig** command.

User response:
Ensure that all adapter interfaces are properly functioning. Ensure that the correct GPFS configuration file is specified on the command line. If the node indeed is not a member of the cluster, use the **mmaddnode** command instead.

6027-1386 Unexpected value for Gpfs object: *value*.

Explanation:
A function received a value that is not allowed for the Gpfs object.

User response:
Perform problem determination.

6027-1388 File system *fileSystem* is not known to the GPFS cluster.

Explanation:
The file system was not found in the GPFS cluster.

User response:
If the file system was specified as part of a GPFS command, reissue the command with a valid file system.

6027-1390 Node *node* does not belong to the GPFS cluster, or was specified as input multiple times.

Explanation:
Nodes that are not valid were specified.

User response:
Verify the list of nodes. All specified nodes must belong to the GPFS cluster, and each node can be specified only once.

6027-1393 Incorrect node designation specified: *type*.

Explanation:
A node designation that is not valid was specified. Valid values are **client** or **manager**.

User response:
Correct the command line and reissue the command.

6027-1394 Operation not allowed for the local cluster.

Explanation:
The requested operation cannot be performed for the local cluster.

User response:
Specify the name of a remote cluster.

6027-1397 Existing virtual shared disk *diskName* cannot be renamed.

Explanation:
A disk descriptor passed to the command specified an existing virtual shared disk and a desired disk name. An existing virtual shared disk cannot be renamed.

User response:
Change the disk descriptor to specify a disk that is not an existing virtual shared disk or to not specify a desired disk name, and then reissue the command.

6027-1450 Could not allocate storage.

Explanation:
Sufficient memory cannot be allocated to run the **mmsanrepairfs** command.

User response:
Increase the amount of memory available.

6027-1500 Open of *device* failed with error:

Explanation:
The **open()** of a device failed. Operation of the file system may continue unless this device is needed for operation. If this is a replicated disk device, it will often not be needed. If this is a block or character device for another subsystem (such as **/dev/VSD0**) then GPFS will discontinue operation.

User response:
Problem diagnosis will depend on the subsystem that the device belongs to. For example, device **/dev/VSD0** belongs to the IBM Virtual Shared Disk subsystem, and problem determination should follow guidelines in that subsystem's documentation. If this is a normal disk device, take needed repair action on the specified disk.

6027-1501 Volume label of disk *name* is *name*, should be *uid*.

Explanation:
The UID in the disk descriptor does not match the expected value from the file system descriptor. This could occur if a disk was overwritten by another application or if the IBM Virtual Shared Disk subsystem incorrectly identified the disk.

User response:
Check the disk configuration.

6027-1502 Volume label of disk *diskName* is corrupt.

Explanation:
The disk descriptor has a bad magic number, version, or checksum. This could occur if a disk was overwritten by another application or if the IBM Virtual Shared Disk subsystem incorrectly identified the disk.

User response:
Check the disk configuration.

6027-1503 Completed adding disks to file system *name*.

Explanation:
The **mmadddisk** command successfully completed.

User response:
None. Informational message only.

6027-1504 File *name* could not be run:

Explanation:
Failed in trying to run an external program.

User response:
Check file existence and access permissions.

6027-1505 Could not get minor number for *name*.

Explanation:
Could not obtain a minor number for the specified block or character device.

User response:
Problem diagnosis will depend on the subsystem that the device belongs to. For example, device **/dev/VSD0** belongs to the IBM Virtual Shared Disk subsystem and problem determination should follow guidelines in that subsystem's documentation.

6027-1507 READ_KEYS ioctl failed with errno=*returnCode*, tried *timesTried* times. Related values are scsi_status=*scsiStatusValue*, sense_key=*senseKeyValue*, scsi_asc=*scsiAscValue*, scsi_ascq=*scsiAscqValue*.

Explanation:
A **READ_KEYS ioctl** call failed with the **errno** value and related values shown.

User response:
Check the reported **errno** value and try to correct the problem. If the problem persists, contact the IBM Support Center.

6027-1508 Registration failed with errno=*returnCode*, tried *timesTried* times. Related values are: scsi_status=*scsiStatusValue*, sense_key=*senseKeyValue*, scsi_asc=*scsiAscValue*, scsi_ascq=*scsiAscqValue*.

Explanation:
A **REGISTER ioctl** call failed with the **errno** value and related values shown.

User response:
Check the reported **errno** value and try to correct the problem. If the problem persists, contact the IBM Support Center.

6027-1509 READRES ioctl failed with errno=*returnCode*, tried *timesTried* times. Related values are: scsi_status=*scsiStatusValue*, sense_key=*senseKeyValue*, scsi_asc=*scsiAscValue*, scsi_ascq=*scsiAscqValue*.

Explanation:
A **READRES ioctl** call failed with the **errno** value and related values shown.

User response:
Check the reported **errno** value and try to correct the problem. If the problem persists, contact the IBM Support Center.

6027-1510 Error mounting file system *stripeGroup* **on** *mountPoint*; *errorQualifier* **(***gpfsErrno***).**

Explanation:
An error occurred while attempting to mount a GPFS file system on Windows.

User response:
Examine the error details, previous errors, and the GPFS message log to identify the cause.

6027-1511 Error unmounting file system *stripeGroup*; *errorQualifier* **(***gpfsErrno***).**

Explanation:
An error occurred while attempting to unmount a GPFS file system on Windows.

User response:
Examine the error details, previous errors, and the GPFS message log to identify the cause.

6027-1512 WMI query for *queryType* **failed;** *errorQualifier* **(***gpfsErrno***).**

Explanation:
An error occurred while running a WMI query on Windows.

User response:
Examine the error details, previous errors, and the GPFS message log to identify the cause.

6027-1530 Attention: *parameter* **is set to** *value***.**

Explanation:
A configuration parameter is temporarily assigned a new value.

User response:
Check the **mmfs.cfg** file. Use the **mmchconfig** command to set a valid value for the parameter.

6027-1531 *parameter value*

Explanation:
The configuration parameter was changed from its default value.

User response:
Check the **mmfs.cfg** file.

6027-1532 Attention: *parameter* **(***value***) is not valid in conjunction with** *parameter* **(***value***).**

Explanation:
A configuration parameter has a value that is not valid in relation to some other parameter. This can also happen when the default value for some parameter is not sufficiently large for the new, user set value of a related parameter.

User response:
Check the **mmfs.cfg** file.

6027-1533 *parameter* **cannot be set dynamically.**

Explanation:
The **mmchconfig** command encountered a configuration parameter that cannot be set dynamically.

User response:
Check the **mmchconfig** command arguments. If the parameter must be changed, use the **mmshutdown**, **mmchconfig**, and **mmstartup** sequence of commands.

6027-1534 *parameter* **must have a value.**

Explanation:
The **tsctl** command encountered a configuration parameter that did not have a specified value.

User response:
Check the **mmchconfig** command arguments.

6027-1535 Unknown config name: *parameter*

Explanation:
The **tsctl** command encountered an unknown configuration parameter.

User response:
Check the **mmchconfig** command arguments.

6027-1536 *parameter* **must be set using the tschpool command.**

Explanation:
The **tsctl** command encountered a configuration parameter that must be set using the **tschpool** command.

User response:
Check the **mmchconfig** command arguments.

6027-1537 Connect failed to *ipAddress*: *reason*

Explanation:
An attempt to connect sockets between nodes failed.

User response:
Check the reason listed and the connection to the indicated IP address.

6027-1538 Connect in progress to *ipAddress*

Explanation:
Connecting sockets between nodes.

User response:
None. Information message only.

6027-1539 Connect progress select failed to *ipAddress*: *reason*

Explanation:
An attempt to connect sockets between nodes failed.

User response:
Check the reason listed and the connection to the indicated IP address.

6027-1540 Try and buy license has expired!

Explanation:
Self explanatory.

User response:
Purchase a GPFS license to continue using GPFS.

6027-1541 Try and buy license expires in *number* **days.**

Explanation:
Self-explanatory.

User response:
When the **Try and Buy** license expires, you will need to purchase a GPFS license to continue using GPFS.

6027-1542 Old shared memory exists but it is not valid nor cleanable.

Explanation:
A new GPFS daemon started and found existing shared segments. The contents were not recognizable, so the GPFS daemon could not clean them up.

User response:

1. Stop the GPFS daemon from trying to start by issuing the **mmshutdown** command for the nodes having the problem.
2. Find the owner of the shared segments with keys from `0x9283a0ca` through `0x9283a0d1`. If a non-GPFS program owns these segments, GPFS cannot run on this node.
3. If these segments are left over from a previous GPFS daemon:
   a. Remove them by issuing:
      `ipcrm -m shared_memory_id`
   b. Restart GPFS by issuing the **mmstartup** command on the affected nodes.

6027-1543 error propagating *parameter*.

Explanation:
mmfsd could not propagate a configuration parameter value to one or more nodes in the cluster.

User response:
Contact the IBM Support Center.

6027-1544 Sum of prefetchthreads (*value***) and worker1threads (***value***) exceeds** *value*. **Reducing them to** *value* **and** *value*.

Explanation:
The sum of prefetchthreads and worker1threads exceeds the permitted value.

User response:
Accept the calculated values or reduce the individual settings using **mmchconfig prefetchthreads=newvalue** or **mmchconfig worker1threads=newvalue** or both. After using **mmchconfig**, the new settings will not take affect until the GPFS daemon is restarted.

6027-1545 The GPFS product that you are attempting to run is not a fully functioning version. This probably means that this is an update version and not the full product version. Install the GPFS full product version first, then apply any applicable update version before attempting to start GPFS.

Explanation:
GPFS requires a fully licensed GPFS installation.

User response:
Verify installation of licensed GPFS, or purchase and install a licensed version of GPFS.

6027-1546 Attention: *parameter* **size of** *value* **is too small. New value is** *value*.

Explanation:
A configuration parameter is temporarily assigned a new value.

User response:
Check the **mmfs.cfg** file. Use the **mmchconfig** command to set a valid value for the parameter.

6027-1547 Error initializing daemon: performing shutdown.

Explanation:
GPFS kernel extensions are not loaded, and the daemon cannot initialize. GPFS may have been started incorrectly.

User response:
Check GPFS log for errors resulting from kernel extension loading. Ensure that GPFS is started with the **mmstartup** command.

6027-1548 Error: daemon and kernel extension do not match.

Explanation:

The GPFS kernel extension loaded in memory and the daemon currently starting do not appear to have come from the same build.

User response:
Ensure that the kernel extension was reloaded after upgrading GPFS. See “GPFS modules cannot be loaded” on page 44 for details.

6027-1549 Attention: custom-built kernel extension; the daemon and kernel extension do not match.

Explanation:
The GPFS kernel extension loaded in memory does not come from the same build as the starting daemon. The kernel extension appears to have been built from the kernel open source package.

User response:
None.

6027-1550 Error: Unable to establish a session with an Active Directory server. ID remapping via Microsoft Identity Management for Unix will be unavailable.

Explanation:
GPFS tried to establish an LDAP session with an Active Directory server (normally the domain controller host), and has been unable to do so.

User response:
Ensure the domain controller is available.

6027-1555 Mount point and device name cannot be equal: *name*

Explanation:
The specified mount point is the same as the absolute device name.

User response:
Enter a new device name or absolute mount point path name.

6027-1556 Interrupt received.

Explanation:
A GPFS administration command received an interrupt.

User response:
None. Informational message only.

6027-1557 You must first generate an authentication key file. Run: mmauth genkey new.

Explanation:
Before setting a cipher list, you must generate an authentication key file.

User response:
Run the specified command to establish an authentication key for the nodes in the cluster.

6027-1558 Disk name *diskName* **already exists in the cluster configuration file as the name of an NSD or of a disk underlying an NSD. If the cited disk no longer exists, remove the associated NSD from the cluster configuration file by issuing: mmdelnsd** *nsdName***.**

Explanation:
The command cannot create the specified disk because a disk with the same name is already recorded in the cluster configuration file.

User response:
If the specified disk no longer exists, issue the **mmdelnsd** command to remove the associated NSD from the cluster configuration file.

6027-1559 The -i option failed. Changes will take effect after GPFS is restarted.

Explanation:
The **-i** option on the **mmchconfig** command failed. The changes were processed successfully, but will take effect only after the GPFS daemons are restarted.

User response:
Check for additional error messages. Correct the problem and reissue the command.

6027-1560 This GPFS cluster contains file systems. You cannot delete the last node.

Explanation:
An attempt has been made to delete a GPFS cluster that still has one or more file systems associated with it.

User response:
Before deleting the last node of a GPFS cluster, delete all file systems that are associated with it. This applies to both local and remote file systems.

6027-1561 Attention: Failed to remove node-specific changes.

Explanation:
The internal **mmfixcfg** routine failed to remove node-specific configuration settings, if any, for one or more of the nodes being deleted. This is of consequence only if the **mmchconfig** command was indeed used to establish node specific settings and these nodes are later added back into the cluster.

User response:
If you add the nodes back later, ensure that the configuration parameters for the nodes are set as desired.

6027-1562 *command* **command cannot be executed. Either none of the nodes in the cluster are reachable, or GPFS is down on all of the nodes.**

Explanation:
The command that was issued needed to perform an operation on a remote node, but none of the nodes in the cluster were reachable, or GPFS was not accepting commands on any of the nodes.

User response:
Ensure that the affected nodes are available and all authorization requirements are met. Correct any problems and reissue the command.

6027-1563 **Attention: The file system may no longer be properly balanced.**

Explanation:
The restripe phase of the **mmadddisk** or **mmdeldisk** command failed.

User response:
Determine the cause of the failure and run the **mmrestripefs -b** command.

6027-1564 **To change the authentication key for the local cluster, run: mmauth genkey.**

Explanation:
The authentication keys for the local cluster must be created only with the specified command.

User response:
Run the specified command to establish a new authentication key for the nodes in the cluster.

6027-1565 *disk* **not found in file system** *fileSystem***.**

Explanation:
A disk specified for deletion or replacement does not exist.

User response:
Specify existing disks for the indicated file system.

6027-1566 **Remote cluster** *clusterName* **is already defined.**

Explanation:
A request was made to add the cited cluster, but the cluster is already known to GPFS.

User response:
None. The cluster is already known to GPFS.

6027-1567 *fileSystem* **from cluster** *clusterName* **is already defined.**

Explanation:
A request was made to add the cited file system from the cited cluster, but the file system is already known to GPFS.

User response:
None. The file system is already known to GPFS.

6027-1568 *command* **command failed. Only** *parameterList* **changed.**

Explanation:
The **mmchfs** command failed while making the requested changes. Any changes to the attributes in the indicated parameter list were successfully completed. No other file system attributes were changed.

User response:
Reissue the command if you want to change additional attributes of the file system. Changes can be undone by issuing the **mmchfs** command with the original value for the affected attribute.

6027-1569 **The volume group** *volumeName* **exists, but partition information cannot be determined. Perhaps it needs to be varied on?**

Explanation:
The **mmvsdhelper** subroutine has successfully found the requested volume group, but attempts to get the free physical partitions with the **lsvg** command are failing.

User response:
Ensure that the volume group is varied on with the **varyonvg** command.

6027-1570 **virtual shared disk support is not installed.**

Explanation:
The command detected that IBM Virtual Shared Disk support is not installed on the node on which it is running.

User response:
Install IBM Virtual Shared Disk support.

6027-1571 *commandName* **does not exist or failed; automount mounting may not work.**

Explanation:
One or more of the GPFS file systems were defined with the automount attribute but the requisite automount command is missing or failed.

User response:
Correct the problem and restart GPFS. Or use the **mount** command to explicitly mount the file system.

6027-1572 The command must run on a node that is part of the cluster.

Explanation:
The node running the **mmcrcluster** command (this node) must be a member of the GPFS cluster.

User response:
Issue the command from a node that will belong to the cluster.

6027-1573 Command completed: No changes made.

Explanation:
Informational message.

User response:
Check the preceding messages, correct any problems, and reissue the command.

6027-1574 Permission failure. The command requires root authority to execute.

Explanation:
The command, or the specified command option, requires root authority.

User response:
Log on as 'root' and reissue the command.

6027-1578 File *fileName* does not contain node names.

Explanation:
The specified file does not contain valid node names.

User response:
Node names must be specified one per line. The name **localhost** and lines that start with '#' character are ignored.

6027-1579 File *fileName* does not contain data.

Explanation:
The specified file does not contain data.

User response:
Verify that you are specifying the correct file name and reissue the command.

6027-1580 Failed to obtain Kerberos credentials; ksrvtgt *name* command rc=*value*. Continuing.

Explanation:
An attempt to obtain **SPbgAdm** Kerberos credentials has failed. Processing continues, but there may be an authentication failure later on.

User response:
Check the preceding messages. Ensure that **_PRINCIPLE root.SPbgAdm** is defined in the file **/etc/sysctl.mmcmd.acl**. If the problem persists, go to the *Parallel System Support Programs for AIX: Diagnosis Guide* and search on *diagnosing per node key management (PNKM) problems*. Follow the problem determination and repair actions specified.

6027-1585 Primary server *nodeName* specified on the disk descriptor for previously-created virtual shared disk *diskName* does not match the primary server (*nodeName*) currently in effect for the virtual shared disk.

Explanation:
The primary server that was specified on the disk descriptor for the cited previously-created virtual shared disk is different than the primary server currently in effect for that virtual shared disk.

User response:
Either correct or omit the primary server specified on the descriptor and try the command again.

6027-1586 Backup server *nodeName* specified on the disk descriptor for previously-created virtual shared disk *diskName* does not match the backup server (*nodeName*) currently in effect for the virtual shared disk.

Explanation:
The backup server that was specified on the disk descriptor for the cited previously-created virtual shared disk is different than the backup server currently in effect for that virtual shared disk.

User response:
Either correct or omit the backup server specified on the descriptor and try the command again.

6027-1587 Unable to determine the local device name for disk *nsdName* on node *nodeName*.

Explanation:
GPFS was unable to determine the local device name for the specified GPFS disk.

User response:
Determine why the specified disk on the specified node could not be accessed and correct the problem. Possible reasons include: connectivity problems, authorization problems, fenced disk, and so forth.

6027-1588 Unknown GPFS execution environment: *value*

Explanation:
A GPFS administration command (prefixed by **mm**) was asked to operate on an unknown GPFS cluster type.

The only supported GPFS cluster type is **lc**. This message may also be generated if there is corruption in the GPFS system files.

User response:
Verify that the correct level of GPFS is installed on the node. If this is a cluster environment, make sure the node has been defined as a member of the GPFS cluster with the help of the **mmcrcluster** or the **mmaddnode** command. If the problem persists, contact the IBM Support Center.

6027-1590 *nodeName* **cannot be reached.**

Explanation:
A command needs to issue a remote function on a particular node but the node is not reachable.

User response:
Determine why the node is unreachable, correct the problem, and reissue the command.

6027-1591 **Attention: Unable to retrieve GPFS cluster files from node** *nodeName***.**

Explanation:
A command could not retrieve the GPFS cluster files from a particular node. An attempt will be made to retrieve the GPFS cluster files from a backup node.

User response:
None. Informational message only.

6027-1592 **Unable to retrieve GPFS cluster files from node** *nodeName***.**

Explanation:
A command could not retrieve the GPFS cluster files from a particular node.

User response:
Correct the problem and reissue the command.

6027-1594 **Run the** *command* **command until successful.**

Explanation:
The command could not complete normally. The GPFS cluster data may be left in a state that precludes normal operation until the problem is corrected.

User response:
Check the preceding messages, correct the problems, and issue the specified command until it completes successfully.

6027-1595 **No nodes were found that matched the input specification.**

Explanation:
No nodes were found in the GPFS cluster that matched those specified as input to a GPFS command.

User response:
Determine why the specified nodes were not valid, correct the problem, and reissue the GPFS command.

6027-1596 **The same node was specified for both the primary and the secondary server.**

Explanation:
A command would have caused the primary and secondary GPFS cluster configuration server nodes to be the same.

User response:
Specify a different primary or secondary node.

6027-1597 **Node** *node* **is specified more than once.**

Explanation:
The same node appears more than once on the command line or in the input file for the command.

User response:
All specified nodes must be unique. Note that even though two node identifiers may appear different on the command line or in the input file, they may still refer to the same node.

6027-1598 **Node** *nodeName* **was not added to the cluster. The node appears to already belong to a GPFS cluster.**

Explanation:
A GPFS cluster command found that a node to be added to a cluster already has GPFS cluster files on it.

User response:
Use the **mmlscluster** command to verify that the node is in the correct cluster. If it is not, follow the procedure in “Node cannot be added to the GPFS cluster” on page 54.

6027-1599 **The level of GPFS on node** *nodeName* **does not support the requested action.**

Explanation:
A GPFS command found that the level of the GPFS code on the specified node is not sufficient for the requested action.

User response:
Install the correct level of GPFS.

6027-1600 **Make sure that the following nodes are available:** *nodeList*

Explanation:
A GPFS command was unable to complete because nodes critical for the success of the operation were not reachable or the command was interrupted.

User response:

This message will normally be followed by a message telling you which command to issue as soon as the problem is corrected and the specified nodes become available.

6027-1602 *nodeName* **is not a member of this cluster.**

Explanation:
A command found that the specified node is not a member of the GPFS cluster.

User response:
Correct the input or add the node to the GPFS cluster and reissue the command.

6027-1603 **The following nodes could not be added to the GPFS cluster:** *nodeList***. Correct the problems and use the mmaddnode command to add these nodes to the cluster.**

Explanation:
The **mmcrcluster** or the **mmaddnode** command was unable to add the listed nodes to a GPFS cluster.

User response:
Correct the problems and add the nodes to the cluster using the **mmaddnode** command.

6027-1604 **Information cannot be displayed. Either none of the nodes in the cluster are reachable, or GPFS is down on all of the nodes.**

Explanation:
The command needed to perform an operation on a remote node, but none of the nodes in the cluster were reachable, or GPFS was not accepting commands on any of the nodes.

User response:
Ensure that the affected nodes are available and all authorization requirements are met. Correct any problems and reissue the command.

6027-1610 **Disk** *diskName* **is the only disk in file system** *fileSystem***. You cannot replace a disk when it is the only remaining disk in the file system.**

Explanation:
The **mmrpldisk** command was issued, but there is only one disk in the file system.

User response:
Add a second disk and reissue the command.

6027-1613 **WCOLL (working collective) environment variable not set.**

Explanation:
The **mmdsh** command was invoked without explicitly specifying the nodes on which the command is to run by means of the **-F** or **-L** options, and the WCOLL environment variable has not been set.

User response:
Change the invocation of the **mmdsh** command to use the **-F** or **-L** options, or set the WCOLL environment variable before invoking the **mmdsh** command.

6027-1614 **Cannot open file** *fileName***. Error string was:** *errorString***.**

Explanation:
The **mmdsh** command was unable to successfully open a file.

User response:
Determine why the file could not be opened and correct the problem.

6027-1615 *nodeName* **remote shell process had return code** *value***.**

Explanation:
A child remote shell process completed with a nonzero return code.

User response:
Determine why the child remote shell process failed and correct the problem.

6027-1616 **Caught SIG** *signal* **- terminating the child processes.**

Explanation:
The **mmdsh** command has received a signal causing it to terminate.

User response:
Determine what caused the signal and correct the problem.

6027-1617 **There are no available nodes on which to run the command.**

Explanation:
The **mmdsh** command found that there are no available nodes on which to run the specified command. Although nodes were specified, none of the nodes were reachable.

User response:
Determine why the specified nodes were not available and correct the problem.

6027-1618 Unable to pipe. Error string was: *errorString***.**

Explanation:
The **mmdsh** command attempted to open a pipe, but the pipe command failed.

User response:
Determine why the call to pipe failed and correct the problem.

6027-1619 Unable to redirect *outputStream***. Error string was:** *string***.**

Explanation:
The **mmdsh** command attempted to redirect an output stream using open, but the open command failed.

User response:
Determine why the call to open failed and correct the problem.

6027-1623 *command***: Mounting file systems ...**

Explanation:
This message contains progress information about the **mmmount** command.

User response:
None. Informational message only.

6027-1625 *option* **cannot be used with attribute** *name***.**

Explanation:
An attempt was made to change a configuration attribute and requested the change to take effect immediately (**-i** or **-I** option). However, the specified attribute does not allow the operation.

User response:
If the change must be made now, leave off the **-i** or **-I** option. Then recycle the nodes to pick up the new value.

6027-1626 Command is not supported in the *type* **environment.**

Explanation:
A GPFS administration command (**mm...**) is not supported in the specified environment.

User response:
Verify if the task is needed in this environment, and if it is, use a different command.

6027-1627 The following nodes are not aware of the configuration server change: *nodeList***. Do not start GPFS on the above nodes until the problem is resolved.**

Explanation:
The **mmchcluster** command cannot propagate the new cluster configuration servers to the specified nodes.

User response:
Correct the problems and run the **mmchcluster -p LATEST** command before starting GPFS on the specified nodes.

6027-1628 Cannot determine basic environment information. Not enough nodes are available.

Explanation:
The **mmchcluster** command was unable to retrieve the GPFS cluster data files. Usually, this is due to too few nodes being available.

User response:
Correct any problems and ensure that as many of the nodes in the cluster are available as possible. Reissue the command. If the problem persists, record the above information and contact the IBM Support Center.

6027-1629 Error retrieving data from *nodeName* **to** *nodeName***.**

Explanation:
A GPFS command is unable to correctly copy data (checksum error).

User response:
Correct any communication problems and reissue the command.

6027-1630 The GPFS cluster data on *nodeName* **is back level.**

Explanation:
A GPFS command attempted to commit changes to the GPFS cluster configuration data, but the data on the server is already at a higher level. This can happen if the GPFS cluster configuration files were altered outside the GPFS environment, or if the **mmchcluster** command did not complete successfully.

User response:
Correct any problems and reissue the command. If the problem persists, issue the **mmrefresh -f -a** command.

6027-1631 The commit process failed.

Explanation:
A GPFS administration command (**mm...**) cannot commit its changes to the GPFS cluster configuration data.

User response:
Examine the preceding messages, correct the problem, and reissue the command. If the problem persists, perform problem determination and contact the IBM Support Center.

6027-1632 The GPFS cluster configuration data on *nodeName* is different than the data on *nodeName*.

Explanation:
The GPFS cluster configuration data on the primary cluster configuration server node is different than the data on the secondary cluster configuration server node. This can happen if the GPFS cluster configuration files were altered outside the GPFS environment or if the **mmchcluster** command did not complete successfully.

User response:
Correct any problems and issue the **mmrefresh -f -a** command. If the problem persists, perform problem determination and contact the IBM Support Center.

6027-1633 Failed to create a backup copy of the GPFS cluster data on *nodeName*.

Explanation:
Commit could not create a correct copy of the GPFS cluster configuration data.

User response:
Check the preceding messages, correct any problems, and reissue the command. If the problem persists, perform problem determination and contact the IBM Support Center.

6027-1634 The GPFS cluster configuration server node *nodeName* cannot be removed.

Explanation:
An attempt was made to delete a GPFS cluster configuration server node.

User response:
You cannot remove a cluster configuration server node unless all nodes in the GPFS cluster are being deleted. Before deleting a cluster configuration server node, you must use the **mmchcluster** command to transfer its function to another node in the GPFS cluster.

6027-1636 Error found while checking disk descriptor: *descriptor*.

Explanation:
A disk descriptor was found to be unsatisfactory in some way.

User response:
Check the preceding messages, if any, and correct the condition that caused the disk descriptor to be rejected.

6027-1637 *command* quitting. None of the specified nodes are valid.

Explanation:
A GPFS command found that none of the specified nodes passed the required tests.

User response:
Determine why the nodes were not accepted, fix the problems, and reissue the command.

6027-1638 *Command*: There are no unassigned nodes in the cluster.

Explanation:
A GPFS command in a cluster environment needs unassigned nodes, but found there are none.

User response:
Verify whether there are any unassigned nodes in the cluster. If there are none, either add more nodes to the cluster using the **mmaddnode** command, or delete some nodes from the cluster using the **mmdelnode** command, and then reissue the command.

6027-1639 Command failed. Examine previous error messages to determine cause.

Explanation:
A GPFS command failed due to previously-reported errors.

User response:
Check the previous error messages, fix the problems, and then reissue the command.

6027-1642 *command*: Starting GPFS ...

Explanation:
Progress information for the **mmstartup** command.

User response:
None. Informational message only.

6027-1643 The number of quorum nodes exceeds the maximum (*number*) allowed.

Explanation:
An attempt was made to add more quorum nodes to a cluster than the maximum number allowed.

User response:
Reduce the number of quorum nodes, and reissue the command.

6027-1644 Attention: The number of quorum nodes exceeds the suggested maximum (*number*).

Explanation:
The number of quorum nodes in the cluster exceeds the maximum suggested number of quorum nodes.

User response:
Informational message. Consider reducing the number of quorum nodes to the maximum suggested number of quorum nodes for improved performance.

6027-1645 Node *nodeName* is fenced out from disk *diskName*.

Explanation:
A GPFS command attempted to access the specified disk, but found that the node attempting the operation was fenced out from the disk.

User response:
Check whether there is a valid reason why the node should be fenced out from the disk. If there is no such reason, unfence the disk and reissue the command.

6027-1647 Unable to find disk with NSD volume id *NSD volume id*.

Explanation:
A disk with the specified NSD volume id cannot be found.

User response:
Specify a correct disk NSD volume id.

6027-1648 GPFS was unable to obtain a lock from node *nodeName*.

Explanation:
GPFS failed in its attempt to get a lock from another node in the cluster.

User response:
Verify that the reported node is reachable. Examine previous error messages, if any. Fix the problems and then reissue the command.

6027-1661 Failed while processing disk descriptor *descriptor* on node *nodeName*.

Explanation:
A disk descriptor was found to be unsatisfactory in some way.

User response:
Check the preceding messages, if any, and correct the condition that caused the disk descriptor to be rejected.

6027-1662 Disk descriptor *descriptor* refers to an existing NSD *name*.

Explanation:
The specified disk descriptor refers to an existing NSD.

User response:
Specify another disk that is not an existing NSD.

6027-1663 Disk descriptor *descriptor* should refer to an existing NSD. Use mmcrnsd to create the NSD.

Explanation:
An NSD disk given as input is not known to GPFS.

User response:
Create the NSD. Then rerun the command.

6027-1664 *command*: Processing node *nodeName*

Explanation:
Progress information.

User response:
None. Informational message only.

6027-1665 Issue the command from a node that remains in the cluster.

Explanation:
The nature of the requested change requires the command be issued from a node that will remain in the cluster.

User response:
Run the command from a node that will remain in the cluster.

6027-1666 No disks were found.

Explanation:
A command searched for disks but found none.

User response:
If disks are desired, create some using the **mmcrnsd** command.

6027-1670 Incorrect or missing remote shell command: *name*

Explanation:
The specified remote command does not exist or is not executable.

User response:
Specify a valid command.

6027-1671 Incorrect or missing remote file copy command: *name*

Explanation:
The specified remote command does not exist or is not executable.

User response:
Specify a valid command.

6027-1672 *option value* parameter must be an absolute path name.

Explanation:
The mount point does not begin with '/'.

User response:
Specify the full path for the mount point.

6027-1674 *command*: **Unmounting file systems ...**

Explanation:
This message contains progress information about the **mmumount** command.

User response:
None. Informational message only.

6027-1677 **Disk** *diskName* **is of an unknown type.**

Explanation:
The specified disk is of an unknown type.

User response:
Specify a disk whose type is recognized by GPFS.

6027-1680 **Disk name** *diskName* **is already registered for use by GPFS.**

Explanation:
The cited disk name was specified for use by GPFS, but there is already a disk by that name registered for use by GPFS.

User response:
Specify a different disk name for use by GPFS and reissue the command.

6027-1681 **Node** *nodeName* **is being used as an NSD server.**

Explanation:
The specified node is defined as a server node for some disk.

User response:
If you are trying to delete the node from the GPFS cluster, you must either delete the disk or define another node as its server.

6027-1685 **Processing continues without lock protection.**

Explanation:
The command will continue processing although it was not able to obtain the lock that prevents other GPFS commands from running simultaneously.

User response:
Ensure that no other GPFS command is running. See the command documentation for additional details.

6027-1688 **Command was unable to obtain the lock for the GPFS system data. Unable to reach the holder of the lock** *nodeName*. **Check the preceding messages, if any. Follow the procedure outlined in the GPFS Problem Determination Guide.**

Explanation:
A command requires the lock for the GPFS system data but was not able to obtain it.

User response:
Check the preceding messages, if any. Follow the procedure in the *GPFS: Problem Determination Guide* for what to do when the GPFS system data is locked. Then reissue the command.

6027-1689 **vpath disk** *diskName* **is not recognized as an IBM SDD device.**

Explanation:
The **mmvsdhelper** command found that the specified disk is a vpath disk, but it is not recognized as an IBM SDD device.

User response:
Ensure the disk is configured as an IBM SDD device. Then reissue the command.

6027-1698 **Disk** *diskName* **belongs to vpath** *vpathName*.

Explanation:
The specified disk belongs to a vpath device.

User response:
Reissue the command specifying the vpath device, or specify another disk that does not belong to a vpath device.

6027-1699 **Remount failed for file system** *fileSystem*. **Error code** *errorCode*.

Explanation:
The specified file system was internally unmounted. An attempt to remount the file system failed with the specified error code.

User response:
Check the daemon log for additional error messages. Ensure that all file system disks are available and reissue the **mount** command.

6027-1700 **Failed to load LAPI library.** *functionName* **not found. Changing communication protocol to TCP.**

Explanation:
The GPFS daemon failed to load **liblapi_r.a** dynamically.

User response:
Verify installation of **liblapi_r.a**.

6027-1701 **mmfsd waiting to connect to mmspsecserver. Setting up to retry every** *number* **seconds for** *number* **minutes.**

Explanation:

The GPFS daemon failed to establish a connection with the **mmspsecserver** process.

User response:
None. Informational message only.

6027-1702 Process *pid* **failed at** *functionName* **call, socket** *socketName*, **errno** *value*

Explanation:
Either The **mmfsd** daemon or the **mmspsecserver** process failed to create or set up the communication socket between them.

User response:
Determine the reason for the error.

6027-1703 The *processName* **process encountered error:** *errorString*.

Explanation:
Either the **mmfsd** daemon or the **mmspsecserver** process called the error log routine to log an incident.

User response:
None. Informational message only.

6027-1704 mmspsecserver (pid *number*) **ready for service.**

Explanation:
The **mmspsecserver** process has created all the service threads necessary for **mmfsd**.

User response:
None. Informational message only.

6027-1705 *command*: **incorrect number of connections (***number*), **exiting...**

Explanation:
The **mmspsecserver** process was called with an incorrect number of connections. This will happen only when the **mmspsecserver** process is run as an independent program.

User response:
Retry with a valid number of connections.

6027-1706 mmspsecserver: parent program is not "mmfsd", exiting...

Explanation:
The **mmspsecserver** process was invoked from a program other than **mmfsd**.

User response:
None. Informational message only.

6027-1707 mmfsd connected to mmspsecserver

Explanation:
The **mmfsd** daemon has successfully connected to the **mmspsecserver** process through the communication socket.

User response:
None. Informational message only.

6027-1708 The mmfsd daemon failed to fork mmspsecserver. Failure reason *explanation*

Explanation:
The **mmfsd** daemon failed to fork a child process.

User response:
Check the GPFS installation.

6027-1709 Accepted and connected to *ipAddress*.

Explanation:
The local **mmfsd** daemon has successfully accepted and connected to a remote daemon.

User response:
None. Informational message only.

6027-1710 Connecting to *ipAddress*.

Explanation:
The local **mmfsd** daemon has started a connection request to a remote daemon.

User response:
None. Informational message only.

6027-1711 Connected to *ipAddress*.

Explanation:
The local **mmfsd** daemon has successfully connected to a remote daemon.

User response:
None. Informational message only.

6027-1712 Unexpected zero bytes received from *name*. **Continuing.**

Explanation:
This is an informational message. A socket read resulted in zero bytes being read.

User response:
If this happens frequently, check IP connections.

6027-1715 EINVAL trap from connect call to *ipAddress* **(socket** *name*)

Explanation:
The connect call back to the requesting node failed.

User response:

This is caused by a bug in AIX socket support. Upgrade AIX kernel and TCP client support.

6027-1716 Close connection to *ipAddress*

Explanation:
Connection socket closed.

User response:
None. Informational message only.

6027-1717 The administrator of *name* **does not allow remote mounts.**

Explanation:
The administrator of a file system has not configured remote mount.

User response:
Contact the administrator and request remote mount access.

6027-1718 The administrator of *name* **does not require secure connections. Unregister the target clusters key using: mmremotecluster update**

Explanation:
Connection socket closed.

User response:
None. Informational message only.

6027-1724 The key used by the cluster named *clusterName* **has changed. Contact the administrator to obtain the new key and register it using: mmremotecluster update**

Explanation:
The administrator of the cluster has changed the key used for authentication.

User response:
Contact the administrator to obtain the new key and register it using **mmremotecluster update**.

6027-1725 The key used by the cluster named *clusterName* **has changed. Contact the administrator to obtain the new key and register it using: mmauth update**

Explanation:
The administrator of the cluster has changed the key used for authentication.

User response:
Contact the administrator to obtain the new key and register it using **mmauth update**.

6027-1726 The administrator of the cluster named *clusterName* **requires authentication. Contact the administrator to obtain the cluster's key and register the key using: mmremotecluster update**

Explanation:
The administrator of the cluster requires authentication.

User response:
Contact the administrator to obtain the cluster's key and register it using: **mmremotecluster update**.

6027-1727 The administrator of the cluster named *clusterName* **does not require authentication. Unregister the cluster's key using: mmremotecluster update**

Explanation:
The administrator of the cluster does not require authentication.

User response:
Unregister the clusters key using: **mmremotecluster update**.

6027-1728 Remote mounts are not enabled within the cluster named *clusterName***. Contact the administrator and request that they enable remote mounts.**

Explanation:
The administrator of the cluster has not enabled remote mounts.

User response:
Contact the administrator and request remote mount access.

6027-1729 The cluster named *clusterName* **has not authorized this cluster to mount file systems. Contact the cluster administrator and request access.**

Explanation:
The administrator of the cluster has not authorized this cluster to mount file systems.

User response:
Contact the administrator and request access.

6027-1730 Unsupported cipherList *cipherList* **requested.**

Explanation:
The target cluster requested a **cipherList** not supported by the installed version of OpenSSL.

User response:
Install a version of OpenSSL that supports the required **cipherList** or contact the administrator of the target cluster and request that a supported **cipherList** be assigned to this remote cluster.

6027-1731 Unsupported cipherList *cipherList* requested.

Explanation:
The target cluster requested a **cipherList** that is not supported by the installed version of OpenSSL.

User response:
Either install a version of OpenSSL that supports the required **cipherList** or contact the administrator of the target cluster and request that a supported **cipherList** be assigned to this remote cluster.

6027-1732 Remote mounts are not enabled within this cluster.

Explanation:
Remote mounts cannot be performed in this cluster.

User response:
See the *GPFS: Advanced Administration Guide* for instructions about enabling remote mounts. In particular, make sure the keys have been generated and a **cipherlist** has been set.

6027-1733 OpenSSL dynamic lock support could not be loaded.

Explanation:
One of the functions required for dynamic lock support was not included in the version of the OpenSSL library that GPFS is configured to use.

User response:
If this functionality is required, shut down the daemon, install a version of OpenSSL with the desired functionality, and configure GPFS to use it. Then restart the daemon.

6027-1734 OpenSSL engine support could not be loaded.

Explanation:
One of the functions required for engine support was not included in the version of the OpenSSL library that GPFS is configured to use.

User response:
If this functionality is required, shut down the daemon, install a version of OpenSSL with the desired functionality, and configure GPFS to use it. Then restart the daemon.

6027-1735 Close connection to *ipAddress*. Attempting reconnect.

Explanation:
Connection socket closed. The GPFS daemon will attempt to reestablish the connection.

User response:
None. Informational message only.

6027-1736 Reconnected to *ipAddress*

Explanation:
The local **mmfsd** daemon has successfully reconnected to a remote daemon following an unexpected connection break.

User response:
None. Informational message only.

6027-1737 Close connection to *ipAddress* (*errorString*).

Explanation:
Connection socket closed.

User response:
None. Informational message only.

6027-1738 Close connection to *ipAddress* (*errorString*). Attempting reconnect.

Explanation:
Connection socket closed.

User response:
None. Informational message only.

6027-1739 Accept socket connection failed: err *value*.

Explanation:
The Accept socket connection received an unexpected error.

User response:
None. Informational message only.

6027-1740 Timed out waiting for a reply from node *ipAddress*.

Explanation:
A message that was sent to the specified node did not receive a response within the expected time limit.

User response:
None. Informational message only.

6027-1741 Error code *value* received from node *ipAddress*.

Explanation:
When a message was sent to the specified node to check its status, an error occurred and the node could not handle the message.

User response:
None. Informational message only.

6027-1742 Message ID *value* was lost by node *ipAddress*.

Explanation:
During a periodic check of outstanding messages, a problem was detected where the destination node no longer has any knowledge of a particular message.

User response:
None. Informational message only.

6027-1803 Global NSD disk, *name*, not found.

Explanation:
A client tried to open a globally-attached NSD disk, but a scan of all disks failed to find that NSD.

User response:
Ensure that the globally-attached disk is available on every node that references it.

6027-1804 I/O to NSD disk, *name*, fails. No such NSD locally found.

Explanation:
A server tried to perform I/O on an NSD disk, but a scan of all disks failed to find that NSD.

User response:
Make sure that the NSD disk is accessible to the client. If necessary, break a reservation.

6027-1805 Rediscovered NSD server access to *name*.

Explanation:
A server rediscovered access to the specified disk.

User response:
None.

6027-1806 A Persistent Reserve could not be established on device name (*deviceName*).

Explanation:
GPFS is using Persistent Reserve on this disk, but was unable to establish a reserve for this node.

User response:
Perform disk diagnostics.

6027-1807 NSD *nsdName* is using Persistent Reserve, this will require an NSD server on an *osName* node.

Explanation:
A client tried to open a globally-attached NSD disk, but the disk is using Persistent Reserve. An *osName* NSD server is needed. GPFS only supports Persistent Reserve on certain operating systems.

User response:
Use the **mmchnsd** command to add an *osName* NSD server for the NSD.

6027-1850 Successfully subscribed to *adapter* group.

Explanation:
GPFS has successfully subscribed to the adapter group.

User response:
None. Informational message only.

6027-1852 Adapter *adapter* is down. Wait for *number* sec.

Explanation:
The GPFS network is down.

User response:
None. Retries will be attempted every five seconds for a total of five minutes.

6027-1853 Waiting for *adapter* adapter group subscription.

Explanation:
GPFS could not subscribe to the adapter group.

User response:
None. Retries will be attempted every five seconds for a total of five minutes.

6027-1867 Node join: phase II node recovery in progress. Wait *value* secs.

Explanation:
Node join protocol is delayed until phase II of previous node failure recovery protocol is complete.

User response:
None.

6027-1868 Rejected node join protocol. Phase II node failure recovery still in progress, delays ended.

Explanation:
Node join protocol is rejected after a number of internal delays and phase II node failure protocol is still in progress.

User response:
None.

6027-1869 Phase II node recovery complete. *number* node(s) processed.

Explanation:
Informational. The asynchronous part (phase II) of node failure recovery has completed.

User response:
None.

6027-1871 Core quorum nodes and single node quorum cannot coexist.

Explanation:
A node has been specified as a quorum node in the cluster configuration, and single node quorum is also specified.

User response:
Either modify the single node quorum attribute or node designation using **mmchconfig**.

6027-1881 Cannot find failover group *groupName***.**

Explanation:
The failover group to be deleted does not exist.

User response:
Retry the command with a valid failover group.

6027-1900 Failed to stat *pathName***.**

Explanation:
A **stat()** call failed for the specified object.

User response:
Correct the problem and reissue the command.

6027-1901 *pathName* **is not a GPFS file system object.**

Explanation:
The specified path name does not resolve to an object within a mounted GPFS file system.

User response:
Correct the problem and reissue the command.

6027-1902 The policy file cannot be determined.

Explanation:
The command was not able to retrieve the policy rules associated with the file system.

User response:
Examine the preceding messages and correct the reported problems. Establish a valid policy file with the **mmchpolicy** command or specify a valid policy file on the command line.

6027-1903 *path* **must be an absolute path name.**

Explanation:
The path name did not begin with a **/**.

User response:
Specify the absolute path name for the object.

6027-1904 *number* **and** *number* **already exists.**

Explanation:
A device with the cited major and minor numbers already exists.

User response:
Check the preceding messages for detailed information.

6027-1905 *name* **was not created by GPFS or could not be refreshed.**

Explanation:
The attributes (device type, major/minor number) of the specified file system device name are not as expected.

User response:
Check the preceding messages for detailed information on the current and expected values. These errors are most frequently caused by the presence of **/dev** entries that were created outside the GPFS environment. Resolve the conflict by renaming or deleting the offending entries. Reissue the command letting GPFS create the **/dev** entry with the appropriate parameters.

6027-1908 The *option* **option is not allowed for remote file systems.**

Explanation:
The specified option can be used only for locally-owned file systems.

User response:
Correct the command line and reissue the command.

6027-1909 There are no available free disks. Disks must be prepared prior to invoking *command***. Define the disks using the** *command* **command.**

Explanation:
The currently executing command (**mmcrfs**, **mmadddisk**, **mmrpldisk**) requires disks to be defined for use by GPFS using one of the GPFS disk creation commands: **mmcrnsd**, **mmcrvsd**.

User response:
Create disks and reissue the failing command.

6027-1910 Node *nodeName* **is not a quorum node.**

Explanation:
The **mmchmgr** command was asked to move the cluster manager to a nonquorum node. Only one of the quorum nodes can be a cluster manager.

User response:
Designate the node to be a quorum node, specify a different node on the command line, or allow GPFS to choose the new cluster manager node.

6027-1911 File system *fileSystem* belongs to cluster *clusterName*. The *option* option is not allowed for remote file systems.

Explanation:
The specified option can be used only for locally-owned file systems.

User response:
Correct the command line and reissue the command.

6027-1920 *subsystem* not active on nodes: *nodeList*.

Explanation:
A GPFS command requires the specified subsystem to be up and in active state.

User response:
Correct the problems and reissue the command.

6027-1922 IP aliasing is not supported (*node*). Specify the main device.

Explanation:
IP aliasing is not supported.

User response:
Specify a node identifier that resolves to the IP address of a main device for the node.

6027-1927 The requested disks are not known to GPFS.

Explanation:
GPFS could not find the requested NSDs in the cluster.

User response:
Reissue the command, specifying known disks.

6027-1929 *cipherlist* is not a valid cipher list.

Explanation:
The cipher list must be set to a value supported by GPFS. All nodes in the cluster must support a common cipher.

User response:
See the GPFS Frequently Asked Questions at publib.boulder.ibm.com/infocenter/ clresctr/topic/ com.ibm.cluster.gpfs.doc/gpfs_faqs/ gpfsclustersfaq.html for a list of the supported ciphers.

6027-1930 Disk *diskName* belongs to file system *fileSystem*.

Explanation:
A GPFS administration command (**mm...**) found that the requested disk to be deleted still belongs to a file system.

User response:
Check that the correct disk was requested. If so, delete the disk from the file system before proceeding.

6027-1931 The following disks are not known to GPFS: *diskNames*.

Explanation:
A GPFS administration command (**mm...**) found that the specified disks are not known to GPFS.

User response:
Verify that the correct disks were requested.

6027-1932 No disks were specified that could be deleted.

Explanation:
A GPFS administration command (**mm...**) determined that no disks were specified that could be deleted.

User response:
Examine the preceding messages, correct the problems, and reissue the command.

6027-1933 Disk *diskName* has been removed from the GPFS cluster configuration data but the NSD volume id was not erased from the disk. To remove the NSD volume id, issue mmdelnsd -p *NSDvolumeid*.

Explanation:
A GPFS administration command (**mm...**) successfully removed the specified disk from the GPFS cluster configuration data, but was unable to erase the NSD volume id from the disk.

User response:
Issue the specified command to remove the NSD volume id from the disk.

6027-1934 Disk *diskName* has been removed from the GPFS cluster configuration data but the NSD volume id was not erased from the disk. To remove the NSD volume id, issue: mmdelnsd -p *NSDvolumeid* -N *nodeList*.

Explanation:
A GPFS administration command (**mm...**) successfully removed the specified disk from the GPFS cluster configuration data but was unable to erase the NSD volume id from the disk.

User response:
Issue the specified command to remove the NSD volume id from the disk.

6027-1936 **Node** *nodeName* **cannot support Persistent Reserve on disk** *diskName* **because it is not an AIX node. The disk will be used as a non-PR disk.**

Explanation:
A non-AIX node was specified as an NSD server for the disk. The disk will be used as a non-PR disk.

User response:
None. Informational message only.

6027-1937 **A node was specified more than once as an NSD server in disk descriptor** *descriptor***.**

Explanation:
A node was specified more than once as an NSD server in the disk descriptor shown.

User response:
Change the disk descriptor to eliminate any redundancies in the list of NSD servers.

6027-1938 *configParameter* **is an incorrect parameter. Line in error:** *configLine***. The line is ignored; processing continues.**

Explanation:
The specified parameter is not valid and will be ignored.

User response:
None. Informational message only.

6027-1939 **Line in error:** *line***.**

Explanation:
The specified line from a user-provided input file contains errors.

User response:
Check the preceding messages for more information. Correct the problems and reissue the command.

6027-1940 **Unable to set reserve policy** *policy* **on disk** *diskName* **on node** *nodeName***.**

Explanation:
The specified disk should be able to support Persistent Reserve, but an attempt to set up the registration key failed.

User response:
Correct the problem and reissue the command.

6027-1941 **Cannot handle multiple interfaces for host** *hostName***.**

Explanation:
Multiple entries were found for the given hostname or IP address either in **/etc/hosts** or by the **host** command.

User response:
Make corrections to **/etc/hosts** and reissue the command.

6027-1942 **Unexpected output from the 'host -t a** *name*' **command:**

Explanation:
A GPFS administration command (**mm...**) received unexpected output from the **host -t a** command for the given host.

User response:
Issue the **host -t a** command interactively and carefully review the output, as well as any error messages.

6027-1943 **Host** *name* **not found.**

Explanation:
A GPFS administration command (**mm...**) could not resolve a host from **/etc/hosts** or by using the **host** command.

User response:
Make corrections to **/etc/hosts** and reissue the command.

6027-1944 **Desired disk name** *diskName* **is longer than 13 characters.**

Explanation:
The cited disk name is not valid because it is longer than the maximum allowed length of 13 characters.

User response:
Specify a disk name whose length is 13 characters or less and reissue the command.

6027-1945 **Disk name** *diskName* **is not allowed. Names beginning with gpfs are reserved for use by GPFS.**

Explanation:
The cited disk name is not allowed because it begins with **gpfs**.

User response:
Specify a disk name that does not begin with **gpfs** and reissue the command.

6027-1947 **File** *fileName* **not found. Recover the file, or generate and commit a new key using: mmauth genkey.**

Explanation:
The cited file was not found.

User response:
Recover the file, or generate a new key by running: **mmauth genkey new**, followed by the **mmauth genkey commit** command.

6027-1948 Disk *diskName* is too large.

Explanation:
The specified disk is too large.

User response:
Specify a smaller disk and reissue the command.

6027-1949 Unexpected error obtaining the local environment update lock.

Explanation:
GPFS was unable to obtain the local environment update lock.

User response:
Examine previous error messages, if any. Correct any problems and reissue the command. If the problem persists, perform problem determination and contact the IBM Support Center.

6027-1950 Local update lock is busy.

Explanation:
More than one process is attempting to update the GPFS environment at the same time.

User response:
Repeat the command. If the problem persists, verify that there are no blocked processes.

6027-1951 Failed to obtain the local environment update lock.

Explanation:
GPFS was unable to obtain the local environment update lock for more than 30 seconds.

User response:
Examine previous error messages, if any. Correct any problems and reissue the command. If the problem persists, perform problem determination and contact the IBM Support Center.

6027-1962 Permission denied for disk *diskName*

Explanation:
The user does not have permission to access disk *diskName*.

User response:
Correct the permissions and reissue the command.

6027-1963 Disk *diskName* was not found.

Explanation:
The specified disk was not found.

User response:
Specify an existing disk and reissue the command.

6027-1964 I/O error on *diskName*

Explanation:
An I/O error occurred on the specified disk.

User response:
Check for additional error messages. Check the error log for disk hardware problems.

6027-1965 *logicalVolume* is not a valid logical volume.

Explanation:
The specified logical volume is not a valid logical volume with a corresponding volume group.

User response:
Reissue the command using a valid logical volume.

6027-1967 Disk *diskName* belongs to back-level file system *fileSystem* or the state of the disk is not ready. Use mmchfs -V to convert the file system to the latest format. Use mmchdisk to change the state of a disk.

Explanation:
The specified disk cannot be initialized for use as a tiebreaker disk. Possible reasons are suggested in the message text.

User response:
Use the **mmlsfs** and **mmlsdisk** commands to determine what action is needed to correct the problem.

6027-1968 Failed while processing disk *diskName*.

Explanation:
An error was detected while processing the specified disk.

User response:
Examine prior messages to determine the reason for the failure. Correct the problem and reissue the command.

6027-1969 Device *device* already exists on node *nodeName*

Explanation:
This device already exists on the specified node.

User response:
None.

6027-1970 Disk *diskName* has no space for the quorum data structures. Specify a different disk as tiebreaker disk.

Explanation:
There is not enough free space in the file system descriptor for the tiebreaker disk data structures.

User response:
Specify a different disk as a tiebreaker disk.

6027-1971 Disk *lvName* (pvid *pvid*) is not known on node *nodeName*

Explanation:
The specified disk is not known on the above node.

User response:
Check the disk and node names and reissue the command.

6027-1972 Import of volume group *vgName* on node *nodeName* failed.

Explanation:
Import of the specified volume group on the specified node failed.

User response:
Check for additional error messages. Check both names and reissue the command.

6027-1973 Volume group *vgName* is not known on node *nodeName*.

Explanation:
The above volume group is not defined on the specified node.

User response:
Check both names and reissue the command.

6027-1974 None of the quorum nodes can be reached.

Explanation:
Ensure that the quorum nodes in the cluster can be reached. At least one of these nodes is required for the command to succeed.

User response:
Ensure that the quorum nodes are available and reissue the command.

6027-1975 The descriptor file contains more than one descriptor.

Explanation:
The descriptor file must contain only one descriptor.

User response:
Correct the descriptor file.

6027-1976 The descriptor file contains no descriptor.

Explanation:
The descriptor file must contain only one descriptor.

User response:
Correct the descriptor file.

6027-1977 Failed validating disk *diskName*. Error code *errorCode*.

Explanation:
GPFS control structures are not as expected.

User response:
Contact the IBM Support Center.

6027-1984 Name *name* is not allowed. It is longer than the maximum allowable length (*length*).

Explanation:
The cited name is not allowed because it is longer than the cited maximum allowable length.

User response:
Specify a name whose length does not exceed the maximum allowable length, and reissue the command.

6027-1985 mmfskxload: The format of the GPFS kernel extension is not correct for this version of AIX.

Explanation:
This version of AIX is incompatible with the current format of the GPFS kernel extension.

User response:
Contact your system administrator to check the AIX version and GPFS kernel extension.

6027-1986 *junctionName* does not resolve to a directory in *deviceName*. The junction must be within the specified file system.

Explanation:
The cited junction path name does not belong to the specified file system.

User response:
Correct the junction path name and reissue the command.

6027-1987 Name *name* is not allowed.

Explanation:
The cited name is not allowed because it is a reserved word or a prohibited character.

User response:
Specify a different name and reissue the command.

6027-1988 File system *fileSystem* is not mounted.

Explanation:
The cited file system is not currently mounted on this node.

User response:

Ensure that the file system is mounted and reissue the command.

6027-1991 Vpath disk *diskName* has an underlying hdisk that already belongs to a volume group.

Explanation:
The specified vpath disk cannot be used because one or more of its underlying hdisks already belongs to a volume group.

User response:
Remove the underlying hdisks from the volume group, or use a different vpath disk.

6027-1993 File *fileName* either does not exist or has an incorrect format.

Explanation:
The specified file does not exist or has an incorrect format.

User response:
Check whether the input file specified actually exists.

6027-1994 Did not find any match with the input disk address.

Explanation:
The **mmfileid** command returned without finding any disk addresses that match the given input.

User response:
None. Informational message only.

6027-1995 Device *deviceName* is not mounted on node *nodeName*.

Explanation:
The specified device is not mounted on the specified node.

User response:
Mount the specified device on the specified node and reissue the command.

6027-1996 Command was unable to determine whether file system *fileSystem* is mounted.

Explanation:
The command was unable to determine whether the cited file system is mounted.

User response:
Examine any prior error messages to determine why the command could not determine whether the file system was mounted, resolve the problem if possible, and then reissue the command. If you cannot resolve the problem, reissue the command with the daemon down on all nodes of the cluster. This will ensure that the file system is not mounted, which may allow the command to proceed.

6027-1997 Backup control file *fileName* from a previous backup does not exist.

Explanation:
The **mmbackup** command was asked to do an incremental or a resume backup, but the control file from a previous backup could not be found.

User response:
Restore the named file to the file system being backed up and reissue the command, or else do a full backup.

6027-1998 Line *lineNumber* of file *fileName* is incorrect:

Explanation:
A line in the specified file passed to the command had incorrect syntax. The line with the incorrect syntax is displayed next, followed by a description of the correct syntax for the line.

User response:
Correct the syntax of the line and reissue the command.

6027-1999 Syntax error. The correct syntax is: *string*.

Explanation:
The specified input passed to the command has incorrect syntax.

User response:
Correct the syntax and reissue the command.

6027-2000 Could not clear fencing for disk *physicalDiskName*.

Explanation:
The fencing information on the disk could not be cleared.

User response:
Make sure the disk is accessible by this node and retry.

6027-2002 Disk *physicalDiskName* of type *diskType* is not supported for fencing.

Explanation:
This disk is not a type that supports fencing.

User response:
None.

6027-2004 None of the specified nodes belong to this GPFS cluster.

Explanation:
The nodes specified do not belong to the GPFS cluster.

User response:
Choose nodes that belong to the cluster and try the command again.

6027-2007 **Unable to display fencing for disk** *physicalDiskName***.**

Explanation:
Cannot retrieve fencing information for this disk.

User response:
Make sure that this node has access to the disk before retrying.

6027-2008 **For the logical volume specification -l** *lvName* **to be valid** *lvName* **must be the only logical volume in the volume group. However, volume group** *vgName* **contains logical volumes.**

Explanation:
The command is being run on a logical volume that belongs to a volume group that has more than one logical volume.

User response:
Run this command only on a logical volume where it is the only logical volume in the corresponding volume group.

6027-2009 *logicalVolume* **is not a valid logical volume.**

Explanation:
logicalVolume does not exist in the ODM, implying that logical name does not exist.

User response:
Run the command on a valid logical volume.

6027-2010 *vgName* **is not a valid volume group name.**

Explanation:
vgName passed to the command is not found in the ODM, implying that *vgName* does not exist.

User response:
Run the command on a valid volume group name.

6027-2011 **For the hdisk specification -h** *physicalDiskName* **to be valid** *physicalDiskName* **must be the only disk in the volume group. However, volume group** *vgName* **contains disks.**

Explanation:
The hdisk specified belongs to a volume group that contains other disks.

User response:
Pass an hdisk that belongs to a volume group that contains only this disk.

6027-2012 *physicalDiskName* **is not a valid physical volume name.**

Explanation:
The specified name is not a valid physical disk name.

User response:
Choose a correct physical disk name and retry the command.

6027-2013 *pvid* **is not a valid physical volume id.**

Explanation:
The specified value is not a valid physical volume ID.

User response:
Choose a correct physical volume ID and retry the command.

6027-2014 **Node** *node* **does not have access to disk** *physicalDiskName***.**

Explanation:
The specified node is not able to access the specified disk.

User response:
Choose a different node or disk (or both), and retry the command. If both the node and disk name are correct, make sure that the node has access to the disk.

6027-2015 **Node** *node* **does not hold a reservation for disk** *physicalDiskName***.**

Explanation:
The node on which this command is run does not have access to the disk.

User response:
Run this command from another node that has access to the disk.

6027-2016 **SSA fencing support is not present on this node.**

Explanation:
This node does not support SSA fencing.

User response:
None.

6027-2017 **Node ID** *nodeId* **is not a valid SSA node ID. SSA node IDs must be a number in the range of 1 to 128.**

Explanation:
You specified a node ID outside of the acceptable range.

User response:
Choose a correct node ID and retry the command.

6027-2018 The SSA node id is not set.

Explanation:
The SSA node ID has not been set.

User response:
Set the SSA node ID.

6027-2019 Unable to retrieve the SSA node id.

Explanation:
A failure occurred while trying to retrieve the SSA node ID.

User response:
None.

6027-2020 Unable to set fencing for disk *physicalDiskName***.**

Explanation:
A failure occurred while trying to set fencing for the specified disk.

User response:
None.

6027-2021 Unable to clear PR reservations for disk *physicalDiskNam***.**

Explanation:
Failed to clear Persistent Reserve information on the disk.

User response:
Make sure the disk is accessible by this node before retrying.

6027-2022 Could not open disk *physicalDiskName***, errno** *value***.**

Explanation:
The specified disk cannot be opened.

User response:
Examine the **errno** value and other messages to determine the reason for the failure. Correct the problem and reissue the command.

6027-2023 retVal = *value***, errno =** *value* **for key** *value***.**

Explanation:
An **ioctl** call failed with stated return code, **errno** value, and related values.

User response:
Check the reported **errno** and correct the problem if possible. Otherwise, contact the IBM Support Center.

6027-2024 ioctl failed with rc=*returnCode***, errno=***errnoValue***. Related values are scsi_status=***scsiStatusValue***, sense_key=***senseKeyValue***, scsi_asc=***scsiAscValue***, scsi_ascq=***scsiAscqValue***.**

Explanation:
An **ioctl()** call failed with stated return code, **errno** value, and related values.

User response:
Check the reported **errno** and correct the problem if possible. Otherwise, contact the IBM Support Center.

6027-2025 READ_KEYS ioctl failed with errno=*returnCode***, tried** *timesTried* **times. Related values are scsi_status=***scsiStatusValue***, sense_key=***senseKeyValue***, scsi_asc=***scsiAscValue***, scsi_ascq=***scsiAscqValue***.**

Explanation:
A **READ_KEYS ioctl** call failed with stated **errno** value, and related values.

User response:
Check the reported **errno** and correct the problem if possible. Otherwise, contact the IBM Support Center.

6027-2026 READRES ioctl failed with errno=*returnCode***, tried** *timesTried* **times. Related values are: scsi_status=***scsiStatusValue***, sense_key=***senseKeyValue***, scsi_asc=***scsiAscValue***, scsi_ascq=***scsiAscqValue***.**

Explanation:
A **REGISTER ioctl** call failed with stated **errno** value, and related values.

User response:
Check the reported **errno** and correct the problem if possible. Otherwise, contact the IBM Support Center.

6027-2027 READRES ioctl failed with errno=*returnCode***, tried** *timesTried* **times. Related values are: scsi_status=***scsiStatusValue***, sense_key=***senseKeyValue***, scsi_asc=***scsiAscValue***, scsi_ascq=***scsiAscqValue***.**

Explanation:
A **READRES ioctl** call failed with stated **errno** value, and related values.

User response:
Check the reported **errno** and correct the problem if possible. Otherwise, contact the IBM Support Center.

6027-2028 Could not open disk device *diskDeviceName***.**

Explanation:
A problem occurred on a disk open.

User response:
Ensure that the disk is accessible and not fenced out, and then reissue the command.

6027-2029 Could not close disk device *diskDeviceName***.**

Explanation:
A problem occurred on a disk close.

User response:
None.

6027-2030 ioctl failed with DSB=*value* **and result=***value* **reason:** *explanation***.**

Explanation:
An **ioctl** call failed with stated return code, **errno** value, and related values.

User response:
Check the reported **errno** and correct the problem, if possible. Otherwise, contact the IBM Support Center.

6027-2031 ioctl failed with nonzero return code.

Explanation:
An **ioctl** failed with a nonzero return code.

User response:
Correct the problem, if possible. Otherwise, contact the IBM Support Center.

6027-2049 Cannot pin a page pool of size *value* **bytes.**

Explanation:
A GPFS page pool cannot be pinned into memory on this machine.

User response:
Increase the physical memory size of the machine.

6027-2050 Pagepool has size *actualValue* **bytes instead of the requested** *requestedValue* **bytes.**

Explanation:
The configured GPFS page pool is too large to be allocated or pinned into memory on this machine. GPFS will work properly, but with reduced capacity for caching user data.

User response:
To prevent this message from being generated when the GPFS daemon starts, reduce the page pool size using the **mmchconfig** command.

6027-2100 Incorrect range *value***-***value* **specified.**

Explanation:
The range specified to the command is incorrect. The first parameter value must be less than or equal to the second parameter value.

User response:
Correct the address range and reissue the command.

6027-2101 Insufficient free space in *fileSystem* **(***storage* **minimum required).**

Explanation:
There is not enough free space in the specified file system or directory for the command to successfully complete.

User response:
Correct the problem and reissue the command.

6027-2102 Node *nodeName* **is not mmremotefs to run the command.**

Explanation:
The specified node is not available to run a command. Depending on the command, a different node may be tried.

User response:
Determine why the specified node is not available and correct the problem.

6027-2103 Directory *dirName* **does not exist**

Explanation:
The specified directory does not exist.

User response:
Reissue the command specifying an existing directory.

6027-2104 The GPFS release level could not be determined on nodes: *nodeList***.**

Explanation:
The command was not able to determine the level of the installed GPFS code on the specified nodes.

User response:
Reissue the command after correcting the problem.

6027-2105 The following nodes must be upgraded to GPFS release *productVersion* **or higher:** *nodeList*

Explanation:
The command requires that all nodes be at the specified GPFS release level.

User response:
Correct the problem and reissue the command.

6027-2106 Ensure the nodes are available and run: *command***.**

Explanation:
The command could not complete normally.

User response:
Check the preceding messages, correct the problems, and issue the specified command until it completes successfully.

6027-2107 Upgrade the lower release level nodes and run: *command***.**

Explanation:
The command could not complete normally.

User response:
Check the preceding messages, correct the problems, and issue the specified command until it completes successfully.

6027-2108 The GPFS release level in effect for the cluster remains unchanged. Old level: *daemonVersion* **(Release** *productVersion***).**

Explanation:
Informational message.

User response:
Check the preceding messages, correct the problems, and reissue the command if necessary.

6027-2109 The GPFS release level in effect for the cluster cannot be determined. Assumed: *daemonVersion* **(Release** *productVersion***).**

Explanation:
Informational message.

User response:
Check the preceding messages, correct the problems, and reissue the command if necessary.

6027-2110 The GPFS release level in effect for the cluster will be changed. Old level: *daemonVersion* **(Release** *productVersion***). New level:** *daemonVersion* **(Release** *productVersion***).**

Explanation:
Informational message.

User response:
None.

6027-2111 The cluster contains nodes that are at different GPFS release levels. Lowest level: *daemonVersion* **(Release** *productVersion***). Highest level:** *daemonVersion* **(Release** *productVersion***).**

Explanation:
Informational message.

User response:
None.

6027-2112 -V option requires all nodes in the cluster to be at the same GPFS release level.

Explanation:
The **mmchfs** command either detected that not all nodes in the cluster are at the same GPFS release level, or was not able to confirm this.

User response:
Examine the preceding messages and correct any problems. Upgrade back-level nodes and run the **mmchconfig release=LATEST** command before rerunning the **mmchfs -V** command.

6027-2113 Not able to associate *diskName* **on node** *nodeName* **with any known GPFS disk.**

Explanation:
A command could not find a GPFS disk that matched the specified disk and node values passed as input.

User response:
Correct the disk and node values passed as input and reissue the command.

6027-2114 The *subsystem* **subsystem is already active.**

Explanation:
The user attempted to start a subsystem that was already active.

User response:
None. Informational message only.

6027-2115 Unable to resolve address range for disk *diskName* **on node** *nodeName***.**

Explanation:
A command could not perform address range resolution for the specified disk and node values passed as input.

User response:
Correct the disk and node values passed as input and reissue the command.

6027-2500 mmsanrepairfs already in progress for *name*.

Explanation:
This is an output from **mmsanrepairfs** when another **mmsanrepairfs** command is already running.

User response:
Wait for the currently running command to complete and reissue the command.

6027-2501 Could not allocate storage.

Explanation:
Sufficient memory could not be allocated to run the **mmsanrepairfs** command.

User response:
Increase the amount of memory available.

6027-2502 Error in repairing inode *number*, error *number*.

Explanation:
This is an output from **mmsanrepairfs** when an error is found while scanning the inode file to repair a specific inode.

User response:
Run **mmfsck** to fix the inode and rerun **mmsanrepairfs**.

6027-2600 Cannot create a new snapshot until an existing one is deleted. File system *fileSystem* has a limit of *number* online snapshots.

Explanation:
The file system has reached its limit of online snapshots

User response:
Delete an existing snapshot, then issue the create snapshot command again.

6027-2576 Error: Daemon *value* kernel *value* PAGE_SIZE mismatch.

Explanation:
The GPFS kernel extension loaded in memory does not have the same **PAGE_SIZE** value as the GPFS daemon **PAGE_SIZE** value that was returned from the POSIX **sysconf** API.

User response:
Verify that the kernel header files used to build the GPFS portability layer are the same kernel header files used to build the running kernel.

6027-2601 Snapshot directory *dirName* already exists.

Explanation:
This message is issued by the **tscrsnapshot** command.

User response:
Delete an existing file or directory and reissue the command.

6027-2602 Unable to delete snapshot *snapshotName* from file system *fileSystem*. rc=*returnCode*.

Explanation:
This message is issued by the **tscrsnapshot** command.

User response:
Delete the snapshot using the **tsdelsnapshot** command.

6027-2603 Unable to get permission to create snapshot, rc=*returnCode*.

Explanation:
This message is issued by the **tscrsnapshot** command.

User response:
Reissue the command.

6027-2604 Unable to quiesce all nodes, rc=*returnCode*.

Explanation:
This message is issued by the **tscrsnapshot** command.

User response:
Restart failing nodes or switches and reissue the command.

6027-2605 Unable to resume all nodes, rc=*returnCode*.

Explanation:
This message is issued by the **tscrsnapshot** command.

User response:
Restart failing nodes or switches.

6027-2606 Unable to sync all nodes, rc=*returnCode*.

Explanation:
This message is issued by the **tscrsnapshot** command.

User response:
Restart failing nodes or switches and reissue the command.

6027-2607 Cannot create new snapshot until an existing one is deleted. Fileset *filesetName* has a limit of *number* snapshots.

Explanation:
The fileset has reached its limit of snapshots.

User response:
None.

6027-2609 Fileset named *filesetName* does not exist.

Explanation:
One of the filesets listed does not exist.

User response:
Specify only existing fileset names.

6027-2610 File system *fileSystem* does not contain snapshot *snapshotName* err = *number*.

Explanation:
An incorrect snapshot name was specified.

User response:
Select a valid snapshot and issue the command again.

6027-2611 Cannot delete snapshot *snapshotName* which is *snapshotState* err = *number*.

Explanation:
The snapshot cannot be deleted because it is in the cited transition state due to an in-progress snapshot operation.

User response:
Wait for the in-progress operation to complete and then reissue the snapshot command.

6027-2612 Snapshot named *snapshotName* does not exist.

Explanation:
A snapshot to be listed does not exist.

User response:
Specify only existing snapshot names.

6027-2613 Cannot restore snapshot. *fileSystem* is mounted on *number* node(s) and in use on *number* node(s).

Explanation:
This message is issued by the **tsressnapshot** command.

User response:
Unmount the file system and reissue the restore command.

6027-2614 File system *fileSystem* does not contain snapshot *snapshotName* err = *number*.

Explanation:
An incorrect snapshot name was specified.

User response:
Specify a valid snapshot and issue the command again.

6027-2615 Cannot restore snapshot *snapshotName* which is *snapshotState* err = *number*.

Explanation:
The snapshot specified is not in a valid state.

User response:
Specify a snapshot that is in a valid state and issue the command again.

6027-2616 Restoring snapshot *snapshotName* requires *quotaTypes* quotas to be enabled.

Explanation:
The snapshot being restored requires quotas to be enabled, since they were enabled when the snapshot was created.

User response:
Issue the recommended **mmchfs** command to enable quotas.

6027-2617 You must run: mmchfs *fileSystem* -Q yes.

Explanation:
The snapshot being restored requires quotas to be enabled, since they were enabled when the snapshot was created.

User response:
Issue the cited **mmchfs** command to enable quotas.

6027-2618 Restoring snapshot *snapshotName* in file system *fileSystem* requires *quotaTypes* quotas to be enabled.

Explanation:
The snapshot being restored in the cited file system requires quotas to be enabled, since they were enabled when the snapshot was created.

User response:
Issue the **mmchfs** command to enable quotas.

6027-2619 Restoring snapshot *snapshotName* requires *quotaTypes* quotas to be disabled.

Explanation:
The snapshot being restored requires quotas to be disabled, since they were not enabled when the snapshot was created.

User response:
Issue the cited **mmchfs** command to disable quotas.

6027-2620 **You must run: mmchfs** *fileSystem* **-Q no.**

Explanation:
The snapshot being restored requires quotas to be disabled, since they were not enabled when the snapshot was created.

User response:
Issue the cited **mmchfs** command to disable quotas.

6027-2621 **Restoring snapshot** *snapshotName* **in file system** *fileSystem* **requires** *quotaTypes* **quotas to be disabled.**

Explanation:
The snapshot being restored in the cited file system requires quotas to be disabled, since they were disabled when the snapshot was created.

User response:
Issue the **mmchfs** command to disable quotas.

6027-2622 **Error restoring inode** *number* **err** *number***.**

Explanation:
The online snapshot was corrupted.

User response:
Restore the file from an offline snapshot.

6027-2623 **Error deleting snapshot** *fileSystem* **in file system** *fileSystem* **err** *number***.**

Explanation:
The cited snapshot could not be deleted during file system recovery.

User response:
Run the **mmfsck** command to recover any lost data blocks.

6027-2624 **Previous snapshot** *snapshotName* **is not valid and must be deleted before a new snapshot may be created.**

Explanation:
The cited previous snapshot is not valid and must be deleted before a new snapshot may be created.

User response:
Delete the previous snapshot using the **mmdelsnapshot** command, and then reissue the original snapshot command.

6027-2625 **Previous snapshot** *snapshotName* **must be restored before a new snapshot may be created.**

Explanation:
The cited previous snapshot must be restored before a new snapshot may be created.

User response:
Run **mmrestorefs** on the previous snapshot, and then reissue the original snapshot command.

6027-2626 **Previous snapshot** *snapshotName* **is not valid and must be deleted before another snapshot may be deleted.**

Explanation:
The cited previous snapshot is not valid and must be deleted before another snapshot may be deleted.

User response:
Delete the previous snapshot using the **mmdelsnapshot** command, and then reissue the original snapshot command.

6027-2627 **Previous snapshot** *snapshotName* **is not valid and must be deleted before another snapshot may be restored.**

Explanation:
The cited previous snapshot is not valid and must be deleted before another snapshot may be restored.

User response:
Delete the previous snapshot using the **mmdelsnapshot** command, and then reissue the original snapshot command.

6027-2628 **More than one snapshot is marked for restore.**

Explanation:
More than one snapshot is marked for restore.

User response:
Restore the previous snapshot and then reissue the original snapshot command.

6027-2629 **Offline snapshot being restored.**

Explanation:
An offline snapshot is being restored.

User response:
When the restore of the offline snapshot completes, reissue the original snapshot command.

6027-2630 *Program* **failed, error** *number***.**

Explanation:
The **tssnaplatest** command encountered an error and **printErrnoMsg** failed.

User response:
Correct the problem shown and reissue the command.

6027-2631 Attention: Snapshot *snapshotName* was being restored to *fileSystem*.

Explanation:
A file system in the process of a snapshot restore cannot be mounted except under a restricted mount.

User response:
None. Informational message only.

6027-2632 Mount of *fileSystem* failed: snapshot *snapshotName* must be restored before it can be mounted.

Explanation:
A file system in the process of a snapshot restore cannot be mounted for read only or read/write access.

User response:
Run the **mmrestorefs** command to complete the restoration, then reissue the **mount** command.

6027-2633 Attention: Disk configuration for *fileSystem* has changed while tsdf was running.

Explanation:
The disk configuration for the cited file system changed while the **tsdf** command was running.

User response:
Reissue the **mmdf** command.

6027-2634 Attention: *number* of *number* regions in *fileSystem* were unavailable for free space.

Explanation:
Some regions could not be accessed during the **tsdf** run. Typically, this is due to utilities such **mmdefragfs** or **mmfsck** running concurrently.

User response:
Reissue the **mmdf** command.

6027-2635 The free space data is not available. Reissue the command without the -q option to collect it.

Explanation:
The existing free space information for the file system is currently unavailable.

User response:
Reissue the **mmdf** command.

6027-2636 Disks in storage pool *storagePool* must have disk usage type 'dataOnly'.

Explanation:
A non-system storage pool cannot hold metadata or descriptors.

User response:
Modify the command's disk descriptors and reissue the command.

6027-2637 The file system must contain at least one disk for metadata.

Explanation:
The disk descriptors for this command must include one and only one storage pool that is allowed to contain metadata.

User response:
Modify the command's disk descriptors and reissue the command.

6027-2638 Maximum of *number* storage pools allowed.

Explanation:
The cited limit on the number of storage pools that may be defined has been exceeded.

User response:
Modify the command's disk descriptors and reissue the command.

6027-2639 Incorrect fileset name *filesetName*.

Explanation:
The fileset name provided in the command invocation is incorrect.

User response:
Correct the fileset name and reissue the command.

6027-2640 Incorrect path to fileset junction *filesetJunction*.

Explanation:
The path to the cited fileset junction is incorrect.

User response:
Correct the junction path and reissue the command.

6027-2641 Incorrect fileset junction name *filesetJunction*.

Explanation:
The cited junction name is incorrect.

User response:
Correct the junction name and reissue the command.

6027-2642 Specify one and only one of FilesetName or -J JunctionPath.

Explanation:
The change fileset and unlink fileset commands accept either a fileset name or the fileset's junction path to uniquely identify the fileset. The user failed to provide either of these, or has tried to provide both.

User response:
Correct the command invocation and reissue the command.

6027-2643 Cannot create a new fileset until an existing one is deleted. File system *fileSystem* has a limit of *maxNumber* filesets.

Explanation:
An attempt to create a fileset for the cited file system failed because it would exceed the cited limit.

User response:
Remove unneeded filesets and reissue the command.

6027-2644 Comment exceeds maximum length of *maxNumber* characters.

Explanation:
The user-provided comment for the new fileset exceeds the maximum allowed length.

User response:
Shorten the comment and reissue the command.

6027-2645 Fileset *filesetName* already exists.

Explanation:
An attempt to create a fileset failed because the specified fileset name already exists.

User response:
Select a unique name for the fileset and reissue the command.

6027-2647 Fileset *filesetName* must be unlinked to be deleted.

Explanation:
The cited fileset must be unlinked before it can be deleted.

User response:
Unlink the fileset, and then reissue the delete command.

6027-2648 Filesets have not been enabled for file system *fileSystem*.

Explanation:
The current file system format version does not support filesets, so fileset commands cannot be run.

User response:
Change the file system format version by issuing **mmchfs -V**.

Note: This is a request to upgrade the file system and activate functions that are incompatible with the previous release of GPFS.

6027-2649 Fileset *filesetName* contains user files, but can be deleted with the -f option.

Explanation:
An attempt was made to delete a non-empty fileset.

User response:
Remove all files and directories from the fileset, or specify the **-f** option to the delete fileset command.

6027-2650 Fileset information is not available.

Explanation:
A fileset command failed to read file system metadata file. The file system may be corrupted.

User response:
Run the **mmfsck** command to recover the file system.

6027-2651 Fileset *filesetName* cannot be unlinked.

Explanation:
The user tried to unlink the root fileset, or is not authorized to unlink the selected fileset.

User response:
None. The fileset cannot be unlinked.

6027-2652 Fileset at *junctionPath* cannot be unlinked.

Explanation:
The user tried to unlink the root fileset, or is not authorized to unlink the selected fileset.

User response:
None. The fileset cannot be unlinked.

6027-2653 Failed to unlink fileset *filesetName* from *filesetName*.

Explanation:
An attempt was made to unlink a fileset that is linked to a parent fileset that is being deleted.

User response:
Delete or unlink the children, and then delete the parent fileset.

6027-2654 Fileset *filesetName* cannot be deleted while other filesets are linked to it.

Explanation:
The fileset to be deleted has other filesets linked to it, and cannot be deleted without using the **-f** flag, or unlinking the child filesets.

User response:
Delete or unlink the children, and then delete the parent fileset.

6027-2655 Fileset *filesetName* cannot be deleted.

Explanation:
The user is not allowed to delete the root fileset.

User response:
None. The fileset cannot be deleted.

6027-2656 Unable to quiesce fileset at all nodes.

Explanation:
An attempt to quiesce the fileset at all nodes failed.

User response:
Check communication hardware and reissue the command.

6027-2657 Fileset *filesetName* has open files.

Explanation:
An attempt was made to unlink a fileset that has open files.

User response:
Close the open files and then reissue command, or use the **-f** option on the unlink command to force the open files to close.

6027-2658 Fileset *filesetName* cannot be linked into a snapshot at *pathName*.

Explanation:
The user specified a directory within a snapshot for the junction to a fileset, but snapshots cannot be modified.

User response:
Select a directory within the active file system, and reissue the command.

6027-2659 Fileset *filesetName* is already linked.

Explanation:
The user specified a fileset that was already linked.

User response:
Unlink the fileset and then reissue the link command.

6027-2660 Fileset *filesetName* cannot be linked.

Explanation:
The fileset could not be linked. This typically happens when the fileset is in the process of being deleted.

User response:
None.

6027-2661 Fileset junction *pathName* already exists.

Explanation:
A file or directory already exists at the specified junction.

User response:
Select a new junction name or a new directory for the link and reissue the link command.

6027-2662 Directory *pathName* for junction has too many links.

Explanation:
The directory specified for the junction has too many links.

User response:
Select a new directory for the link and reissue the command.

6027-2663 Fileset *filesetName* cannot be changed.

Explanation:
The user specified a fileset to **tschfileset** that cannot be changed.

User response:
None. You cannot change the attributes of the root fileset.

6027-2664 Fileset at *pathName* cannot be changed.

Explanation:
The user specified a fileset to **tschfileset** that cannot be changed.

User response:
None. You cannot change the attributes of the root fileset.

6027-2665 mmfileid already in progress for *name*.

Explanation:
An **mmfileid** command is already running.

User response:
Wait for the currently running command to complete, and issue the new command again.

6027-2666 mmfileid can only handle a maximum of *max disk addresses* disk addresses.

Explanation:
The **mmfileid** command can currently handle only 256 disk addresses as input.

User response:
Provide less than 256 disk addresses to the command.

6027-2667 Allowing block allocation that makes a file ill-replicated due to insufficient {*disk space* | *failure groups*} and puts data at risk.

Explanation:
The **partialReplicaAllocation** file system option allows allocation to succeed even when all replica blocks cannot be allocated. The file was marked as not replicated correctly and the data may be at risk if one of the remaining disks fails.

User response:
None. Informational message only.

6027-2670 Fileset name '*filesetName*' not found.

Explanation:
The fileset name that was specified with the command invocation was not found.

User response:
Correct the fileset name and reissue the command.

6027-2671 Fileset command on '*fileSystem*' failed, snapshot *snapshotName* must be restored first.

Explanation:
The file system is being restored either from an offline backup or a snapshot, and the restore operation has not finished. Fileset commands cannot be run.

User response:
Run the **mmrestorefs** command to complete the snapshot restore operation or to finish the offline restore, then reissue the fileset command.

6027-2672 Junction parent directory inode number *inodeNumber* is not valid.

Explanation:
An inode number passed to **tslinkfileset** is not valid.

User response:
Check the **mmlinkfileset** command arguments for correctness. If a valid junction path was provided, contact the IBM Support Center.

6027-2673 Duplicate owners of an allocation region (index *indexNumber*, region *regionNumber*, pool *poolNumber*) were detected for file system *fileSystem*: nodes *nodeName* and *nodeName*.

Explanation:
The allocation region should not have duplicate owners.

User response:
Contact the IBM Support Center.

6027-2674 The owner of an allocation region (index *indexNumber*, region *regionNumber*, pool *poolNumber*) that was detected for file system *fileSystem*: node *nodeName* is not valid.

Explanation:
The file system had detected a problem with the ownership of an allocation region. This may result in a corrupted file system and loss of data. One or more nodes may be terminated to prevent any further damage to the file system.

User response:
Unmount the file system and run the **kwdmmfsck** command to repair the file system.

6027-2675 Only file systems with NFSv4 ACL semantics enabled can be mounted on this platform.

Explanation:
A user is trying to mount a file system on Microsoft Windows, but the ACL semantics disallow NFSv4 ACLs.

User response:
Enable NFSv4 ACL semantics using the **mmchfs** command (**-k** option)

6027-2676 Only file systems with NFSv4 locking semantics enabled can be mounted on this platform.

Explanation:
A user is trying to mount a file system on Microsoft Windows, but the POSIX locking semantics are in effect.

User response:
Enable NFSv4 locking semantics using the **mmchfs** command (**-D** option)

6027-2700 A node join was rejected. This could be due to incompatible daemon versions, failure to find the node in the configuration database, or no configuration manager found.

Explanation:
A request to join nodes was explicitly rejected.

User response:
Verify that compatible versions of GPFS are installed on all nodes. Also, verify that the joining node is in the configuration database.

6027-2701 The mmpmon command file is empty.

Explanation:
The **mmpmon** command file is empty.

User response:
Check file size, existence, and access permissions.

6027-2702 Unexpected mmpmon response from file system daemon.

Explanation:
An unexpected response was received to an **mmpmon** request.

User response:
Ensure that the **mmfsd** daemon is running. Check the error log. Ensure that all GPFS software components are at the same version.

6027-2703 Unknown mmpmon command *command***.**

Explanation:
An unknown **mmpmon** command was read from the input file.

User response:
Correct the command and rerun.

6027-2704 Permission failure. The command requires root authority to execute.

Explanation:
The **mmpmon** command was issued with a nonzero UID.

User response:
Log on as root and reissue the command.

6027-2705 Could not establish connection to file system daemon.

Explanation:
The connection between a GPFS command and the **mmfsd** daemon could not be established. The daemon may have crashed, or never been started, or (for **mmpmon**) the allowed number of simultaneous connections has been exceeded.

User response:
Ensure that the **mmfsd** daemon is running. Check the error log. For **mmpmon**, ensure that the allowed number of simultaneous connections has not been exceeded.

6027-2706 Recovered *number* **nodes.**

Explanation:
The asynchronous part (phase 2) of node failure recovery has completed.

User response:
None. Informational message only.

6027-2707 Node join protocol waiting *value* **seconds for node recovery.**

Explanation:
Node join protocol is delayed until phase 2 of previous node failure recovery protocol is complete.

User response:
None. Informational message only.

6027-2708 Rejected node join protocol. Phase two of node failure recovery appears to still be in progress.

Explanation:
Node join protocol is rejected after a number of internal delays and phase two node failure protocol is still in progress.

User response:
None. Informational message only.

6027-2709 Configuration manager node *nodeName* **not found in the node list.**

Explanation:
The specified node was not found in the node list.

User response:
Add the specified node to the node list and reissue the command.

6027-2710 Node *nodeName* **is being expelled due to expired lease.**

Explanation:
The nodes listed did not renew their lease in a timely fashion and will be expelled from the cluster.

User response:
Check the network connection between this node and the node specified above.

6027-2711 File system table full.

Explanation:
The **mmfsd** daemon cannot add any more file systems to the table because it is full.

User response:
None. Informational message only.

6027-2712 Option '*optionName***' has been deprecated.**

Explanation:
The option that was specified with the command is no longer supported. A warning message is generated to indicate that the option has no effect.

User response:
Correct the command line and then reissue the command.

6027-2713 Permission failure. The command requires *SuperuserName* **authority to execute.**

Explanation:

The command, or the specified command option, requires administrative authority.

User response:
Log on as a user with administrative privileges and reissue the command.

6027-2714 Could not appoint node *nodeName* as cluster manager.

Explanation:
The **mmchmgr -c** command generates this message if the specified node cannot be appointed as a new cluster manager.

User response:
Make sure that the specified node is a quorum node and that GPFS is running on that node.

6027-2715 Could not appoint a new cluster manager.

Explanation:
The **mmchmgr -c** command generates this message when a node is not available as a cluster manager.

User response:
Make sure that GPFS is running on a sufficient number of quorum nodes.

6027-2716 Challenge response received; canceling disk election.

Explanation:
The node has challenged another node, which won the previous election, and detected a response to the challenge.

User response:
None. Informational message only.

6027-2717 Node *nodeName* is already a cluster manager or another node is taking over as the cluster manager.

Explanation:
The **mmchmgr -c** command generates this message if the specified node is already the cluster manager.

User response:
None. Informational message only.

6027-2718 Incorrect port range: GPFSCMDPORTRANGE='*range*'. Using default.

Explanation:
The GPFS command port range format is *lllll*[*-hhhhh*], where *lllll* is the low port value and *hhhhh* is the high port value. The valid range is 1 to 65535.

User response:
None. Informational message only.

6027-2719 The files provided do not contain valid quota entries.

Explanation:
The quota file provided does not have valid quota entries.

User response:
Check that the file being restored is a valid GPFS quota file.

6027-2722 Node limit of *number* has been reached. Ignoring *nodeName*.

Explanation:
The number of nodes that have been added to the cluster is greater than some cluster members can handle.

User response:
Delete some nodes from the cluster using the **mmdelnode** command, or shut down GPFS on nodes that are running older versions of the code with lower limits.

6027-2723 This node (*nodeName*) is now cluster manager for *clusterName*.

Explanation:
This is an informational message when a new cluster manager takes over.

User response:
None. Informational message only.

6027-2724 *reasonString* Probing cluster *clusterName*.

Explanation:
This is an informational message when a lease request has not been renewed.

User response:
None. Informational message only.

6027-2725 Node *nodeName* lease renewal is overdue. Pinging to check if it is alive.

Explanation:
This is an informational message on the cluster manager when a lease request has not been renewed.

User response:
None. Informational message only.

6027-2726 Recovered *number* nodes for file system *fileSystem*.

Explanation:
The asynchronous part (phase 2) of node failure recovery has completed.

User response:

| None. Informational message only.

6027-2800 Available memory exceeded on request to allocate *number* bytes. Trace point *sourceFile-tracePoint*.

Explanation:
The available memory was exceeded during an allocation request made from the cited source file and trace point.

User response:
Try shutting down and then restarting GPFS. If the problem recurs, contact the IBM Support Center.

6027-2801 Policy set syntax version *versionString* not supported.

Explanation:
The policy rules do not comply with the supported syntax.

User response:
Rewrite the policy rules, following the documented, supported syntax and keywords.

6027-2802 Object name '*poolName_or_filesetName*' is not valid.

Explanation:
The cited name is not a valid GPFS object, names an object that is not valid in this context, or names an object that no longer exists.

User response:
Correct the input to identify a GPFS object that exists and is valid in this context.

6027-2803 Policy set must start with VERSION.

Explanation:
The policy set does not begin with **VERSION** as required.

User response:
Rewrite the policy rules, following the documented, supported syntax and keywords.

6027-2804 Unexpected SQL result code - *sqlResultCode*.

Explanation:
This could be an IBM programming error.

User response:
Check that your SQL expressions are correct and supported by the current release of GPFS. If the error recurs, contact the IBM Support Center.

6027-2805 Loaded policy '*policyFileName or stripeGroupName*': *summaryOfPolicyRules*.

Explanation:
The cited loaded policy has the cited policy rules.

User response:
None. Informational message only.

6027-2806 Error while validating policy '*policyFileName or stripeGroupName*': rc=*errorCode*: *errorDetailsString*.

Explanation:
An error occurred while validating the cited policy.

User response:
Correct the policy rules, heeding the error details in this message and other messages issued immediately before or after this message. Use the **mmchpolicy** command to install a corrected policy rules file.

6027-2807 Error in evaluation of placement policy for file *fileName*: *errorDetailsString*.

Explanation:
An error occurred while evaluating the installed placement policy for a particular new file. Although the policy rules appeared to be syntactically correct when the policy was installed, evidently there is a problem when certain values of file attributes occur at runtime.

User response:
Determine which file names and attributes trigger this error. Correct the policy rules, heeding the error details in this message and other messages issued immediately before or after this message. Use the **mmchpolicy** command to install a corrected policy rules file.

6027-2808 In rule '*ruleName*' (*ruleNumber*), '*wouldBePoolName*' is not a valid pool name.

Explanation:
The cited name that appeared in the cited rule is not a valid pool name. This may be because the cited name was misspelled or removed from the file system.

User response:
Correct or remove the rule.

6027-2809 Validated policy '*policyFileName or stripeGroupName*': *summaryOfPolicyRules*.

Explanation:
The cited validated policy has the cited policy rules.

User response:
None. Informational message only.

6027-2810 There are *numberOfPools* storage pools but the policy file is missing or empty.

Explanation:
The cited number of storage pools are defined, but the policy file is missing or empty.

User response:
You should probably install a policy with placement rules using the **mmchpolicy** command, so that at least some of your data will be stored in your nonsystem storage pools.

6027-2811 Policy has no storage pool placement rules!

Explanation:
The policy has no storage pool placement rules.

User response:
You should probably install a policy with placement rules using the **mmchpolicy** command, so that at least some of your data will be stored in your nonsystem storage pools.

6027-2812 Keyword '*keywordValue*' begins a second *clauseName* clause - only one is allowed.

Explanation:
The policy rule should only have one clause of the indicated type.

User response:
Correct the rule and reissue the policy command.

6027-2813 This '*ruleName*' rule is missing a *clauseType* required clause.

Explanation:
The policy rule must have a clause of the indicated type.

User response:
Correct the rule and reissue the policy command.

6027-2814 This '*ruleName*' rule is of unknown type or not supported.

Explanation:
The policy rule set seems to have a rule of an unknown type or a rule that is unsupported by the current release of GPFS.

User response:
Correct the rule and reissue the policy command.

6027-2815 The value '*value*' is not supported in a '*clauseType*' clause.

Explanation:
The policy rule clause seems to specify an unsupported argument or value that is not supported by the current release of GPFS.

User response:
Correct the rule and reissue the policy command.

6027-2816 Policy rules employ features that would require a file system upgrade.

Explanation:
One or more policy rules have been written to use new features that cannot be installed on a back-level file system.

User response:
Install the latest GPFS software on all nodes and upgrade the file system or change your rules. (Note that **LIMIT** was introduced in GPFS Release 3.2.)

6027-2817 Error on popen/pclose (*command_string*): rc=*return_code_from_popen_or_pclose*

Explanation:
The execution of the *command_string* by **popen**/**pclose** resulted in an error.

User response:
To correct the error, do one or more of the following:

Check that the standard **m4** macro processing command is installed on your system as **/usr/bin/m4**.

Or:

Set the **MM_M4_CMD** environment variable.

Or:

Correct the macro definitions in your policy rules file.

If the problem persists, contact the IBM Support Center.

6027-2818 A problem occurred during m4 processing of policy rules. rc = *return_code_from_popen_pclose_or_m4*

Explanation:
An attempt to expand the policy rules with an **m4** subprocess yielded some warnings or errors or the **m4** macro wrote some output to standard error. Details or related messages may follow this message.

User response:
To correct the error, do one or more of the following:

Check that the standard **m4** macro processing command is installed on your system as **/usr/bin/m4**.

Or:

Set the **MM_M4_CMD** environment variable.

Or:

Correct the macro definitions in your policy rules file.

If the problem persists, contact the IBM Support Center.

6027-2819 **Error opening temp file** *temp_file_name*: *errorString*

Explanation:
An error occurred while attempting to open the specified temporary work file.

User response:
Check that the path name is defined and accessible. Check the file and then reissue the command.

6027-2820 **Error reading temp file** *temp_file_name*: *errorString*

Explanation:
An error occurred while attempting to read the specified temporary work file.

User response:
Check that the path name is defined and accessible. Check the file and then reissue the command.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any of IBM's intellectual property rights may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
USA

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Intellectual Property Law
Mail Station P300

2455 South Road,
Poughkeepsie, NY 12601-5400
USA

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment or a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to the application programming interfaces for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at ″Copyright and trademark information″ at www.ibm.com/legal/copytrade.shtml

Intel®, Intel Inside® (logos), MMX and Pentium® are trademarks of Intel Corporation in the United States, other countries, or both.

Java™ and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Red Hat, the Red Hat “Shadow Man” logo, and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc., in the United States and other countries.

UNIX is a registered trademark of the Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Glossary

This glossary defines technical terms and abbreviations used in GPFS documentation. If you do not find the term you are looking for, refer to the index of the appropriate book or view the IBM Glossary of Computing Terms, located on the Internet at: http://www-306.ibm.com/software/globalization/terminology/index.jsp.

B

block utilization. The measurement of the percentage of used subblocks per allocated blocks.

C

cluster. A loosely-coupled collection of independent systems (nodes) organized into a network for the purpose of sharing resources and communicating with each other. See also *GPFS cluster*.

cluster configuration data. The configuration data that is stored on the cluster configuration servers.

cluster manager. The node that monitors node status using disk leases, detects failures, drives recovery, and selects file system managers. The cluster manager is the node with the lowest node number among the quorum nodes that are operating at a particular time.

control data structures. Data structures needed to manage file data and metadata cached in memory. Control data structures include hash tables and link pointers for finding cached data; lock states and tokens to implement distributed locking; and various flags and sequence numbers to keep track of updates to the cached data.

D

Data Management Application Program Interface (DMAPI). The interface defined by the Open Group's XDSM standard as described in the publication *System Management: Data Storage Management (XDSM) API Common Application Environment (CAE) Specification C429*, The Open Group ISBN 1-85912-190-X.

deadman switch timer. A kernel timer that works on a node that has lost its disk lease and has outstanding I/O requests. This timer ensures that the node cannot complete the outstanding I/O requests (which would risk causing file system corruption), by causing a panic in the kernel.

disk descriptor. A definition of the type of data that the disk contains and the failure group to which this disk belongs. See also *failure group*.

disposition. The session to which a data management event is delivered. An individual disposition is set for each type of event from each file system.

disk leasing. A method for controlling access to storage devices from multiple host systems. Any host that wants to access a storage device configured to use disk leasing registers for a lease; in the event of a perceived failure, a host system can deny access, preventing I/O operations with the storage device until the preempted system has reregistered.

domain. A logical grouping of resources in a network for the purpose of common management and administration.

F

failback. Cluster recovery from failover following repair. See also *failover*.

failover. (1) The process of transferring all control of the ESS to a single cluster in the ESS when the other cluster in the ESS fails. See also *cluster*. (2) The routing of all transactions to a second controller when the first controller fails. See also *cluster*. (3) The assumption of file system duties by another node when a node fails.

failure group. A collection of disks that share common access paths or adapter connection, and could all become unavailable through a single hardware failure.

fileset. A hierarchical grouping of files managed as a unit for balancing workload across a cluster.

file-management policy. A set of rules defined in a policy file that GPFS uses to manage file migration and file deletion. See also *policy*.

file-placement policy. A set of rules defined in a policy file that GPFS uses to manage the initial placement of a newly created file. See also *policy*.

file system descriptor. A data structure containing key information about a file system. This information includes the disks assigned to the file system (*stripe group*), the current state of the file system, and pointers to key files such as quota files and log files.

file system descriptor quorum. The number of disks needed in order to write the file system descriptor correctly.

file system manager. The provider of services for all the nodes using a single file system. A file system manager processes changes to the state or description of the file system, controls the regions of disks that are allocated to each node, and controls token management and quota management.

fragment. The space allocated for an amount of data too small to require a full block. A fragment consists of one or more subblocks.

G

GPFS cluster. A cluster of nodes defined as being available for use by GPFS file systems.

GPFS portability layer. The interface module that each installation must build for its specific hardware platform and Linux distribution.

GPFS recovery log. A file that contains a record of metadata activity, and exists for each node of a cluster. In the event of a node failure, the recovery log for the failed node is replayed, restoring the file system to a consistent state and allowing other nodes to continue working.

I

ill-placed file. A file assigned to one storage pool, but having some or all of its data in a different storage pool.

ill-replicated file. A file with contents that are not correctly replicated according to the desired setting for that file. This situation occurs in the interval between a change in the file's replication settings or suspending one of its disks, and the restripe of the file.

indirect block. A block containing pointers to other blocks.

IBM Virtual Shared Disk. The subsystem that allows application programs running on different nodes to access a logical volume as if it were local to each node. The logical volume is local to only one of the nodes (the server node).

inode. The internal structure that describes the individual files in the file system. There is one inode for each file.

J

journaled file system (JFS). A technology designed for high-throughput server environments, which are important for running intranet and other high-performance e-business file servers.

junction.

A special directory entry that connects a name in a directory of one fileset to the root directory of another fileset.

K

kernel. The part of an operating system that contains programs for such tasks as input/output, management and control of hardware, and the scheduling of user tasks.

L

logical volume. A collection of physical partitions organized into logical partitions, all contained in a single volume group. Logical volumes are expandable and can span several physical volumes in a volume group.

Logical Volume Manager (LVM). A set of system commands, library routines, and other tools that allow the user to establish and control logical volume (LVOL) storage. The LVM maps data between the logical view of storage space and the physical disk drive module (DDM).

M

metadata. A data structures that contain access information about file data. These include: inodes, indirect blocks, and directories. These data structures are not accessible to user applications.

metanode. The one node per open file that is responsible for maintaining file metadata integrity. In most cases, the node that has had the file open for the longest period of continuous time is the metanode.

mirroring. The process of writing the same data to multiple disks at the same time. The mirroring of data protects it against data loss within the database or within the recovery log.

multi-tailed. A disk connected to multiple nodes.

N

namespace. Space reserved by a file system to contain the names of its objects.

Network File System (NFS). A protocol, developed by Sun Microsystems, Incorporated, that allows any host in a network to gain access to another host or netgroup and their file directories.

Network Shared Disk (NSD). A component for cluster-wide disk naming and access.

NSD volume ID. A unique 16 digit hex number that is used to identify and access all NSDs.

node. An individual operating-system image within a cluster. Depending on the way in which the computer system is partitioned, it may contain one or more nodes.

node descriptor. A definition that indicates how GPFS uses a node. Possible functions include: manager node, client node, quorum node, and nonquorum node

node number. A number that is generated and maintained by GPFS as the cluster is created, and as nodes are added to or deleted from the cluster.

node quorum. The minimum number of nodes that must be running in order for the daemon to start.

node quorum with tiebreaker disks. A form of quorum that allows GPFS to run with as little as one quorum node available, as long as there is access to a majority of the quorum disks.

non-quorum node. A node in a cluster that is not counted for the purposes of quorum determination.

P

policy. A list of file-placement and service-class rules that define characteristics and placement of files. Several policies can be defined within the configuration, but only one policy set is active at one time.

policy rule. A programming statement within a policy that defines a specific action to be preformed.

pool. A group of resources with similar characteristics and attributes.

portability. The ability of a programming language to compile successfully on different operating systems without requiring changes to the source code.

primary GPFS cluster configuration server. In a GPFS cluster, the node chosen to maintain the GPFS cluster configuration data.

private IP address. A IP address used to communicate on a private network.

public IP address. A IP address used to communicate on a public network.

Q

quorum node. A node in the cluster that is counted to determine whether a quorum exists.

quota. The amount of disk space and number of inodes assigned as upper limits for a specified user, group of users, or fileset.

quota management. The allocation of disk blocks to the other nodes writing to the file system, and comparison of the allocated space to quota limits at regular intervals.

R

Redundant Array of Independent Disks (RAID). A collection of two or more disk physical drives that present to the host an image of one or more logical disk drives. In the event of a single physical device failure, the data can be read or regenerated from the other disk drives in the array due to data redundancy.

recovery. The process of restoring access to file system data when a failure has occurred. Recovery can involve reconstructing data or providing alternative routing through a different server.

replication. The process of maintaining a defined set of data in more than one location. Replication involves copying designated changes for one location (a source) to another (a target), and synchronizing the data in both locations.

rule. A list of conditions and actions that are triggered when certain conditions are met. Conditions include attributes about an object (file name, type or extension, dates, owner, and groups), the requesting client, and the container name associated with the object.

S

SAN-attached. Disks that are physically attached to all nodes in the cluster using Serial Storage Architecture (SSA) connections or using fibre channel switches

secondary GPFS cluster configuration server. In a GPFS cluster, the node chosen to maintain the GPFS cluster configuration data in the event that the primary GPFS cluster configuration server fails or becomes unavailable.

Secure Hash Algorithm digest (SHA digest). A character string used to identify a GPFS security key.

Serial Storage Architecture (SSA). An American National Standards Institute (ANSI) standard, implemented by IBM, for a high-speed serial interface that provides point-to-point connection for peripherals, such as storage arrays.

session failure. The loss of all resources of a data management session due to the failure of the daemon on the session node.

session node. The node on which a data management session was created.

Small Computer System Interface (SCSI). An ANSI-standard electronic interface that allows personal computers to communicate with peripheral hardware, such as disk drives, tape drives, CD-ROM drives, printers, and scanners faster and more flexibly than previous interfaces.

snapshot. A copy of changed data in the active files and directories of a file system with the exception of the inode number, which is changed to allow application programs to distinguish between the snapshot and the active files and directories.

source node. The node on which a data management event is generated.

SSA. See Serial Storage Architecture.

stand-alone client. The node in a one-node cluster.

storage area network (SAN). A dedicated storage network tailored to a specific environment, combining servers, storage products, networking products, software, and services.

storage pool. A grouping of storage space consisting of volumes, logical unit numbers (LUNs), or addresses that share a common set of administrative characteristics.

stripe group. The set of disks comprising the storage assigned to a file system.

striping. A storage process in which information is split into blocks (a fixed amount of data) and the blocks are written to (or read from) a series of disks in parallel.

subblock. The smallest unit of data accessible in an I/O operation, equal to one thirty-second of a data block.

system storage pool. A storage pool containing file system control structures, reserved files, directories, symbolic links, special devices, as well as the metadata associated with regular files, including indirect blocks and extended attributes The **system storage pool** can also contain user data.

T

token management. A system for controlling file access in which each application performing a read or write operation is granted some form of access to a specific block of file data. Token management provides data consistency and controls conflicts. Token management has two components: the token management server, and the token management function.

token management function. A component of token management that requests tokens from the token management server. The token management function is located on each cluster node.

token management server. A component of token management that controls tokens relating to the operation of the file system. The token management server is located at the file system manager node.

twin-tailed. A disk connected to two nodes.

U

user storage pool. A storage pool containing the blocks of data that make up user files.

V

virtual file system (VFS). A remote file system that has been mounted so that it is accessible to the local user.

virtual shared disk. See *IBM Virtual Shared Disk*.

virtual node (vnode). The structure that contains information about a file system object in an virtual file system (VFS).

Accessibility features for GPFS

Accessibility features help users who have a disability, such as restricted mobility or limited vision, to use information technology products successfully.

Accessibility features

The following list includes the major accessibility features in GPFS:

- Keyboard-only operation
- Interfaces that are commonly used by screen readers
- Keys that are discernible by touch but do not activate just by touching them
- Industry-standard devices for ports and connectors
- The attachment of alternative input and output devices

The **IBM Cluster Information Center**, and its related publications, are accessibility-enabled. The accessibility features of the information center are described at http://publib.boulder.ibm.com/infocenter/clresctr/vxrx/index.jsp?topic=/com.ibm.cluster.addinfo.doc/access.html.

Keyboard navigation

This product uses standard Microsoft Windows navigation keys.

IBM and accessibility

See the IBM Human Ability and Accessibility Center for more information about the commitment that IBM has to accessibility:

`http://www.ibm.com/able`

Index

Special characters

A

C

D

E

F

G

H

I

J

K

L

M

N

O

P

Q

R

S

T

U

V

W

Reader's comments - We'd like to hear from you

General Parallel File System
Problem Determination Guide
Version 3 Release 2.1

Publication No. GA76-0415-02

We appreciate your comments about this publication. Please comment on specific errors or omissions, accuracy, organization, subject matter, or completeness of this book. The comments you send should pertain to only the information in this manual or product and the way in which the information is presented.

For technical questions and information about products and prices, please contact your IBM branch office, your IBM business partner, or your authorized remarketer.

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you. IBM or any other organizations will only use the personal information that you supply to contact you about the issues that you state on this form.

Comments:

Thank you for your support.

Submit your comments using one of these channels:

- Send your comments to the address on the reverse side of this form.
- Send your comments via e-mail to: mhvrcfs@us.ibm.com

If you would like a response from IBM, please fill in the following information:

Name	Address
Company or Organization	
Phone No.	E-mail address

Readers' Comments — We'd Like to Hear from You
GA76-0415-02

Cut or Fold
Along Line

Fold and Tape | **Please do not staple** | Fold and Tape

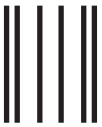

NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

IBM Corporation
Department 58HA, Mail Station P181
2455 South Road
Poughkeepsie, NY
12601-5400

Fold and Tape | **Please do not staple** | Fold and Tape

Cut or Fold
Along Line

Program Number: 5724-N94
5765-G66

GA76-0415-02