

IBM Spectrum Scale  
Version 4 Release 2.0

*Problem Determination Guide*





IBM Spectrum Scale  
Version 4 Release 2.0

*Problem Determination Guide*



**Note**

Before using this information and the product it supports, read the information in “Notices” on page 319.

This edition applies to version 4 release 2 of the following products, and to all subsequent releases and modifications until otherwise indicated in new editions:

- IBM Spectrum Scale ordered through Passport Advantage® (product number 5725-Q01)
- IBM Spectrum Scale ordered through AAS/eConfig (product number 5641-GPF)
- IBM Spectrum Scale for Linux on z Systems (product number 5725-S28)

Significant changes or additions to the text and illustrations are indicated by a vertical line (|) to the left of the change.

IBM welcomes your comments; see the topic “How to send your comments” on page xii. 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.

© Copyright IBM Corporation 2014, 2016.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

# Contents

<b>Tables</b> . . . . .	<b>vii</b>
-------------------------	------------

## **About this information** . . . . . **ix**

Prerequisite and related information . . . . .	xi
Conventions used in this information . . . . .	xi
How to send your comments . . . . .	xii

## **Summary of changes** . . . . . **xiii**

## **Chapter 1. Logs, dumps, and traces** . . . **1**

GPFS logs . . . . .	1
Creating a master GPFS log file . . . . .	2
Protocol services logs . . . . .	3
SMB logs . . . . .	3
NFS logs. . . . .	4
Object logs . . . . .	6
The IBM Spectrum Scale HDFS transparency log . . . . .	8
Protocol authentication log files . . . . .	9
CES monitoring and troubleshooting . . . . .	11
CES tracing and debug data collection . . . . .	13
The operating system error log facility . . . . .	19
MMFS_ABNORMAL_SHUTDOWN . . . . .	20
MMFS_DISKFAIL . . . . .	20
MMFS_ENVIRON . . . . .	20
MMFS_FSSTRUCT . . . . .	20
MMFS_GENERIC . . . . .	20
MMFS_LONGDISKIO . . . . .	21
MMFS_QUOTA . . . . .	21
MMFS_SYSTEM_UNMOUNT . . . . .	22
MMFS_SYSTEM_WARNING . . . . .	22
Error log entry example . . . . .	22
Using the gpfs.snap command . . . . .	23
Data gathered by gpfs.snap on all platforms . . . . .	23
Data gathered by gpfs.snap on AIX . . . . .	24
Data gathered by gpfs.snap on Linux. . . . .	25
Data gathered by gpfs.snap on Windows . . . . .	25
Data gathered by gpfs.snap for a master snapshot . . . . .	25
Data gathered by gpfs.snap on Linux for protocols . . . . .	26
mmdumpperfdata command . . . . .	31
mmfsadm command . . . . .	33
Trace facility . . . . .	34
Generating GPFS trace reports . . . . .	34
Best practices for setting up core dumps on a client system . . . . .	38

## **Chapter 2. Troubleshooting options available in GUI** . . . . . **41**

## **Chapter 3. GPFS cluster state information** . . . . . **43**

The <b>mmafmctl Device getstate</b> command . . . . .	43
The <b>mmdiag</b> command . . . . .	43
The <b>mmgetstate</b> command . . . . .	43

The <b>mmclscluster</b> command . . . . .	44
The <b>mmclsconfig</b> command . . . . .	45
The <b>mmrefresh</b> command . . . . .	45
The <b>mmrestore</b> command . . . . .	46
The <b>mmxpeinode</b> command . . . . .	46

## **Chapter 4. GPFS file system and disk information** . . . . . **49**

Restricted mode mount . . . . .	49
Read-only mode mount . . . . .	49
The <b>lsdf</b> command . . . . .	50
The <b>mmmount</b> command . . . . .	50
The <b>mmapplypolicy -L</b> command . . . . .	51
<b>mmapplypolicy -L 0</b> . . . . .	52
<b>mmapplypolicy -L 1</b> . . . . .	52
<b>mmapplypolicy -L 2</b> . . . . .	53
<b>mmapplypolicy -L 3</b> . . . . .	54
<b>mmapplypolicy -L 4</b> . . . . .	55
<b>mmapplypolicy -L 5</b> . . . . .	55
<b>mmapplypolicy -L 6</b> . . . . .	56
The <b>mmcheckquota</b> command . . . . .	57
The <b>mmclsnsd</b> command . . . . .	57
The <b>mmwindisk</b> command . . . . .	58
The <b>mmfileid</b> command . . . . .	59
The <b>SHA digest</b> . . . . .	61

## **Chapter 5. Resolving deadlocks** . . . . . **63**

Automated deadlock detection . . . . .	63
Automated deadlock data collection . . . . .	65
Automated deadlock breakup . . . . .	66
Deadlock breakup on demand . . . . .	67
Cluster overload detection . . . . .	68

## **Chapter 6. Other problem determination tools** . . . . . **71**

## **Chapter 7. Installation and configuration issues** . . . . . **73**

Installation and configuration problems . . . . .	73
What to do after a node of a GPFS cluster crashes and has been reinstalled . . . . .	74
Problems with the <b>/etc/hosts</b> file . . . . .	74
Linux configuration considerations . . . . .	74
Protocol authentication problem determination . . . . .	75
Problems with running commands on other nodes . . . . .	75
GPFS cluster configuration data files are locked . . . . .	76
Recovery from loss of GPFS cluster configuration data file . . . . .	77
Automatic backup of the GPFS cluster data. . . . .	78
Error numbers specific to GPFS applications calls . . . . .	78
GPFS modules cannot be loaded on Linux . . . . .	79
GPFS daemon will not come up . . . . .	79

Steps to follow if the GPFS daemon does not come up . . . . .	80
Unable to start GPFS after the installation of a new release of GPFS . . . . .	81
GPFS error messages for shared segment and network problems . . . . .	82
Error numbers specific to GPFS application calls when the daemon is unable to come up . . . . .	82
GPFS daemon went down . . . . .	83
IBM Spectrum Scale failures due to a network failure . . . . .	84
Kernel panics with a 'GPFS dead man switch timer has expired, and there's still outstanding I/O requests' message . . . . .	85
Quorum loss . . . . .	85
Delays and deadlocks . . . . .	86
Node cannot be added to the GPFS cluster . . . . .	87
Remote node expelled after remote file system successfully mounted . . . . .	87
Disaster recovery issues . . . . .	88
Disaster recovery setup problems . . . . .	88
Other problems with disaster recovery . . . . .	89
GPFS commands are unsuccessful . . . . .	89
GPFS error messages for unsuccessful GPFS commands. . . . .	90
Application program errors . . . . .	91
GPFS error messages for application program errors . . . . .	92
Troubleshooting Windows problems . . . . .	92
Home and .ssh directory ownership and permissions . . . . .	92
Problems running as Administrator . . . . .	92
GPFS Windows and SMB2 protocol (CIFS serving) . . . . .	93
OpenSSH connection delays . . . . .	93
File protocol authentication setup issues. . . . .	93

## Chapter 8. File system issues . . . . . 95

File system will not mount . . . . .	95
GPFS error messages for file system mount problems . . . . .	97
Error numbers specific to GPFS application calls when a file system mount is not successful . . . . .	98
Automount file system will not mount . . . . .	98
Remote file system will not mount . . . . .	100
Mount failure due to client nodes joining before NSD servers are online . . . . .	103
File system will not unmount . . . . .	104
File system forced unmount . . . . .	105
Additional failure group considerations . . . . .	106
GPFS error messages for file system forced unmount problems . . . . .	107
Error numbers specific to GPFS application calls when a file system has been forced to unmount . . . . .	107
Unable to determine whether a file system is mounted . . . . .	108
GPFS error messages for file system mount status . . . . .	108
Multiple file system manager failures . . . . .	108
GPFS error messages for multiple file system manager failures . . . . .	108

Error numbers specific to GPFS application calls when file system manager appointment fails . . . . .	109
Discrepancy between GPFS configuration data and the on-disk data for a file system. . . . .	109
Errors associated with storage pools, filesets and policies . . . . .	109
A NO_SPACE error occurs when a file system is known to have adequate free space . . . . .	110
Negative values occur in the 'predicted pool utilizations', when some files are 'ill-placed' . . . . .	111
Policies - usage errors . . . . .	111
Errors encountered with policies . . . . .	112
Filesets - usage errors. . . . .	113
Errors encountered with filesets . . . . .	114
Storage pools - usage errors . . . . .	114
Errors encountered with storage pools . . . . .	115
Failures using the mmbakup command . . . . .	116
GPFS error messages for mmbakup errors . . . . .	116
TSM error messages . . . . .	116
Snapshot problems . . . . .	116
Problems with locating a snapshot . . . . .	116
Problems not directly related to snapshots. . . . .	116
Snapshot usage errors . . . . .	117
Snapshot status errors . . . . .	117
Errors encountered when restoring a snapshot . . . . .	118
Snapshot directory name conflicts . . . . .	118
Failures using the mmpmon command . . . . .	119
Setup problems using mmpmon . . . . .	119
Incorrect output from mmpmon . . . . .	120
Abnormal termination or hang in mmpmon . . . . .	120
NFS issues . . . . .	121
NFS client with stale inode data . . . . .	121
NFS V4 problems . . . . .	121
Determining the health of integrated SMB server . . . . .	122
Problems working with Samba . . . . .	123
Data integrity . . . . .	124
Error numbers specific to GPFS application calls when data integrity may be corrupted . . . . .	124
Messages queueing in AFM . . . . .	124

## Chapter 9. Disk issues . . . . . 127

NSD and underlying disk subsystem failures. . . . .	127
Error encountered while creating and using NSD disks . . . . .	127
Displaying NSD information . . . . .	128
NSD creation fails with a message referring to an existing NSD . . . . .	130
GPFS has declared NSDs as down . . . . .	130
Unable to access disks . . . . .	131
Guarding against disk failures. . . . .	132
Disk media failure. . . . .	132
Disk connectivity failure and recovery . . . . .	135
Partial disk failure. . . . .	136
GPFS has declared NSDs built on top of AIX logical volumes as down . . . . .	136
Verify logical volumes are properly defined for GPFS use. . . . .	136
Check the volume group on each node. . . . .	137
Volume group varyon problems . . . . .	137
Disk accessing commands fail to complete due to problems with some non-IBM disks . . . . .	138

Persistent Reserve errors . . . . .	138
Understanding Persistent Reserve . . . . .	138
Checking Persistent Reserve . . . . .	139
Clearing a leftover Persistent Reserve reservation . . . . .	139
Manually enabling or disabling Persistent Reserve . . . . .	140
GPFS is not using the underlying multipath device	141

**Chapter 10. Encryption issues . . . . . 143**

Unable to add encryption policies . . . . .	143
Receiving “Permission denied” message . . . . .	143
“Value too large” failure when creating a file . . . . .	143
Mount failure for a file system with encryption rules . . . . .	143
“Permission denied” failure of key rewrap . . . . .	143

**Chapter 11. Other problem determination hints and tips . . . . . 145**

Which physical disk is associated with a logical volume? . . . . .	145
Which nodes in my cluster are quorum nodes? . . . . .	145
What is stored in the /tmp/mmfs directory and why does it sometimes disappear? . . . . .	146
Why does my system load increase significantly during the night? . . . . .	146
What do I do if I receive message 6027-648? . . . . .	147
Why can't I see my newly mounted Windows file system? . . . . .	147
Why is the file system mounted on the wrong drive letter? . . . . .	147
Why does the offline mmfscck command fail with "Error creating internal storage"? . . . . .	147
Why do I get timeout executing function error message? . . . . .	147

Questions related to active file management . . . . .	148
Questions related to File Placement Optimizer (FPO) . . . . .	148

**Chapter 12. Reliability, Availability, and Serviceability (RAS) events . . . . 151**

**Chapter 13. Contacting IBM support center . . . . . 167**

Information to be collected before contacting the IBM Support Center . . . . .	167
How to contact the IBM Support Center . . . . .	169

**Chapter 14. Message severity tags . . . . . 171**

**Chapter 15. Messages. . . . . 173**

**Accessibility features for IBM Spectrum Scale . . . . . 317**

Accessibility features . . . . .	317
Keyboard navigation . . . . .	317
IBM and accessibility . . . . .	317

**Notices . . . . . 319**

Trademarks . . . . .	321
Terms and conditions for product documentation	321
IBM Online Privacy Statement . . . . .	322

**Glossary . . . . . 323**

**Index . . . . . 329**





---

## Tables

1. IBM Spectrum Scale library information units . . . . .	ix	8. Events for the GPFS component . . . . .	153
2. Conventions . . . . .	xii	9. Events for the KEYSTONE component . . . . .	154
3. Core object log files in /var/log/swift . . . . .	7	10. Events for the NFS component . . . . .	155
4. Additional object log files in /var/log/swift . . . . .	8	11. Events for the Network component . . . . .	159
5. General system log files in /var/adm/ras . . . . .	8	12. Events for the Object component . . . . .	160
6. Authentication log files . . . . .	9	13. Events for the SMB component . . . . .	165
7. Events for the AUTH component . . . . .	151	14. Message severity tags ordered by priority . . . . .	171



---

## About this information

This edition applies to IBM Spectrum Scale™ version 4.2 for AIX®, Linux, and Windows.

IBM Spectrum Scale is a file management infrastructure, based on IBM® General Parallel File System (GPFS™) technology, that provides unmatched performance and reliability with scalable access to critical file data.

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

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

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

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

To find out which version of IBM Spectrum Scale is running on a particular Windows node, open the **Programs and Features** control panel. The IBM Spectrum Scale installed program name includes the version number.

## Which IBM Spectrum Scale information unit provides the information you need?

The IBM Spectrum Scale library consists of the information units listed in Table 1.

To use these information units effectively, you must be familiar with IBM Spectrum Scale and the AIX, Linux, or Windows operating system, or all of them, depending on which operating systems are in use at your installation. Where necessary, these information units provide some background information relating to AIX, Linux, or Windows; however, more commonly they refer to the appropriate operating system documentation.

**Note:** Throughout this documentation, the term “Linux” refers to all supported distributions of Linux, unless otherwise specified.

*Table 1. IBM Spectrum Scale library information units*

Information unit	Type of information	Intended users
<i>IBM Spectrum Scale: Administration and Programming Reference</i>	This information unit explains how to do the following: <ul style="list-style-type: none"><li>• Use the commands, programming interfaces, and user exits unique to GPFS</li><li>• Manage clusters, file systems, disks, and quotas</li><li>• Export a GPFS file system using the Network File System (NFS) protocol</li></ul>	System administrators or programmers of GPFS systems

Table 1. IBM Spectrum Scale library information units (continued)

Information unit	Type of information	Intended users
<i>IBM Spectrum Scale: Advanced Administration Guide</i>	<p>This information unit explains how to use the following advanced features of GPFS:</p> <ul style="list-style-type: none"> <li>• Accessing GPFS file systems from other GPFS clusters</li> <li>• Policy-based data management for GPFS</li> <li>• Creating and maintaining snapshots of GPFS file systems</li> <li>• Establishing disaster recovery for your GPFS cluster</li> <li>• Monitoring GPFS I/O performance with the <b>mmpmon</b> command</li> <li>• Miscellaneous advanced administration topics</li> </ul>	System administrators or programmers seeking to understand and use the advanced features of GPFS
<i>IBM Spectrum Scale: Concepts, Planning, and Installation Guide</i>	<p>This information unit provides information about the following topics:</p> <ul style="list-style-type: none"> <li>• Introducing GPFS</li> <li>• GPFS architecture</li> <li>• Planning concepts for GPFS</li> <li>• Installing GPFS</li> <li>• Migration, coexistence and compatibility</li> <li>• Applying maintenance</li> <li>• Configuration and tuning</li> <li>• Uninstalling GPFS</li> </ul>	System administrators, analysts, installers, planners, and programmers of GPFS clusters who are very experienced with the operating systems on which each GPFS cluster is based

Table 1. IBM Spectrum Scale library information units (continued)

Information unit	Type of information	Intended users
<i>IBM Spectrum Scale: Data Management API Guide</i>	<p>This information unit describes the Data Management Application Programming Interface (DMAPI) for GPFS.</p> <p>This implementation is based on The Open Group's System Management: Data Storage Management (XDMS) API Common Applications Environment (CAE) Specification C429, The Open Group, ISBN 1-85912-190-X specification. The implementation is compliant with the standard. Some optional features are not implemented.</p> <p>The XDMS DMAPI model is intended mainly for a single-node environment. Some of the key concepts, such as sessions, event delivery, and recovery, required enhancements for a multiple-node environment such as GPFS.</p> <p>Use this information if you intend to write application programs to do the following:</p> <ul style="list-style-type: none"> <li>• Monitor events associated with a GPFS file system or with an individual file</li> <li>• Manage and maintain GPFS file system data</li> </ul>	Application programmers who are experienced with GPFS systems and familiar with the terminology and concepts in the XDMS standard
<i>IBM Spectrum Scale: Problem Determination Guide</i>	This information unit contains explanations of GPFS error messages and explains how to handle problems you may encounter with GPFS.	System administrators of GPFS systems who are experienced with the subsystems used to manage disks and who are familiar with the concepts presented in the <i>IBM Spectrum Scale: Concepts, Planning, and Installation Guide</i>

## Prerequisite and related information

For updates to this information, see IBM Spectrum Scale in IBM Knowledge Center ([www.ibm.com/support/knowledgecenter/STXKQY/ibmspectrumscale\\_welcome.html](http://www.ibm.com/support/knowledgecenter/STXKQY/ibmspectrumscale_welcome.html)).

For the latest support information, see the IBM Spectrum Scale FAQ in IBM Knowledge Center ([www.ibm.com/support/knowledgecenter/STXKQY/gpfsclustersfaq.html](http://www.ibm.com/support/knowledgecenter/STXKQY/gpfsclustersfaq.html)).

## Conventions used in this information

Table 2 on page xii describes the typographic conventions used in this information. UNIX file name conventions are used throughout this information.

**Note: Users of IBM Spectrum Scale for Windows** must be aware that on Windows, UNIX-style file names need to be converted appropriately. For example, the GPFS cluster configuration data is stored in the `/var/mmfs/gen/mmsdrfs` file. On Windows, the UNIX namespace starts under the `%SystemDrive%\cygwin64` directory, so the GPFS cluster configuration data is stored in the `C:\cygwin64\var\mmfs\gen\mmsdrfs` file.

Table 2. Conventions

Convention	Usage
<b>bold</b>	<p><b>Bold</b> words or characters represent system elements that you must use literally, such as commands, flags, values, and selected menu options.</p> <p>Depending on the context, <b>bold</b> typeface sometimes represents path names, directories, or file names.</p>
<b><u>bold underlined</u></b>	<p><b><u>bold underlined</u></b> keywords are defaults. These take effect if you do not specify a different keyword.</p>
constant width	<p>Examples and information that the system displays appear in constant-width typeface.</p> <p>Depending on the context, constant-width typeface sometimes represents path names, directories, or file names.</p>
<i>italic</i>	<p><i>Italic</i> words or characters represent variable values that you must supply.</p> <p><i>Italics</i> are also used for information unit titles, for the first use of a glossary term, and for general emphasis in text.</p>
<key>	<p>Angle brackets (less-than and greater-than) enclose the name of a key on the keyboard. For example, &lt;Enter&gt; refers to the key on your terminal or workstation that is labeled with the word <i>Enter</i>.</p>
\	<p>In command examples, a backslash indicates that the command or coding example continues on the next line. For example:</p> <pre>mkcondition -r IBM.FileSystem -e "PercentTotUsed &gt; 90" \ -E "PercentTotUsed &lt; 85" -m p "FileSystem space used"</pre>
{item}	<p>Braces enclose a list from which you must choose an item in format and syntax descriptions.</p>
[item]	<p>Brackets enclose optional items in format and syntax descriptions.</p>
<Ctrl-x>	<p>The notation &lt;Ctrl-x&gt; indicates a control character sequence. For example, &lt;Ctrl-c&gt; means that you hold down the control key while pressing &lt;c&gt;.</p>
item...	<p>Ellipses indicate that you can repeat the preceding item one or more times.</p>
	<p>In <i>synopsis</i> statements, vertical lines separate a list of choices. In other words, a vertical line means <i>Or</i>.</p> <p>In the left margin of the document, vertical lines indicate technical changes to the information.</p>

## 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 IBM Spectrum Scale documentation, send your comments to the following e-mail address:

**[mhvrcfs@us.ibm.com](mailto:mhvrcfs@us.ibm.com)**

Include the publication title and order number, and, if applicable, the specific location of the information about which you have comments (for example, a page number or a table number).

To contact the IBM Spectrum Scale development organization, send your comments to the following e-mail address:

**[gpfs@us.ibm.com](mailto:gpfs@us.ibm.com)**

---

## Summary of changes

This topic summarizes changes to the IBM Spectrum Scale licensed program and the IBM Spectrum Scale library. 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 information.

### Summary of changes for IBM Spectrum Scale version 4 release 2 as updated, November 2015

Changes to this release of the IBM Spectrum Scale licensed program and the IBM Spectrum Scale library include the following:

#### Cluster Configuration Repository (CCR): Backup and restore

You can backup and restore a cluster that has Cluster Configuration Repository (CCR) enabled. In the **mmsdrbackup** user exit, the type of backup that is created depends on the configuration of the cluster. If the Cluster Configuration Repository (CCR) is enabled, then a CCR backup is created. Otherwise, a **mmsdrfs** backup is created. In the **mmsdrrestore** command, if the configuration file is a Cluster Configuration Repository (CCR) backup file, then you must specify the **-a** option. All the nodes in the cluster are restored.

#### Changes in IBM Spectrum Scale for object storage

##### Object capabilities

Object capabilities describe the object protocol features that are configured in the IBM Spectrum Scale cluster such as unified file and object access, multi-region object deployment, and S3 API emulation. For more information, see the following topics:

- *Object capabilities in IBM Spectrum Scale: Concepts, Planning, and Installation Guide*
- *Managing object capabilities in IBM Spectrum Scale: Administration and Programming Reference*

##### Storage policies for object storage

Storage policies enable segmenting of the object storage within a single cluster for various use cases. Currently, OpenStack Swift supports storage policies that allow you to define the replication settings and location of objects in a cluster. IBM Spectrum Scale enhances storage policies to add compression and unified file and object access functions for object storage. For more information, see the following topics:

- *Storage policies for object storage in IBM Spectrum Scale: Concepts, Planning, and Installation Guide*
- *Mapping of storage policies to filesets in IBM Spectrum Scale: Administration and Programming Reference*
- *Administering storage policies for object storage in IBM Spectrum Scale: Administration and Programming Reference*

##### Multi-region object deployment

The main purpose of the object protocol is to enable the upload and download of object data. When clients have a fast connection to the cluster, the network delay is minimal. However, when client access to object data is over a WAN or a high-latency network, the network can introduce an unacceptable delay and affect quality-of-service metrics. To improve that response time, you can create a replica of the data in a cluster closer to the clients using the active-active multi-region replication support in OpenStack Swift. Multi-region can also be used to distribute the object load over several clusters to reduce contention in the file system. For more information, see the following topics:

- *Overview of multi-region object deployment in IBM Spectrum Scale: Concepts, Planning, and Installation Guide*
- *Planning for multi-region object deployment in IBM Spectrum Scale: Concepts, Planning, and Installation Guide*
- *Enabling multi-region object deployment initially in IBM Spectrum Scale: Concepts, Planning, and Installation Guide*
- *Adding a region in a multi-region object deployment in IBM Spectrum Scale: Administration and Programming Reference*
- *Administering a multi-region object deployment environment in IBM Spectrum Scale: Administration and Programming Reference*

### **Unified file and object access**

Unified file and object access allows users to access the same data as an object and as a file. Data can be stored and retrieved through IBM Spectrum Scale for object storage or as files from POSIX, NFS, and SMB interfaces. For more information, see the following topics:

- *Unified file and object access overview in IBM Spectrum Scale: Concepts, Planning, and Installation Guide*
- *Planning for unified file and object access in IBM Spectrum Scale: Concepts, Planning, and Installation Guide*
- *Installing and using unified file and object access in IBM Spectrum Scale: Concepts, Planning, and Installation Guide*
- *Unified file and object access in IBM Spectrum Scale in IBM Spectrum Scale: Administration and Programming Reference*

### **S3 access control lists (ACLs) support**

IBM Spectrum Scale for object storage supports S3 access control lists (ACLs) on buckets and objects. For more information, see *Managing OpenStack access control lists using S3 API emulation in IBM Spectrum Scale: Administration and Programming Reference*.

### **Changes in IBM Spectrum Scale for Linux on z Systems™**

- Compression support
- AFM-based Async Disaster Recovery (AFM DR) support
- IBM Spectrum Protect™ Backup-Archive and Space Management client support
- Support for all editions:
  - Express®
  - Standard
  - Advanced (without encryption)

For more information about current requirements and limitations of IBM Spectrum Scale for Linux on z Systems, see Q2.25 of IBM Spectrum Scale FAQ.

### **Change in AFM-based Async Disaster Recovery (AFM DR)**

- Support for IBM Spectrum Scale for Linux on z Systems

### **File compression**

With file compression, you can reclaim some of the storage space occupied by infrequently accessed files. Run the **mmchattr** command or the **mmapplypolicy** command to identify and compress a few files or many files. Run file compression synchronously or defer it. If you defer it, you can run the **mmrestripefile** or **mmrestripefs** to complete the compression. You can decompress files with the same commands used to compress files. When a compressed file is read it is decompressed on the fly and remains compressed on disk. When a compressed file is overwritten, the parts of the file that overlap with the changed data are decompressed on disk synchronously in the granularity of ten data blocks. File compression in this release is designed to



be used only for compressing cold data or write-once objects and files. Compressing other types of data can result in performance degradation. File compression uses the zlib data compression library and favors saving space over speed.

### GUI servers

The IBM Spectrum Scale system provides a GUI that can be used for managing and monitoring the system. Any server that provides this GUI service is referred to as a GUI server. If you need GUI service in the system, designate at least two nodes as GUI servers in the cluster. A maximum of three nodes can be designated as GUI servers. For more information on installing IBM Spectrum Scale using the GUI, see *IBM Spectrum Scale: Concepts, Planning, and Installation Guide*.

### IBM Spectrum Scale management GUI

The management GUI helps to manage and monitor the IBM Spectrum Scale system. You can perform the following tasks through management GUI:

- Monitoring the performance of the system based on various aspects
- Monitoring system health
- Managing file systems
- Creating filesets and snapshots
- Managing Objects and NFS and SMB data exports
- Creating administrative users and defining roles for the users
- Creating object users and defining roles for them
- Defining default, user, group, and fileset quotas
- Monitoring the capacity details at various levels such as file system, pools, filesets, users, and user groups

### Hadoop Support for IBM Spectrum Scale

IBM Spectrum Scale has been extended to work seamlessly in the Hadoop ecosystem and is available through a feature called File Placement Optimizer (FPO). Storing your Hadoop data using FPO allows you to gain advanced functions and the high I/O performance required for many big data operations. FPO provides Hadoop compatibility extensions to replace HDFS in a Hadoop ecosystem, with no changes required to Hadoop applications.

You can deploy a IBM Spectrum Scale using FPO as a file system platform for big data analytics. The topics in this guide covers a variety of Hadoop deployment architectures, including IBM BigInsights®, Platform Symphony®, or with a Hadoop distribution from another vendor to work with IBM Spectrum Scale.

IBM Spectrum Scale offers two kinds of interfaces for Hadoop applications to access File System data. One is IBM Spectrum Scale connector, which aligns with Hadoop Compatible File System architecture and APIs. The other is HDFS protocol, which provides a HDFS compatible interfaces.

For more information, see the following sections in the *IBM Spectrum Scale: Advanced Administration Guide*:

- *Hadoop support for IBM Spectrum Scale*
- *Configuring FPO*
- *Hadoop connector*
- *HDFS protocol*

### IBM Spectrum Scale installation GUI

You can use the installation GUI to install the IBM Spectrum Scale system. For more information on how to launch the GUI installer, see the *Installing IBM Spectrum Scale using the graphical user interface (GUI)* section in *IBM Spectrum Scale: Concepts, Planning, and Installation Guide*.

### Performance Monitoring Tool using the Installation Kit

The usage statement and optional arguments have changed during the installation of the toolkit. The new usage statement with options is as follows:

```
spectrumscale config perfmon [-h] [-l] [-r {on,off}]
```

For more information, see *IBM Spectrum Scale: Concepts, Planning, and Installation Guide*.

### Protocols cluster disaster recovery (DR)

You can use the **mmcesdr** command to perform DR setup, failover, failback, backup, and restore actions. Protocols cluster DR uses the capabilities of Active File Management based Async Disaster Recovery (AFM DR) to provide a solution that allows an IBM Spectrum Scale cluster to fail over to another cluster and fail back, and backup and restore the protocol configuration information in cases where a secondary cluster is not available. For more information, see *Protocols cluster disaster recovery* in *IBM Spectrum Scale: Advanced Administration Guide*.

### Quality of Service for I/O operations (QoS)

You can use the QoS capability to prevent I/O-intensive, long-running GPFS commands, called *maintenance commands*, from dominating file system performance and significantly delaying normal tasks that also compete for I/O resources. Determine the maximum capacity of your file system in I/O operations per second (IOPS) with the new **mmisqos** command. With the new **mmchqos** command, assign a smaller share of IOPS to the QoS **maintenance** class, which includes all the maintenance commands. Maintenance command instances that are running at the same time compete for the IOPS allocated to the **maintenance** class, and are not allowed to exceed that limit.

### Security mode for new clusters

Starting with IBM Spectrum Scale V4.2, the default security mode for new clusters is AUTHONLY. The **mmcrcluster** command sets the security mode to AUTHONLY when it creates the cluster and automatically generates a public/private key pair for authenticating the cluster. In the AUTHONLY security mode, the sending and receiving nodes authenticate each other with a TLS handshake and then close the TLS connection. Communication continues in the clear. The nodes do not encrypt transmitted data and do not check data integrity.

In IBM Spectrum Scale V4.1 or earlier, the default security mode is EMPTY. If you update a cluster from IBM Spectrum Scale V4.1 to V4.2 or later by running **mmchconfig release=LATEST**, the command checks the security mode. If the mode is EMPTY, the command issues a warning message but does not change the security mode of the cluster.

### Snapshots

You can display information about a global snapshot without displaying information about fileset snapshots with the same name. You can display information about a fileset snapshot without displaying information about other snapshots that have the same name but are snapshots of other filesets.

### spectrumscale Options

The **spectrumscale** command options for installing GPFS and deploying protocols have changed to remove **config enable** and to add **config perf**. For more information, see *IBM Spectrum Scale: Concepts, Planning, and Installation Guide*.

New options have been added to **spectrumscale setup** and **spectrumscale deploy** to disable prompting for the encryption/decryption secret. Note that if **spectrumscale setup --storesecret** is used, passwords will not be secure. New properties have been added to **spectrumscale cofig object** for setting password data instead of doing so through **enable object**. For more information, see *IBM Spectrum Scale: Administration and Programming Reference*.

The **spectrumscale** options for managing share ACLs have been added. For more information, see *IBM Spectrum Scale: Administration and Programming Reference*.

### ssh and scp wrapper scripts

Starting with IBM Spectrum Scale V4.2, a cluster can be configured to use **ssh** and **scp** wrappers. The wrappers allow GPFS to run on clusters where remote root login through **ssh** is disabled. For more information, see the help topic "Running IBM Spectrum Scale without remote root login" in the *IBM Spectrum Scale: Administration and Programming Reference*.

### Documented commands, structures, and subroutines

The following lists the modifications to the documented commands, structures, and subroutines:

### **New commands**

The following commands are new:

- **mmcallhome**
- **mmcesdr**
- **mmchqos**
- **mmlsqos**

### **New structures**

There are no new structures.

### **New subroutines**

There are no new subroutines.

### **Changed commands**

The following commands were changed:

- **mmadddisk**
- **mmaddnode**
- **mmapplypolicy**
- **mmauth**
- **mmbackup**
- **mmces**
- **mmchattr**
- **mmchcluster**
- **mmchconfig**
- **mmchdisk**
- **mmcheckquota**
- **mmchnode**
- **mmcrcluster**
- **mmdefragfs**
- **mmdeldisk**
- **mmdelfileset**
- **mmdelsnapshot**
- **mmdf**
- **mmfileid**
- **mmfsck**
- **mmlsattr**
- **mmlscluster**
- **mmlsconfig**
- **mmlssnapshot**
- **mmnfs**
- **mmobj**
- **mmperfmon**
- **mmprotocoltrace**
- **mmremotefs**
- **mmrestripefile**
- **mmrestripefs**
- **mmrpldisk**
- **mmsdrbackup**

- **mmsdrrestore**
- **mmsmb**
- **mmuserauth**
- **spectrumscale**

**Changed structures**

There are no changed structures.

**Changed subroutines**

There are no changed subroutines.

**Deleted commands**

There are no deleted commands.

**Deleted structures**

There are no deleted structures.

**Deleted subroutines**

There are no deleted subroutines.

**Messages**

The following lists the new, changed, and deleted messages:

**New messages**

6027-2354, 6027-2355, 6027-2356, 6027-2357, 6027-2358, 6027-2359, 6027-2360, 6027-2361,  
6027-2362, 6027-3913, 6027-3914, 6027-3107, 6027-4016, 6027-3317, 6027-3318, 6027-3319,  
6027-3320, 6027-3405, 6027-3406, 6027-3582, 6027-3583, 6027-3584, 6027-3585, 6027-3586,  
6027-3587, 6027-3588, 6027-3589, 6027-3590, 6027-3591, 6027-3592, 6027-3593

**Changed messages**

6027-2299, 6027-887, 6027-888

**Deleted messages**

None.

---

## Chapter 1. Logs, dumps, and traces

The problem determination tools that are provided with IBM Spectrum Scale are intended to be used by experienced system administrators who know how to collect data and run debugging routines.

You can collect various types of logs such as GPFS logs, protocol service logs, operating system logs, and transparent cloud tiering logs. The GPFS™ log is a repository of error conditions that are detected on each node, as well as operational events such as file system mounts. 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 the IBM Spectrum Scale system.

**Note:** The GPFS error logs and messages contain the MMFS prefix to distinguish it from the components of the IBM Multi-Media LAN Server, a related licensed program.

The IBM Spectrum Scale system also provides a system snapshot dump, trace, and other utilities that can be used to obtain detailed information about specific problems.

The information is organized as follows:

- “GPFS logs”
- “The operating system error log facility” on page 19
- “Using the `gpfs.snap` command” on page 23
- “`mmdumpperfdata` command” on page 31
- “`mmfsadm` command” on page 33
- “Trace facility” on page 34

---

### GPFS logs

The GPFS log is a repository of error conditions that are detected on each node, as well as operational events such as file system mounts. The GPFS log is the first place to look when you start debugging the abnormal events. As 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.

The GPFS log can be found in the `/var/adm/ras` directory on each node. The GPFS 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 GPFS log file can be found by using the symbolic file name `/var/adm/ras/mmfs.log.latest`.

The GPFS 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 time stamp and node name appended to the file name.

At GPFS startup, log files that are not accessed during the last 10 days are deleted. If you want to save old log files, copy them elsewhere.

Many GPFS log messages can be sent to `syslog` on Linux. The `systemLogLevel` attribute of the `mmchconfig` command determines the GPFS log messages to be sent to the `syslog`. For more information, see the `mmchconfig` command in the *IBM Spectrum Scale: Administration and Programming Reference*.

This example shows normal operational messages that appear in the GPFS log file on Linux node:

```

Removing old /var/adm/ras/mmfs.log.* files:
Unloading modules from /lib/modules/3.0.13-0.27-default/extra
Unloading module tracedev
Loading modules from /lib/modules/3.0.13-0.27-default/extra
Module          Size Used by
mmfs26          2155186 0
mmfslinux       379348 1 mmfs26
tracedev        48513 2 mmfs26,mmfslinux
Tue Oct 27 11:45:47.149 2015: [I] mmfsd initializing. {Version: 4.2.0.0  Built: Oct 26 2015 15:19:01} ...
Tue Oct 27 11:45:47.150 2015: [I] Tracing in blocking mode
Tue Oct 27 11:45:47.151 2015: [I] Cleaning old shared memory ...
Tue Oct 27 11:45:47.152 2015: [I] First pass parsing mmfs.cfg ...
Tue Oct 27 11:45:47.153 2015: [I] Enabled automated deadlock detection.
Tue Oct 27 11:45:47.154 2015: [I] Enabled automated deadlock debug data collection.
Tue Oct 27 11:45:47.155 2015: [I] Enabled automated expel debug data collection.
Tue Oct 27 11:45:47.156 2015: [I] Initializing the main process ...
Tue Oct 27 11:45:47.169 2015: [I] Second pass parsing mmfs.cfg ...
Tue Oct 27 11:45:47.170 2015: [I] Initializing the page pool ...
Tue Oct 27 11:45:47.500 2015: [I] Initializing the mailbox message system ...
Tue Oct 27 11:45:47.521 2015: [I] Initializing encryption ...
Tue Oct 27 11:45:47.522 2015: [I] Encryption: loaded crypto library: IBM CryptoLite for C v4.10.1.5600 (c4T3/GPFS_LNXPPC64).
Tue Oct 27 11:45:47.523 2015: [I] Initializing the thread system ...
Tue Oct 27 11:45:47.524 2015: [I] Creating threads ...
Tue Oct 27 11:45:47.529 2015: [I] Initializing inter-node communication ...
Tue Oct 27 11:45:47.530 2015: [I] Creating the main SDR server object ...
Tue Oct 27 11:45:47.531 2015: [I] Initializing the sdrServ library ...
Tue Oct 27 11:45:47.532 2015: [I] Initializing the ccrServ library ...
Tue Oct 27 11:45:47.538 2015: [I] Initializing the cluster manager ...
Tue Oct 27 11:45:48.813 2015: [I] Initializing the token manager ...
Tue Oct 27 11:45:48.819 2015: [I] Initializing network shared disks ...
Tue Oct 27 11:45:51.126 2015: [I] Start the ccrServ ...
Tue Oct 27 11:45:51.879 2015: [N] Connecting to 192.168.115.171 js21n07 <c0p1>
Tue Oct 27 11:45:51.880 2015: [I] Connected to 192.168.115.171 js21n07 <c0p1>
Tue Oct 27 11:45:51.897 2015: [I] Node 192.168.115.171 (js21n07) is now the Group Leader.
Tue Oct 27 11:45:51.911 2015: [N] mmfsd ready
Tue Oct 27 11:45:52 EDT 2015: mmcommon mmfsup invoked. Parameters: 192.168.115.220 192.168.115.171 all

```

The **mmcommon logRotate** command can be used to rotate the GPFS log without shutting down and restarting the daemon. After the **mmcommon logRotate** command is issued, **/var/adm/ras/mmfs.log.previous** will contain the messages that occurred since the previous startup of GPFS or the last run of **mmcommon logRotate**. The **/var/adm/ras/mmfs.log.latest** file starts over at the point in time that **mmcommon logRotate** was run.

Depending on the size and complexity of your system configuration, the amount of time to start GPFS varies. If you cannot access a file system that is mounted, examine the log file for error messages.

## Creating a master GPFS log file

The GPFS 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 GPFS log entries from each node in the cluster. The **--gather-logs** option of the **gpfs.snap** command can be used to achieve this:

```
gpfs.snap --gather-logs -d /tmp/logs -N all
```



The system displays output similar to the following example:

```
Jun 3 12:04:34 prt001st001 winbindd[14656]: [2015/06/03 12:04:34.271459, 0] ../lib/util/become_daemon.c:124(daemon_ready)
Jun 3 12:04:34 prt001st001 winbindd[14656]: STATUS=daemon 'winbindd' finished starting up and ready to serve connections
```

Additional logs are available in `/var/adm/ras/log.winbindd*` and `/var/adm/ras/log.wb*`. There are multiple files that get rotated with the “old” suffix, once it becomes the size of a 100 MB.

To capture debug traces for Active Directory authentication, use the following command to enable tracing:

```
mmdsh -N CesNodes /usr/lpp/mmfs/bin/smbcontrol winbindd debug 10
```

To disable tracing for Active Directory authentication, use the following command:

```
mmdsh -N CesNodes /usr/lpp/mmfs/bin/smbcontrol winbindd debug 1
```

**Related concepts:**

“Determining the health of integrated SMB server” on page 122

## NFS logs

The clustered export services (CES) NFS server writes log messages in the `/var/log/ganesha.log` file at runtime.

Operating system's log rotation facility is used to manage NFS logs. The NFS logs are configured and enabled during the NFS server packages installation.

The following example shows a sample log file:

```
# tail -f /var/log/ganesha.log
2015-05-31 17:08:04 : epoch 556b23d2 : cluster1.ibm.com : ganesha.nfsd-27204[main] nfs_rpc_cb_init_ccache
:NFS STARTUP :WARN :gssd_refresh_krb5_machine_credential failed (-1765328160:0)
2015-05-31 17:08:04 : epoch 556b23d2 : cluster1.ibm.com : ganesha.nfsd-27204[main] nfs_Start_threads
:THREAD :EVENT :Starting delayed executor.
2015-05-31 17:08:04 : epoch 556b23d2 : cluster1.ibm.com : ganesha.nfsd-27204[main] nfs_Start_threads
:THREAD :EVENT :gsh_dbusthread was started successfully
2015-05-31 17:08:04 : epoch 556b23d2 : cluster1.ibm.com : ganesha.nfsd-27204[main] nfs_Start_threads
:THREAD :EVENT :admin thread was started successfully
2015-05-31 17:08:04 : epoch 556b23d2 : cluster1.ibm.com : ganesha.nfsd-27204[main] nfs_Start_threads
:THREAD :EVENT :reaper thread was started successfully
2015-05-31 17:08:04 : epoch 556b23d2 : cluster1.ibm.com : ganesha.nfsd-27204[main] nfs_Start_threads
:THREAD :EVENT :General fridge was started successfully
2015-05-31 17:08:04 : epoch 556b23d2 : cluster1.ibm.com : ganesha.nfsd-27204[reaper] nfs_in_grace
:STATE :EVENT :NFS Server Now IN GRACE
2015-05-31 17:08:04 : epoch 556b23d2 : cluster1.ibm.com : ganesha.nfsd-27204[main] nfs_start
:NFS STARTUP :EVENT :-----
2015-05-31 17:08:04 : epoch 556b23d2 : cluster1.ibm.com : ganesha.nfsd-27204[main] nfs_start
:NFS STARTUP :EVENT : NFS SERVER INITIALIZED
2015-05-31 17:08:04 : epoch 556b23d2 : cluster1.ibm.com : ganesha.nfsd-27204[main] nfs_start
:NFS STARTUP :EVENT :-----
2015-05-31 17:09:04 : epoch 556b23d2 : cluster1.ibm.com : ganesha.nfsd-27204[reaper] nfs_in_grace
:STATE :EVENT :NFS Server Now NOT IN GRACE
```

Log levels can be displayed by using the `mmnfs configuration list | grep LOG_LEVEL` command. For example:

```
mmnfs configuration list | grep LOG_LEVEL
```

The system displays output similar to the following example:

```
LOG_LEVEL: EVENT
```

By default, the log level is EVENT. Additionally, the following NFS log levels can also be used; starting from lowest to highest verbosity:



- FATAL
- MAJ
- CRIT
- WARN
- INFO
- DEBUG
- MID\_DEBUG
- FULL\_DEBUG

**Note:** The FULL\_DEBUG level increases the size of the log file. Use it in the production mode only if instructed by the IBM Support.

Increasing the verbosity of the NFS server log impacts the overall NFS I/O performance.

To change the logging to the verbose log level INFO, use the following command:

#### **mmnfs configuration change LOG\_LEVEL=INFO**

The system displays output similar to the following example:

NFS Configuration successfully changed. NFS server restarted on all NFS nodes.

This change is cluster-wide and restarts all NFS instances to activate this setting. The log file now displays more informational messages, for example:

```
2015-06-03 12:49:31 : epoch 556edba9 : cluster1.ibm.com : ganesha.nfsd-21582[main] nfs_rpc_dispatch_threads
:THREAD :INFO :5 rpc dispatcher threads were started successfully
2015-06-03 12:49:31 : epoch 556edba9 : cluster1.ibm.com : ganesha.nfsd-21582[disp] rpc_dispatcher_thread
:DISP :INFO :Entering nfs/rpc dispatcher
2015-06-03 12:49:31 : epoch 556edba9 : cluster1.ibm.com : ganesha.nfsd-21582[disp] rpc_dispatcher_thread
:DISP :INFO :Entering nfs/rpc dispatcher
2015-06-03 12:49:31 : epoch 556edba9 : cluster1.ibm.com : ganesha.nfsd-21582[disp] rpc_dispatcher_thread
:DISP :INFO :Entering nfs/rpc dispatcher
2015-06-03 12:49:31 : epoch 556edba9 : cluster1.ibm.com : ganesha.nfsd-21582[disp] rpc_dispatcher_thread
:DISP :INFO :Entering nfs/rpc dispatcher
2015-06-03 12:49:31 : epoch 556edba9 : cluster1.ibm.com : ganesha.nfsd-21582[main] nfs_Start_threads
:THREAD :EVENT :gsh_dbusthread was started successfully
2015-06-03 12:49:31 : epoch 556edba9 : cluster1.ibm.com : ganesha.nfsd-21582[main] nfs_Start_threads
:THREAD :EVENT :admin thread was started successfully
2015-06-03 12:49:31 : epoch 556edba9 : cluster1.ibm.com : ganesha.nfsd-21582[main] nfs_Start_threads
:THREAD :EVENT :reaper thread was started successfully
2015-06-03 12:49:31 : epoch 556edba9 : cluster1.ibm.com : ganesha.nfsd-21582[main] nfs_Start_threads
:THREAD :EVENT :General fridge was started successfully
2015-06-03 12:49:31 : epoch 556edba9 : cluster1.ibm.com : ganesha.nfsd-21582[reaper] nfs_in_grace
:STATE :EVENT :NFS Server Now IN GRACE
2015-06-03 12:49:32 : epoch 556edba9 : cluster1.ibm.com : ganesha.nfsd-21582[main] nfs_start
:NFS STARTUP :EVENT :-----
2015-06-03 12:49:32 : epoch 556edba9 : cluster1.ibm.com : ganesha.nfsd-21582[main] nfs_start
:NFS STARTUP :EVENT : NFS SERVER INITIALIZED
2015-06-03 12:49:32 : epoch 556edba9 : cluster1.ibm.com : ganesha.nfsd-21582[main] nfs_start
:NFS STARTUP :EVENT :-----
2015-06-03 12:50:32 : epoch 556edba9 : cluster1.ibm.com : ganesha.nfsd-21582[reaper] nfs_in_grace
:STATE :EVENT :NFS Server Now NOT IN GRACE
```

To display the currently configured CES log level, use the following command:

#### **mmces log level**

The system displays output similar to the following example:

CES log level is currently set to 0

The log file is `/var/adm/ras/mmfs.log.latest`. By default, the log level is 0 and other possible values are 1, 2, and 3. To increase the log level, use the following command:

```
mmces log level 1
```

NFS-related log information is written to the standard GPFS log files as part of the overall CES infrastructure. This information relates to the NFS service management and recovery orchestration within CES.

## Object logs

There are a number of locations where messages are logged with the Object protocol.

The core Object services, proxy, account, container, and Object server have their own logging level sets in their respective configuration files. By default, Swift logging is set to show messages at or above the ERROR level, but can be changed to INFO or DEBUG levels if more detailed logging information is required.

By default, the messages logged by these services are saved in the `/var/log/swift` directory.

You can also configure these services to use separate syslog facilities by the `log_facility` parameter in one or all of the Object service configuration files and by updating the rsyslog configuration. These parameters are described in the Swift Deployment Guide ([docs.openstack.org/developer/swift/deployment\\_guide.html](https://docs.openstack.org/developer/swift/deployment_guide.html)) that is available in the OpenStack documentation.

An example of how to set up this configuration can be found in the SAIO - Swift All In One documentation ([docs.openstack.org/developer/swift/development\\_saio.html#optional-setting-up-rsyslog-for-individual-logging](https://docs.openstack.org/developer/swift/development_saio.html#optional-setting-up-rsyslog-for-individual-logging)) that is available in the OpenStack documentation.

**Note:** To configure rsyslog for unique log facilities in the protocol nodes, the administrator needs to ensure that the manual steps mentioned in the preceding link are carried out on each of those protocol nodes.

The Keystone authentication service writes its logging messages to `/var/log/keystone/keystone.log` file. By default, Keystone logging is set to show messages at or above the WARNING level.

For information on how to view or change log levels on any of the Object related services, see the “CES collection and tracing” section in the *IBM Spectrum Scale: Advanced Administration Guide*.

The following commands can be used to determine the health of Object services:

- To see whether there are any nodes in an active (failed) state, run the following command:

```
mmces state cluster OBJ
```

The system displays output similar to this:

NODE	COMPONENT	STATE	EVENTS
prt001st001	OBJECT	HEALTHY	
prt002st001	OBJECT	HEALTHY	
prt003st001	OBJECT	HEALTHY	
prt004st001	OBJECT	HEALTHY	
prt005st001	OBJECT	HEALTHY	
prt006st001	OBJECT	HEALTHY	
prt007st001	OBJECT	HEALTHY	

In this example, all nodes are healthy so no active events are shown.

- To display the history of events generated by the monitoring framework, run the following command:

```
mmces events list OBJ
```

The system displays output similar to this:

Node	Timestamp	Event Name	Severity	Details
node1	2015-06-03 13:30:27.478725+08:08PDT	proxy-server_ok	INFO	proxy process as expected
node1	2015-06-03 14:26:30.567245+08:08PDT	object-server_ok	INFO	object process as expected
node1	2015-06-03 14:26:30.720534+08:08PDT	proxy-server_ok	INFO	proxy process as expected
node1	2015-06-03 14:28:30.689257+08:08PDT	account-server_ok	INFO	account process as expected
node1	2015-06-03 14:28:30.853518+08:08PDT	container-server_ok	INFO	container process as expected
node1	2015-06-03 14:28:31.015307+08:08PDT	object-server_ok	INFO	object process as expected
node1	2015-06-03 14:28:31.177589+08:08PDT	proxy-server_ok	INFO	proxy process as expected
node1	2015-06-03 14:28:49.025021+08:08PDT	postIpChange_info	INFO	IP addresses modified 192.167.12.21_0-1.
node1	2015-06-03 14:28:49.194499+08:08PDT	enable_Address_database_node	INFO	Enable Address Database Node
node1	2015-06-03 14:29:16.483623+08:08PDT	postIpChange_info	INFO	IP addresses modified 192.167.12.22_0-2.
node1	2015-06-03 14:29:25.274924+08:08PDT	postIpChange_info	INFO	IP addresses modified 192.167.12.23_0-3.
node1	2015-06-03 14:29:30.844626+08:08PDT	postIpChange_info	INFO	IP addresses modified 192.167.12.24_0-4.

- To retrieve the OBJ related log entries, query the monitor client and grep for the name of the component you want to filter on, either Object, proxy, account, container, keystone or postgres. For example, to see proxy-server related events, run the following command:

```
mmces events list | grep proxy
```

The system displays output similar to this:

node1	2015-06-01 14:39:49.120912+08:08PDT	proxy-server_failed	ERROR	proxy process should be started but is stopped
node1	2015-06-01 14:44:49.277940+08:08PDT	proxy-server_ok	INFO	proxy process as expected
node1	2015-06-01 16:27:37.923696+08:08PDT	proxy-server_failed	ERROR	proxy process should be started but is stopped
node1	2015-06-01 16:40:39.789920+08:08PDT	proxy-server_ok	INFO	proxy process as expected
node1	2015-06-03 13:28:18.875566+08:08PDT	proxy-server_failed	ERROR	proxy process should be started but is stopped
node1	2015-06-03 13:30:27.478725+08:08PDT	proxy-server_ok	INFO	proxy process as expected
node1	2015-06-03 13:30:57.482977+08:08PDT	proxy-server_failed	ERROR	proxy process should be started but is stopped
node1	2015-06-03 14:26:30.720534+08:08PDT	proxy-server_ok	INFO	proxy process as expected
node1	2015-06-03 14:27:00.759696+08:08PDT	proxy-server_failed	ERROR	proxy process should be started but is stopped
node1	2015-06-03 14:28:31.177589+08:08PDT	proxy-server_ok	INFO	proxy process as expected

- To check the monitor log, grep for the component you want to filter on, either Object, proxy, account, container, keystone or postgres. For example, to see Object-server related log messages:

```
grep object /var/adm/ras/mmcesmonitor.log | head -n 10
```

The system displays output similar to this:

```
2015-06-03T13:59:28.805-08:00 util5.sonasad.almaden.ibm.com D:522632:Thread-9:object:OBJ running command
'systemctl status openstack-swift-proxy'
2015-06-03T13:59:28.916-08:00 util5.sonasad.almaden.ibm.com D:522632:Thread-9:object:OBJ command result1t
ret:3 sout:openstack-swift-proxy.service - OpenStack Object Storage (swift) - Proxy Server
2015-06-03T13:59:28.916-08:00 util5.sonasad.almaden.ibm.com I:522632:Thread-9:object:OBJ openstack-swift-proxy is not started, ret3
2015-06-03T13:59:28.916-08:00 util5.sonasad.almaden.ibm.com D:522632:Thread-9:object:OBJProcessMonitor openstack-swift-proxy failed:
2015-06-03T13:59:28.916-08:00 util5.sonasad.almaden.ibm.com D:522632:Thread-9:object:OBJProcessMonitor memcached started
2015-06-03T13:59:28.917-08:00 util5.sonasad.almaden.ibm.com D:522632:Thread-9:object:OBJ running command
'systemctl status memcached'
2015-06-03T13:59:29.018-08:00 util5.sonasad.almaden.ibm.com D:522632:Thread-9:object:OBJ command result1t
ret:0 sout:memcached.service - Memcached
2015-06-03T13:59:29.018-08:00 util5.sonasad.almaden.ibm.com I:522632:Thread-9:object:OBJ memcached is started and active running
2015-06-03T13:59:29.018-08:00 util5.sonasad.almaden.ibm.com D:522632:Thread-9:object:OBJProcessMonitor memcached succeeded
2015-06-03T13:59:29.018-08:00 util5.sonasad.almaden.ibm.com I:522632:Thread-9:object:OBJ service started checks
after monitor loop, event count:6
```

The following tables list the IBM Spectrum Scale for object storage log files.

*Table 3. Core object log files in /var/log/swift. Core object log files in /var/log/swift*

Log file	Component	Configuration file
account-auditor.log	Account auditor Swift service	account-server.conf
account-auditor.error		
account-reaper.log	Account reaper Swift service	account-server.conf
account-reaper.error		
account-replicator.log	Account replicator Swift service	account-server.conf
account-replicator.error		
account-server.log	Account server Swift service	account-server.conf
account-server.error		
container-auditor.log	Container auditor Swift service	container-server.conf
container-auditor.error		

Table 3. Core object log files in /var/log/swift (continued). Core object log files in /var/log/swift

Log file	Component	Configuration file
container-replicator.log container-replicator.error	Container replicator Swift service	container-server.conf
container-server.log container-server.error	Container server Swift service	container-server.conf
container-updater.log container-updater.error	Container updater Swift service	container-server.conf
object-auditor.log object-auditor.error	Object auditor Swift service	object-server.conf
object-expirer.log object-expirer.error	Object expirer Swift service	object-expirer.conf
object-replicator.log object-replicator.error	Object replicator Swift service	object-server.conf
object-server.log object-server.error	Object server Swift service	object-server.conf object-server-sof.conf
object-updater.log object-updater.error	Object updater Swift service	object-server.conf
proxy-server.log proxy-server.error	Proxy server Swift service	proxy-server.conf

Table 4. Additional object log files in /var/log/swift. Additional object log files in /var/log/swift

Log file	Component	Configuration file
ibmobjectizer.log ibmobjectizer.error	Unified file and object access objectizer service	spectrum-scale-objectizer.conf spectrum-scale-object.conf
policyscheduler.log policyscheduler.error	Object storage policies	spectrum-scale-object- policies.conf
swift.log swift.error	Performance metric collector (pmswift)	

Table 5. General system log files in /var/adm/ras. General system log files in /var/adm/ras

Log file	Component
mmcesmonitor.log	CES framework services monitor
mmfs.log	Various IBM Spectrum Scale command logging

## The IBM Spectrum Scale HDFS transparency log

In IBM Spectrum Scale HDFS transparency, all logs are recorded using log4j. The **log4j.properties** file is under the /usr/lpp/mmfs/hadoop/etc/hadoop directory.

By default, the logs are written under the /usr/lpp/mmfs/hadoop/logs directory.

The following entries can be added into the **log4j.properties** file to turn on the debugging information:

```
log4j.logger.org.apache.hadoop.yarn=DEBUG
log4j.logger.org.apache.hadoop.hdfs=DEBUG
log4j.logger.org.apache.hadoop.gpfs=DEBUG
log4j.logger.org.apache.hadoop.security=DEBUG
```

## Protocol authentication log files

The log files pertaining to protocol authentication are described here.

*Table 6. Authentication log files*

Service name	Log configuration file	Log files	Logging levels
Keystone	/etc/keystone/keystone.conf  /etc/keystone/logging.conf	/var/log/keystone/keystone.log  /var/log/keystone/httpd-error.log  /var/log/keystone/httpd-access.log	In keystone.conf change  1. debug = true- for getting debugging information in log file.  2. verbose = true - for getting Info messages in log file .  By default, these values are false and only warning messages are logged.  Finer grained control of Keystone logging levels can be specified by updating the Keystones logging.conf file. For information on the logging levels in the logging.conf file, see OpenStack logging.conf documentation ( <a href="https://docs.openstack.org/kilo/config-reference/content/section_keystone-logging.conf.html">docs.openstack.org/kilo/config-reference/content/section_keystone-logging.conf.html</a> ).

Table 6. Authentication log files (continued)

Service name	Log configuration file	Log files	Logging levels
SSSD	/etc/sss/ sss.conf	/var/log/sss/sss.log  /var/log/sss/sss_nss.log  /var/log/sss/ sss_LDAPDOMAIN.log (depends upon configuration)  /var/log/sss/ sss_NISDOMAIN.log (depends upon configuration) <b>Note:</b> For more information on SSSD log files, see Red Hat Enterprise Linux documentation.	0x0010: Fatal failures. Issue with invoking or running SSSD.  0x0020: Critical failures. SSSD does not stop functioning. However, this error indicates that at least one major feature of SSSD is not to work properly.  0x0040: Serious failures. A particular request or operation has failed.  0x0080: Minor failures. These are the errors that would percolate down to cause the operation failure of 2.  0x0100: Configuration settings.  0x0200: Function data.  0x0400: Trace messages for operation functions.  0x1000: Trace messages for internal control functions.  0x2000: Contents of function-internal variables that might be interesting.  0x4000: Extremely low-level tracing information. <b>Note:</b> For more information on SSSD log levels, see Red Hat Enterprise Linux documentation.
Winbind	/var/mmfs /ces/smb.conf	/var/adm/ras/log.wb-<DOMAIN>  [Depends upon available domains]  /var/adm/ras/log.winbindd-dc- connect  /var/adm/ras/log.winbindd-idmap  /var/adm/ras/log.winbindd	Log level is an integer. The value can be from 0-10.  The default value for log level is 1.

**Note:** Some of the authentication modules like Keystone services log information also in /var/log/messages.

If you change the log levels, the respective authentication service must be restarted manually on each protocol node. Restarting authentication services might result in disruption of protocol I/O.

## CES monitoring and troubleshooting

You can monitor system health, query events, and perform maintenance and troubleshooting tasks in Cluster Export Services (CES).

### System health monitoring

Each CES node runs a separate GPFS process that monitors the network address configuration of the node. If a conflict between the network interface configuration of the node and the current assignments of the CES address pool is found, corrective action is taken. If the node is unable to detect an address that is assigned to it, the address is reassigned to another node.

Additional monitors check the state of the services that are implementing the enabled protocols on the node. These monitors cover NFS, SMB, Object, and Authentication services that monitor, for example, daemon liveness and port responsiveness. If it is determined that any enabled service is not functioning correctly, the node is marked as failed and its CES addresses are reassigned. When the node returns to normal operation, it returns to the normal (healthy) state and is available to host addresses in the CES address pool.

An additional monitor runs on each protocol node if Microsoft Active Directory (AD), Lightweight Directory Access Protocol (LDAP), or Network Information Service (NIS) user authentication is configured. If a configured authentication server does not respond to test requests, GPFS marks the affected node as failed.

### Querying state and events

Aside from the automatic failover and recovery of CES addresses, two additional outputs are provided by the monitoring that can be queried: events and state.

State can be queried by entering the `mmces state show` command, which shows you the state of each of the CES components. The possible states for a component follow:

#### HEALTHY

The component is working as expected.

#### DISABLED

The component has not been enabled.

#### SUSPENDED

When a CES node is in the suspended state, most components also report suspended.

#### STARTING

The component (or monitor) recently started. This state is a transient state that is updated after the startup is complete.

#### UNKNOWN

Something is preventing the monitoring from determining the state of the component.

#### STOPPED

The component was intentionally stopped. This situation might happen briefly if a service is being restarted due to a configuration change. It might also happen because a user ran the `mmces service stop protocol` command for a node.

#### DEGRADED

There is a problem with the component but not a complete failure. This state does not cause the CES addresses to be reassigned.

#### FAILED

The monitoring detected a significant problem with the component that means it is unable to function correctly. This state causes the CES addresses of the node to be reassigned.

## DEPENDENCY\_FAILED

This state implies that a component has a dependency that is in a failed state. An example would be NFS or SMB reporting `DEPENDENCY_FAILED` because the authentication failed.

Looking at the states themselves can be useful to find out which component is causing a node to fail and have its CES addresses reassigned. To find out why the component is being reported as failed, you can look at events.

The `mmces events` command can be used to show you either events that are currently causing a component to be unhealthy or a list of historical events for the node. If you want to know why a component on a node is in a failed state, use the `mmces events active` invocation. This command gives you a list of any currently active events that are affecting the state of a component, along with a message that describes the problem. This information should provide a place to start when you are trying to find and fix the problem that is causing the failure.

If you want to get a complete idea of what is happening with a node over a longer time period, use the `mmces events list` invocation. By default, this command prints a list of all events that occurred on this node, with a time stamp. This information can be narrowed down by component, time period, and severity. As well as being viewable with the command, all events are also pushed to the syslog.

## Maintenance and troubleshooting

A CES node can be marked as unavailable by the monitoring process. The command `mmces node list` can be used to show the nodes and the current state flags that are associated with it. When unavailable (one of the following node flags are set), the node does not accept CES address assignments. The following possible node states can be displayed:

### Suspended

Indicates that the node is suspended with the `mmces node suspend` command. When suspended, health monitoring on the node is discontinued. The node remains in the suspended state until it is resumed with the `mmces node resume` command.

### Network-down

Indicates that monitoring found a problem that prevents the node from bringing up the CES addresses in the address pool. The state reverts to normal when the problem is corrected. Possible causes for this state are missing or non-functioning network interfaces and network interfaces that are reconfigured so that the node can no longer host the addresses in the CES address pool.

### No-shared-root

Indicates that the CES shared root directory cannot be accessed by the node. The state reverts to normal when the shared root directory becomes available. Possible cause for this state is that the file system that contains the CES shared root directory is not mounted.

**Failed** Indicates that monitoring found a problem with one of the enabled protocol servers. The state reverts to normal when the server returns to normal operation or when the service is disabled.

### Starting up

Indicates that the node is starting the processes that are required to implement the CES services that are enabled in the cluster. The state reverts to normal when the protocol servers are functioning.

Additionally, events that affect the availability and configuration of CES nodes are logged in the GPFS log file `/var/adm/ras/mmfs.log.latest`. The verbosity of the CES logging can be changed with the `mmces log level n` command, where *n* is a number from 0 (less logging) to 4 (more logging). The current log level can be viewed with the `mmfscluster --ces` command.

For more information about CES troubleshooting, see the IBM Spectrum Scale Wiki ([www.ibm.com/developerworks/community/wikis/home/wiki/General Parallel File System \(GPFS\)](http://www.ibm.com/developerworks/community/wikis/home/wiki/General%20Parallel%20File%20System%20(GPFS))).



## CES tracing and debug data collection

You can collect debugging information in Cluster Export Services.

### Data collection (FTDC)

To diagnose the cause of an issue, it might be necessary to gather some extra information from the cluster. This information can then be used to determine the root cause of an issue.

Collection of debugging information, such as configuration files and logs, can be gathered by using the **gpfs.snap** command. This command gathers data about GPFS, operating system information, and information for each of the protocols.

#### GPFS + OS

GPFS configuration and logs plus operating system information such as network configuration or connected drives.

**CES** Generic protocol information such as configured CES nodes.

**NFS** NFS Ganesha configuration and logs.

**SMB** SMB and CTDB configuration and logs.

#### OBJECT

Openstack Swift and Keystone configuration and logs.

#### AUTHENTICATION

Authentication configuration and logs.

#### PERFORMANCE

Dump of the performance monitor database.

Information for each of the enabled protocols is gathered automatically when the **gpfs.snap** command is run. If any protocol is enabled, then information for CES and authentication is gathered.

To gather performance data, add the **--performance** option. The **--performance** option causes **gpfs.snap** to try to collect performance information.

**Note:** Because this process can take up to 30 minutes to run, gather performance data only if necessary.

If data is only required for one protocol or area, the automatic collection can be bypassed. Provided one or more of the following options to the **--protocol** argument: **smb,nfs,object,ces,auth,none**

If the **--protocol** command is provided, automatic data collection is disabled. If **--protocol smb,nfs** is provided to **gpfs.snap**, only NFS and SMB information is gathered and no CES or Authentication data is collected. To disable all protocol data collection, use the argument **--protocol none**.

## Types of tracing

Tracing is logging at a high level. The command for starting and stopping tracing (**mmprotocoltrace**) supports SMB tracing. NFS and Object tracing can be done with a combination of commands.

**SMB** To start SMB tracing, use the **mmprotocoltrace start smb** command. The output looks similar to this example:

```
Starting traces
Trace 'd83235aa-0589-4866-aaf0-2e285aad6f92' created successfully
```

**Note:** Running the **mmprotocoltrace start smb** command without the **-c** option enables tracing for all SMB connections. This configuration can slow performance. Therefore, consider adding the **-c** option to trace connections for specific client IP addresses.

To see the status of the trace command, use the **mmprotocoltrace status smb** command. The output looks similar to this example:

```

Trace ID: d11145ea-9e9a-4fb0-ae8d-7cb48e49ecc2
State: WAITING
User ID: root
Protocol: smb
Start Time: 11:11:37 05/05/2015
End Time: 11:21:37 05/05/2015
Client IPs: []
Origin Node: swift-test-08.stglab.manchester.uk.ibm.com
Nodes:
Node Name: swift-test-07.stglab.manchester.uk.ibm.com
State: WAITING
Trace Location: /dump/ftdc/smb.20150505_111136.trc
Pids: []

```

To stop the trace the command, use the **mmprotocoltrace stop smb** command:

```

Stopping traces
Trace '01239483-be84-wev9-a2d390i9ow02' stopped for smb
Waiting for traces to complete
Waiting for node 'swift-test-07'
Waiting for node 'swift-test-08'
Finishing trace 'd11145ea-9e9a-4fb0-ae8d-7cb48e49ecc2'
Trace tar file has been written to '/tmp/mmfs/smb.20150513_162322.trc/smb.trace.20150513_162542.tar.gz

```

The tar file then includes the log files that contain top-level logs for the time period the trace was running for.

Traces time out after a certain amount of time. By default, this time is 10 minutes. The timeout can be changed by using the **-d** argument when you start the trace. When a trace times out, the first node with the timeout ends the trace and writes the location of the collected data into the **mmprotocoltrace** logs. Each other node writes an information message that states that another node ended the trace.

A full usage message for the **mmprotocoltrace** command is printable by using the **-h** argument.

**NFS** NFS tracing is achieved by increasing the log level, repeating the issue, capturing the log file, and then restoring the log level.

To increase the log level, use the command **mmnfs** configuration change **LOG\_LEVEL=FULL\_DEBUG**.

You can set the log level to the following values: **NULL**, **FATAL**, **MAJ**, **CRIT**, **WARN**, **EVENT**, **INFO**, **DEBUG**, **MID\_DEBUG**, and **FULL\_DEBUG**.

**FULL\_DEBUG** is the most useful for debugging purposes.

After the issue is recreated by running the **gpfs.snab** command either with no arguments or with the **--protocol nfs** argument, the NFS logs are captured. The logs can then be used to diagnose any issues.

To return the log level to normal, use the same command but with a lower logging level (the default is **EVENT**).

## Object

The process for tracing the object protocol is similar to NFS. The Object service consists of multiple processes that can be controlled individually.

The Object services use these logging levels, at increasing severity: **DEBUG**, **INFO**, **AUDIT**, **WARNING**, **ERROR**, **CRITICAL**, and **TRACE**.

### Keystone and Authentication

```

mmobj config change --ccrfile keystone.conf --section DEFAULT --property debug
--value True

```

Finer grained control of Keystone logging levels can be specified by updating the Keystone's **logging.conf** file. For information on the logging levels in the **logging.conf** file, see the OpenStack **logging.conf** documentation ([docs.openstack.org/kilo/config-reference/content/section\\_keystone-logging.conf.html](http://docs.openstack.org/kilo/config-reference/content/section_keystone-logging.conf.html)).

### Swift Proxy Server

```
mmobj config change --ccrfile proxy-server.conf --section DEFAULT --property
log_level --value DEBUG
```

### Swift Account Server

```
mmobj config change --ccrfile account-server.conf --section DEFAULT --property
log_level --value DEBUG
```

### Swift Container Server

```
mmobj config change --ccrfile container-server.conf --section DEFAULT --property
log_level --value DEBUG
```

### Swift Object Server

```
mmobj config change --ccrfile object-server.conf --section DEFAULT --property
log_level --value DEBUG
```

These commands increase the log level for the particular process to the debug level. After you have re-created the problem, run the **gpfs.snap** command with no arguments or with the **--protocol object** argument.

Then, decrease the log levels again by using the commands that are shown previously but with **--value ERROR** instead of **--value DEBUG**.

## Collecting trace information

Use the **mmprotocoltrace** command to collect trace information for debugging system problems or performance issues. For more information, see the **mmprotocoltrace** command in the *IBM Spectrum Scale: Administration and Programming Reference*. This section is divided into the following subsections:

- “Running a typical trace”
- “Trace timeout” on page 16
- “Trace log files” on page 17
- “Trace configuration file” on page 17
- “Resetting the trace system ” on page 18
- “Using advanced options” on page 19

## Running a typical trace

The following steps describe how to run a typical trace. It is assumed that the trace system is reset for the type of trace that you want to run: SMB, Network, or Object. The examples use the SMB trace.

1. Before you start the trace, you can check the configuration settings for the type of trace that you plan to run:

```
mmprotocoltrace config smb
```

The response to this command displays the current settings from the trace configuration file. For more information about this file, see the “Trace configuration file” on page 17 subtopic.

2. Clear the trace records from the previous trace of the same type:

```
mmprotocoltrace clear smb
```

This command responds with an error message if the previous state of a trace node is something other than **DONE** or **FAILED**. If this error occurs, follow the instructions in the “Resetting the trace system ” on page 18 subtopic.

3. Start the new trace:

```
]# mmprotocoltrace start smb
```

The following response is typical:

```
Trace '3f36dbed-b567-4566-9beb-63b6420bbb2d' created successfully for 'smb'
```

4. Check the status of the trace to verify that tracing is active on all the configured nodes:

```
]# mmprotocoltrace status smb
```

The following response is typical:

```
Trace ID:    d11145ea-9e9a-4fb0-ae8d-7cb48e49ecc2
State:      WAITING
User ID:    root
Protocol:   smb
Start Time: 11:11:37 05/05/2015
End Time:   11:21:37 05/05/2015
Client IPs: []
Origin Node: swift-test-08.stglab.manchester.uk.ibm.com
Nodes:
```

```
    Node Name:    swift-test-07.stglab.manchester.uk.ibm.com
    State:        WAITING
```

```
    Node Name:    swift-test-08.stglab.manchester.uk.ibm.com
    State:        WAITING
```

To display more status information, add the **-v** (verbose) option:

```
]# mmprotocoltrace -v status smb
```

If the status of a node is **FAILED**, the node did not start successfully. Look at the logs for the node to determine the problem. After you fix the problem, reset the trace system by following the steps in the “Resetting the trace system ” on page 18 subtopic.

5. If all the nodes started successfully, perform the actions that you want to trace. For example, if you are tracing a client IP address, enter commands that create traffic on that client.
6. Stop the trace:

```
]# mmprotocoltrace stop smb
```

The following response is typical. The last line gives the location of the trace log file:

```
Stopping traces
Trace '01239483-be84-wev9-a2d390i9ow02' stopped for smb
Waiting for traces to complete
Waiting for node 'node1'
Waiting for node 'node2'
Finishing trace '01239483-be84-wev9-a2d390i9ow02'
Trace tar file has been written to '/tmp/mmfs/smb.20150513_162322.trc/smb.trace.20150513_162542.tar.gz'
```

If you do not stop the trace, it continues until the trace duration expires. For more information, see the “Trace timeout” subtopic.

7. Look in the trace log files for the results of the trace. For more information, see the “Trace log files” on page 17 subtopic.

## Trace timeout

If you do not stop a trace manually, the trace runs until its trace duration expires. The default trace duration is 10 minutes, but you can set a different value in the **mmprotocoltrace** command. Each node that participates in a trace starts a timeout process that is set to the trace duration. When a timeout occurs, the process checks the trace status. If the trace is active, the process stops the trace, writes the file location to the log file, and exits. If the trace is not active, the timeout process exits.

If a trace stops because of a timeout, look in the log file of each node to find the location of the trace log file. The log entry is similar to the following entry:

```
2015-08-26T16:53:35.885 W:14150:MainThread:TIMEOUT:
    Trace 'd4643ccf-96c1-467d-93f8-9c71db7333b2' tar file located at
    '/tmp/mmfs/smb.20150826_164328.trc/smb.trace.20150826_165334.tar.gz'
```

## Trace log files

Trace log files are compressed files in the `/var/adm/ras` directory. The contents of a trace log file depends on the type of trace. The product supports three types of tracing: SMB, Network, and Object.

**SMB** SMB tracing captures System Message Block information. The resulting trace log file contains an `smbd.log` file for each node for which information has been collected . A global trace captures information for all the clients that are connected to the SMB server. A targeted trace captures information for the specified IP address.

### Network

Network tracing calls Wireshark's `dumpcap` utility to capture network packets. The resulting trace log file contains a `pcappng` file that is readable by Wireshark and other programs. The file name is similar to `bfn22-10g_all_00001_20150907125015.pcap`.

If the `mmprotocoltrace` command specifies a client IP address, the trace captures traffic between that client and the server. If no IP address is specified, the trace captures traffic across all network interfaces of each participating node.

### Object

The trace log file contains log files for each node, one for each of the object services.

Object tracing sets the log location in the `rsyslog` configuration file. For more information about this file, see the description of the `rsyslogconflocation` configuration parameter in the "Trace configuration file" subtopic.

It is not possible to configure an Object trace by clients so that information for all connections is recorded.

## Trace configuration file

Each node in the cluster has its own trace configuration file, which is stored in the `/var/mmfs/ces` directory. The configuration file contains settings for logging and for each type of tracing:

### [logging]

**filename**

The name of the log file.

**level** The current logging level, which can be `debug`, `info`, `warning`, `error`, or `critical`.

### [smb]

**defaultloglocation**

The default log location that is used by the `reset` command or when current information is not retrievable.

**defaultloglevel**

The default log level that is used by the `reset` command or when current information is not retrievable.

**traceloglevel**

The log level for tracing.

**maxlogsize**

The maximum size of the log file in kilobytes.

**esttracesize**

The estimated trace size in kilobytes.

### [network]

**numoflogfiles**

The maximum number of log files.

**logfilesize**

The maximum size of the log file in kilobytes.

**esttracesize**

The estimated trace size in kilobytes.

**[object]****defaultloglocation**

The default log location that is used by the reset command or when current information is not retrievable.

**defaultloglevel**

The default log level that is used by the reset command or when current information is not retrievable.

**traceloglevel**

The log level for tracing.

**rsyslogconflocation**

The location of the rsyslog configuration file. Rsyslog is a service that is provided by Red Hat, Inc. that redirects log output. The default location is `/etc/rsyslog.d/00-swift.conf..`

**esttracesize**

The estimated trace size in kilobytes.

**Resetting the trace system**

Before you run a new trace, verify that the trace system is reset for the type of trace that you want to run: SMB, Network, or Object. The examples in the following instructions use the SMB trace system. To reset the trace system, follow these steps:

1. Stop the trace if it is still running.
  - a. Check the trace status to see whether the current trace is stopped on all the nodes:

```
mmprotocoltrace status smb
```

If the trace is still running, stop it:

```
mmprotocoltrace stop smb
```

2. Clear the trace records:

```
mmprotocoltrace clear smb
```

If the command is successful, then you have successfully reset the trace system. Skip to the last step in these instructions.

If the command returns an error message, go to the next step.

**Note:** The command responds with an error message if the trace state of a node is something other than **DONE** or **FAILED**. You can verify the trace state of the nodes by running the **status** command:

```
mmprotocoltrace status smb
```

3. Run the clear command again with the **-f** (force) option.

```
mmprotocoltrace -f clear smb
```

4. After a forced clear, the trace system might still be in an invalid state. Run the reset command. For more information about the command, see the “Using advanced options” on page 19.

```
mmprotocoltrace reset smb
```

5. Check the default values in the trace configuration file to verify that they are correct. To display the values in the trace configuration file, run the config command. For more information about the file, see the “Trace configuration file” on page 17 subtopic.

```
mmprotocoltrace config smb
```

6. The trace system is ready. You can now start a new trace.

### Using advanced options

The **reset** command restores the trace system to the default values that are set in the trace configuration file. The command also performs special actions for each type of trace:

- For an SMB trace, the reset removes any IP-specific configuration files and sets the log level and log location to the default values.
- For a Network trace, the reset stops all dumpcap processes.
- For an Object trace, the reset sets the log level to the default value. It then sets the log location to the default location in the rsyslog configuration file, and restarts the rsyslog service.

The following command resets the SMB trace:

```
mmprotocoltrace reset smb
```

The **status** command with the **-v** (verbose) option provides more trace information, including the values of trace variables. The following command returns verbose trace information for the SMB trace:

```
mmprotocoltrace -v status smb
```

---

## 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 Viewer on Windows.

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

Failures in the error log 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 and look for events with a source label of **GPFS** in the **Application** event category.

On Linux, **syslog** may include GPFS log messages and the error logs described in this section. The **systemLogLevel** attribute of the **mmchconfig** command controls which GPFS log messages are sent to **syslog**. For more information, see the **mmchconfig** command in the *IBM Spectrum Scale: Administration and Programming Reference*.

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

- “MMFS\_ABNORMAL\_SHUTDOWN” on page 20
- “MMFS\_DISKFAIL” on page 20
- “MMFS\_ENVIRON” on page 20
- “MMFS\_FSSTRUCT” on page 20
- “MMFS\_GENERIC” on page 20
- “MMFS\_LONGDISKIO” on page 21
- “MMFS\_QUOTA” on page 21
- “MMFS\_SYSTEM\_UNMOUNT” on page 22
- “MMFS\_SYSTEM\_WARNING” on page 22

## **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**

This topic describes about the **MMFS\_DISKFAIL** error log available in IBM Spectrum Scale.

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**

This topic describes about the **MMFS\_ENVIRON** error log available in IBM Spectrum Scale.

**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**

This topic describes about the **MMFS\_FSSTRUCT** error log available in IBM Spectrum Scale.

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.

If the file system is severely damaged, the best course of action is to follow the procedures in “Additional information to collect for file system corruption or **MMFS\_FSSTRUCT** errors” on page 168, and then contact the IBM Support Center.

## **MMFS\_GENERIC**

This topic describes about **MMFS\_GENERIC** error logs available in IBM Spectrum Scale.

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 15, “Messages,” on page 173 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 be collected before contacting the IBM Support Center” on page 167.

## MMFS\_LONGDISKIO

This topic describes about the MMFS\_LONGDISKIO error log available in IBM Spectrum Scale.

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 in IBM Knowledge Center ([www.ibm.com/support/knowledgecenter/ssw\\_aix/welcome](http://www.ibm.com/support/knowledgecenter/ssw_aix/welcome)) and search for *performance management*. To enable or disable, use the **mmchfs -w** command. For more details, contact the IBM Support Center.

The **mmpmon** command can be used to analyze I/O performance on a per-node basis. See *Failures using the mmpmon command* and the *Monitoring GPFS I/O performance with the mmpmon command* topic in the *IBM Spectrum Scale: Advanced Administration Guide*.

## MMFS\_QUOTA

This topic describes about the MMFS\_QUOTA error log available in IBM Spectrum Scale.

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 105).

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 the IBM white paper entitled *UID Mapping for GPFS in a Multi-cluster Environment* in IBM Knowledge Center ([www.ibm.com/support/knowledgecenter/SSFKCN/com.ibm.cluster.gpfs.doc/gpfs\\_uid/uid\\_gpfs.html](http://www.ibm.com/support/knowledgecenter/SSFKCN/com.ibm.cluster.gpfs.doc/gpfs_uid/uid_gpfs.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

After replacing the appropriate quota file, reissue the **mmcheckquota** command to check the file system inode and space usage.

For information about running the **mmcheckquota** command, see “The mmcheckquota command” on page 57.

## MMFS\_SYSTEM\_UNMOUNT

This topic describes about the MMFS\_SYSTEM\_UNMOUNT error log available in IBM Spectrum Scale.

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

See “File system forced unmount” on page 105. Follow the problem determination and repair actions specified.

## MMFS\_SYSTEM\_WARNING

This topic describes about the MMFS\_SYSTEM\_WARNING error log available in IBM Spectrum Scale.

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. If so, 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 topic describes about an example of an error log entry in IBM Spectrum Scale.

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
```

## Using the `gpfs.snap` command

This topic describes about the usage of `gpfs.snap` command in IBM Spectrum Scale.

Running the `gpfs.snap` command with no options is similar to running `gpfs.snap -a`. It collects data from all nodes in the cluster. This invocation creates a file that is made up of multiple `gpfs.snap` snapshots. The file that is created includes a master snapshot of the node from which the `gpfs.snap` command was invoked and non-master snapshots of each of other nodes in the cluster.

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 gathered by `gpfs.snap` on all platforms”
  - “Data gathered by `gpfs.snap` on AIX” on page 24
  - “Data gathered by `gpfs.snap` on Linux” on page 25
  - “Data gathered by `gpfs.snap` on Windows” on page 25
2. Data that is gathered by `gpfs.snap` only in the case of a master snapshot. See “Data gathered by `gpfs.snap` for a master snapshot” on page 25.

When the `gpfs.snap` command runs with no options, data is collected for each of the enabled protocols. You can turn off the collection of all protocol data and specify the type of protocol information to be collected using the `--protocol` option. For more information, see `gpfs.snap command` in *IBM Spectrum Scale: Administration and Programming Reference*.

The following categories of data is collected:

- Data that is always gathered by `gpfs.snap` on Linux for protocols:
  - “Data gathered for SMB on Linux” on page 26
  - “Data gathered for NFS on Linux” on page 27
  - “Data gathered for Object on Linux” on page 27
  - “Data gathered for CES on Linux” on page 28
  - “Data gathered for authentication on Linux” on page 28
  - “Data gathered for performance on Linux” on page 29

## Data 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:
  - `ls -l /user/lpp/mmfs/bin`
  - `mmdevdiscover`
  - `tspreparedisk -S`
  - `mmfsadm dump malloc`

- `mmfsadm dump fs`
- `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 rpc`
- `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/syslog.conf` or `/etc/syslog-ng.conf`
- `/tmp/mmfs/internal*`
- `/tmp/mmfs/trcrpt*`
- `/var/adm/ras/mmfs.log.*`
- `/var/mmfs/gen/*`
- `/var/mmfs/etc/*`
- `/var/mmfs/tmp/*`
- `/var/mmfs/ssl/*` except for `complete.map` and `id_rsa` files

## Data gathered by `gpfs.snap` on AIX

This topic describes about the type of data that is always gathered by the `gpfs.snap` command on the AIX platform.

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

1. The output of these commands:

- `errpt -a`
- `lssrc -a`
- `lslpp -hac`
- `no -a`

2. The contents of these files:

- `/etc/filesystems`
- `/etc/trcfmt`

## Data gathered by `gpfs.snap` on Linux

This topic describes about the type of data that is always gathered by the `gpfs.snap` command on the Linux platform.

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

1. The output of these commands:

- `dmesg`
- `fdisk -l`
- `lsmod`
- `lspci`
- `rpm -qa`
- `rpm --verify gpfs.base`
- `rpm --verify gpfs.docs`
- `rpm --verify gpfs.gpl`
- `rpm --verify gpfs.msg.en_US`

2. The contents of these files:

- `/etc/filesystems`
- `/etc/fstab`
- `/etc/*release`
- `/proc/cpuinfo`
- `/proc/version`
- `/usr/lpp/mmfs/src/config/site.mcr`
- `/var/log/messages*`

## Data gathered by `gpfs.snap` on Windows

This topic describes about the type of data that is always gathered by the `gpfs.snap` command on the Windows platform.

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 gathered by `gpfs.snap` for a master snapshot

This topic describes about the type of data that is always gathered by the `gpfs.snap` command for a master snapshot.

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 gathered by `gpfs.snap` on all platforms” on page 23, “Data gathered by `gpfs.snap` on AIX” on page 24, “Data gathered by `gpfs.snap` on Linux,” and “Data gathered by `gpfs.snap` on Windows” is obtained, as well as this data:

1. The output of these commands:

- `mmauth`
- `mmgetstate -a`

- **mmlscluster**
- **mmlsconfig**
- **mmlsdisk**
- **mmlsfileset**
- **mmlsfs**
- **mmlspolicy**
- **mmlsmgr**
- **mmlsnode -a**
- **mmlsnsd**
- **mmlssnapshot**
- **mmremotecol**
- **mmremotefs**
- **tsstatus**

2. The contents of the `/var/adm/ras/mmfs.log.*` file (on all nodes in the cluster)

## Data gathered by `gpfs.snap` on Linux for protocols

When the `gpfs.snap` command runs with no options, data is collected for each of the enabled protocols.

You can turn off the collection of all protocol data and specify the type of protocol information to be collected using the `--protocol` option..

## Data gathered for SMB on Linux

The following data is always obtained by the `gpfs.snap` command for the server message block (SMB).

1. The output of these commands:

- **ctdb status**
- **ctdb scriptstatus**
- **ctdb ip**
- **ctdb statistics**
- **ctdb uptime**
- **smbstatus**
- **wbinfo -t**
- **rpm -q gpfs.smb**
- **rpm -q samba**
- **net conf list**
- **sharesec --view-all**
- **mmlsperfdata smb2Throughput -n 1440 -b 60**
- **mmlsperfdata smb2IORate -n 1440 -b 60**
- **mmlsperfdata smb2IOLatency -n 1440 -b 60**
- **ls -l /var/ctdb**
- **ls -l /var/ctdb/persistent**
- **tdbtool info for all .tdb files in /var/ctdb/\***
- **tdbtool check for all .tdb files in /var/ctdb/persistent**

2. The content of these files:

- `/var/adm/ras/log.smbd`
- `/var/lib/samba/*`
- `/var/lib/ctdb/*`

- /etc/sysconfig/gpfs-ctdb
- /var/mmfs/ces/smb.conf
- /var/mmfs/ces/smb.ctdb.nodes

## Data gathered for NFS on Linux

The following data is always obtained by the `gpfs.snap` command for NFS.

1. The output of these commands:
  - **mmnfs export list**
  - **mmnfs configuration list**
  - **rpm -qi - for all installed ganesha packages**
2. The content of these files:
  - /var/mmfs/ces/nfs-config/\*
  - /var/log/ganesha.log
  - /var/tmp/abrt/\*

Files stored in the CCR:

- gpfs.ganesha.exports.conf
- gpfs.ganesha.main.conf
- gpfs.ganesha.nfsd.conf
- gpfs.ganesha.log.conf

## Data gathered for Object on Linux

The following data is always obtained by the `gpfs.snap` command for Object protocol.

1. The output of these commands:
  - **swift info**
  - **rpm -qi - for all installed openstack rpms**
2. The content of these files:
  - /var/log/swift/\*
  - /var/log/keystone/\*
  - /var/log/httpd/\*
  - /var/log/messages
  - /etc/httpd/conf/httpd.conf
  - /etc/httpd/conf.d/ssl.conf
  - /etc/httpd/conf.d/wsgi-keystone.conf
  - All files stored in the directory specified in the `spectrum-scale-objectizer.conf` CCR file in the **objectization\_tmp\_dir** parameter.

Files stored in the CCR:

- account-server.conf
- account.builder
- account.ring.gz
- container-server.conf
- container.builder
- container.ring.gz
- object-server.conf
- object\*.builder
- object\*.ring.gz
- container-reconciler.conf

- swift.conf
- spectrum-scale-compression-scheduler.conf
- spectrum-scale-object-policies.conf
- spectrum-scale-objectizer.conf
- spectrum-scale-object.conf
- object-server-sof.conf
- object-expirer.conf
- keystone-paste.ini
- policy\*.json
- sso/certs/ldap\_cacert.pem
- object-expirer.conf
- object-server-sof.conf
- spectrum-scale-compression-scheduler.conf
- spectrum-scale-compression-status.stat
- spectrum-scale-object.conf
- spectrum-scale-object-policies.conf

## Data gathered for CES on Linux

The following data is always obtained by the `gpfs.snap` command for any enabled protocols.

1. The output of these commands:

- **`mmfscluster --ces`**
- **`mmces node list`**
- **`mmces address list`**
- **`mmces service list -a`**
- **`mmccr flist`**

2. The content of these files:

- `/var/adm/ras/mmcesmonitor.log`
- `/var/adm/ras/mmcesmonitor.log.*`
- `/var/adm/ras/ras.db` (Contents exported as csv file)
- All files stored at the `cesSharedRoot + /ces/connections/`
- All files stored at the `cesSharedRoot + /ces/addr/`

Files stored in the CCR:

- `cesiplist`
- `ccr.nodes`
- `ccr.disks`

## Data gathered for authentication on Linux

The following data is always obtained by the `gpfs.snap` command for any enabled protocol.

1. The output of these commands:

- **`mmcesuserauthlsservice`**
- **`mmcesuserauthckservice --data-access-method all --nodes cesNodes`**
- **`mmcesuserauthckservice --data-access-method all --nodes cesNodes --server-reachability`**
- **`systemctl status ybind`**
- **`systemctl status sssd`**
- **`ps aux | grep keystone`**
- **`lsof -i`**



- **sestatus**
  - **systemctl status firewalld**
  - **systemctl status iptables**
2. The content of these files:
- /etc/nsswitch.conf
  - /etc/ypbind.conf
  - /etc/idmapd.conf
  - /etc/sss/\*
  - /etc/krb5.conf
  - /etc/krb5.keytab
  - /etc/firewalld/\*
  - /etc/openldap/certs/\*
  - /etc/keystone/keystone-paste.ini
  - /etc/keystone/logging.conf
  - /etc/keystone/policy.json
  - /etc/keystone/ssl/certs/\*
  - /var/log/keystone/\*
  - /var/log/sss/\*
  - /var/log/secure/\*
  - /var/log/httpd/\*
  - /etc/httpd/conf/httpd.conf
  - /etc/httpd/conf.d/ssl.conf
  - /etc/httpd/conf.d/wsgi-keystone.conf

Files stored in the CCR:

- NSSWITCH\_CONF
- KEYSTONE\_CONF
- YP\_CONF
- SSSD\_CONF
- LDAP\_TLS\_CACERT
- KS\_SIGNING\_CERT
- KS\_SIGNING\_KEY
- KS\_SIGNING\_CACERT
- KS\_SSL\_CERT
- KS\_SSL\_CACERT
- KS\_LDAP\_CACERT
- authccr

## Data gathered for performance on Linux

The following data is always obtained by the `gdfs.snap` command for any enabled protocols.

1. The output of these commands:
- **top -n 1 -b**
  - **mmdiaq --waiters --iohist --threads --stats --memory**
  - **mmfsadm eventsExporter**
  - **mmpmon chms**
  - **mmfsadm dump nsd**
  - **mmfsadm dump mb**

- `mmdumpperfdata -r 86400`

2. The content of these files:

- `/opt/IBM/zimon/*`
- `/var/log/cnlog/zimon/*`

### **Data gathered for core dumps on Linux**

The following data is gathered when running `gpfs.snap` with the `--protocol core` argument:

- If `core_pattern` is set to dump to a file it will gather files matching that pattern.
- If `core_pattern` is set to redirect to `abrt` then everything is gathered from the directory specified in the `abrt.conf` file under `DumpLocation`. If this is not set then `'/var/tmp/abrt'` is used.
- Other core dump mechanisms are not supported by the script.
- Any files in the directory `'/var/adm/ras/cores/'` will also be gathered.

---

## mmdumpperfdata command

Collects and archives the performance metric information.

### Synopsis

```
mmdumpperfdata [--remove-tree] [StartTime EndTime | Duration]
```

### Availability

Available with IBM Spectrum Scale Standard Edition or higher.

### Description

The **mmdumpperfdata** command runs all named queries and computed metrics used in the **mmpfmon query** command for each cluster node, writes the output into CSV files, and archives all the files in a single .tgz file. The file name is in the `iss_perfdump_YYYYMMDD_hhmmss.tgz` format.

The TAR archive file contains a folder for each cluster node and within that folder there is a text file with the output of each named query and computed metric.

If the start and end time, or duration are not given, then by default the last four hours of metrics information is collected and archived.

### Parameters

**--remove-tree or -r**

Removes the folder structure that was created for the TAR archive file.

*StartTime*

Specifies the start timestamp for query in the YYYY-MM-DD[-hh:mm:ss] format.

*EndTime*

Specifies the end timestamp for query in the YYYY-MM-DD[-hh:mm:ss] format.

*Duration*

Specifies the duration in seconds

### Exit status

**0** Successful completion.

**nonzero**

A failure has occurred.

### Security

You must have root authority to run the **mmdumpperfdata** command.

The node on which the command is issued must be able to execute remote shell commands on any other node in the cluster without the use of a password and without producing any extraneous messages. See the following *IBM Spectrum Scale: Administration and Programming Reference* topic: “Requirements for administering a GPFS file system”.

### Examples

1. To archive the performance metric information collected for the default time period of last four hours and also delete the folder structure that the command creates, issue this command:

```
mmdumpperfdata --remove-tree
```

The system displays output similar to this:

Using the following options:

```
tstart :
tend   :
duration: 14400
rem tree: True
Target folder: ./iss_perfdump_20150513_142420
[1/120] Dumping data for node=fsc-hs21-22 and query q=swiftAccThroughput
file: ./iss_perfdump_20150513_142420/fsc-hs21-22/swiftAccThroughput
[2/120] Dumping data for node=fsc-hs21-22 and query q=NetDetails
file: ./iss_perfdump_20150513_142420/fsc-hs21-22/NetDetails
[3/120] Dumping data for node=fsc-hs21-22 and query q=ctdbCallLatency
file: ./iss_perfdump_20150513_142420/fsc-hs21-22/ctdbCallLatency
[4/120] Dumping data for node=fsc-hs21-22 and query q=usage
file: ./iss_perfdump_20150513_142420/fsc-hs21-22/usage
```

2. To archive the performance metric information collected for a specific time period, issue this command:

```
mmdumperfdata --remove-tree 2015-01-25-04:04:04 2015-01-26-04:04:04
```

The system displays output similar to this:

Using the following options:

```
tstart : 2015-01-25 04:04:04
tend   : 2015-01-26 04:04:04
duration:
rem tree: True
Target folder: ./iss_perfdump_20150513_144344
[1/120] Dumping data for node=fsc-hs21-22 and query q=swiftAccThroughput
file: ./iss_perfdump_20150513_144344/fsc-hs21-22/swiftAccThroughput
[2/120] Dumping data for node=fsc-hs21-22 and query q=NetDetails
file: ./iss_perfdump_20150513_144344/fsc-hs21-22/NetDetails
```

3. To archive the performance metric information collected in the last 200 seconds, issue this command:

```
mmdumperfdata --remove-tree 200
```

The system displays output similar to this:

Using the following options:

```
tstart :
tend   :
duration: 200
rem tree: True
Target folder: ./iss_perfdump_20150513_144426
[1/120] Dumping data for node=fsc-hs21-22 and query q=swiftAccThroughput
file: ./iss_perfdump_20150513_144426/fsc-hs21-22/swiftAccThroughput
[2/120] Dumping data for node=fsc-hs21-22 and query q=NetDetails
file: ./iss_perfdump_20150513_144426/fsc-hs21-22/NetDetails
[3/120] Dumping data for node=fsc-hs21-22 and query q=ctdbCallLatency
file: ./iss_perfdump_20150513_144426/fsc-hs21-22/ctdbCallLatency
[4/120] Dumping data for node=fsc-hs21-22 and query q=usage
file: ./iss_perfdump_20150513_144426/fsc-hs21-22/usage
[5/120] Dumping data for node=fsc-hs21-22 and query q=smb2IORate
file: ./iss_perfdump_20150513_144426/fsc-hs21-22/smb2IORate
[6/120] Dumping data for node=fsc-hs21-22 and query q=swiftConLatency
file: ./iss_perfdump_20150513_144426/fsc-hs21-22/swiftConLatency
[7/120] Dumping data for node=fsc-hs21-22 and query q=swiftCon
file: ./iss_perfdump_20150513_144426/fsc-hs21-22/swiftCon
[8/120] Dumping data for node=fsc-hs21-22 and query q=gpfsNSDWaits
file: ./iss_perfdump_20150513_144426/fsc-hs21-22/gpfsNSDWaits
[9/120] Dumping data for node=fsc-hs21-22 and query q=smb2Throughput
file: ./iss_perfdump_20150513_144426/fsc-hs21-22/smb2Throughput
```

## See also

See also the following *IBM Spectrum Scale: Administration and Programming Reference* topic:

- “**mmperfmon** command”.

## Location

/usr/lpp/mmfs/bin

---

## 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, which could 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 **mmtracectl** 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 **mmfsmgr** command:

```
mmfsmgr
```

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)

---

## Trace facility

The IBM Spectrum Scale system includes many different trace points to facilitate rapid problem determination of failures.

IBM Spectrum Scale tracing is based on the 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 *IBM Spectrum Scale: Administration and Programming Reference*.

## Generating GPFS trace reports

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

To configure and use the trace properly:

1. Issue the **mmchconfig dataStructureDump** command to verify that a directory for dumps was created when the cluster was configured. The default location for trace and problem determination data is **/tmp/mmfs**. Use **mmtracectl**, as instructed by the IBM Support Center, 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 *traceRecycle* 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 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 **mmchconfig** command in Step 1. 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 might ask you to modify the GPFS trace levels. Use the **mmtracectl** command to establish the required trace classes and levels of tracing. The syntax to modify trace classes and levels is as follows:

```
mmtracectl --set --trace={io | all | def | "Class Level [Class Level ...]}
```

For example, to tailor the trace level for I/O, issue the following command:

```
mmtracectl --set --trace=io
```

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
```

Other possible values that can be specified for the trace *Class* include:

**afm**

active file management

**alloc**

disk space allocation

**allocmgr**

allocation manager

**basic**

'basic' classes

**brl**

byte range locks

**cksum**

checksum services

**cleanup**

cleanup routines

**cmd**

ts commands

**defrag**

defragmentation

**dentry**

dentry operations

**dentryexit**

daemon routine entry/exit

**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

**ialloc**  
inode allocation

**io** physical I/O

**kentryexit**  
kernel routine entry/exit

**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

**rdma**  
rdma

**sanergy**  
SANergy®



**scsi**  
scsi services

**sec**  
cluster security

**shared**  
shared segments

**smb**  
SMB locks

**sp** SP message handling

**super**  
super\_operations

**tasking**  
tasking system but not Thread operations

**thread**  
operations in Thread class

**tm** token manager

**ts** daemon specific code

**user1**  
miscellaneous tracing and debugging

**user2**  
miscellaneous tracing and debugging

**vbhv1**  
behaviorals

**vnode**  
vnode layer of VFS kernel support

**vnop**  
one line per VNOP with all important information

Values that can be specified for the trace *Class*, relating to vdisks, include:

**vdb**  
vdisk debugger

**vdisk**  
vdisk

**vhosp**  
vdisk hospital

For more information about vdisks and GPFS Native RAID, see *IBM Spectrum Scale RAID: Administration*.

The trace *Level* can be set to 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.

On AIX, the **-aix-trace-buffer-size** option can be used to control the size of the trace buffer in memory.

On Linux nodes only, use the **mmtracectl** command to change the following:

- The trace buffer size in blocking mode.  
For example, to set the trace buffer size in blocking mode to 8K, issue:  
`mmtracectl --set --tracedev-buffer-size=8K`

- The raw data compression level.

For example, to set the trace raw data compression level to the best ratio, issue:

```
mmtracectl --set --tracedev-compression-level=9
```

- The trace buffer size in overwrite mode.

For example, to set the trace buffer size in overwrite mode to 32K, issue:

```
mmtracectl --set --tracedev-overwrite-buffer-size=32K
```

- When to overwrite the old data.

For example, to wait to overwrite the data until the trace data is written to the local disk and the buffer is available again, issue:

```
mmtracectl --set --tracedev-write-mode=blocking
```

**Note:** Before switching between `--tracedev-write-mode=overwrite` and `--tracedev-write-mode=blocking`, or vice versa, run the `mmtracectl --stop` command first. Next, run the `mmtracectl --set --tracedev-write-mode` command to switch to the desired mode. Finally, restart tracing with the `mmtracectl --start` command.

For more information about the `mmtracectl` command, see the *IBM Spectrum Scale: Administration and Programming Reference*.

---

## Best practices for setting up core dumps on a client system

No core dump configuration is set up by IBM Spectrum Scale by default. Core dumps can be configured in a few ways.

### core\_pattern + ulimit

The simplest way is to change the `core_pattern` file at `/proc/sys/kernel/core_pattern` and to enable core dumps using the command `'ulimit -c unlimited'`. Setting it to something like `/var/log/cores/core.%e.%t.%h.%p` will produce core dumps similar to `core.bash.1236975953.node01.2344` in `/var/log/cores`. This will create core dumps for Linux binaries but will not produce information for Java™ or Python exceptions.

### ABRT

ABRT can be used to produce more detailed output as well as output for Java and Python exceptions.

The following packages should be installed to configure abrt:

- abrt (Core package)
- abrt-cli (CLI tools)
- abrt-libs (Libraries)
- abrt-addon-ccpp (C/C++ crash handler)
- abrt-addon-python (Python unhandled exception handler)
- abrt-java-connector (Java crash handler)

This overwrites the values stored in `core_pattern` to pass core dumps to abrt. It then writes this information to the abrt directory configured in `/etc/abrt/abrt.conf`. Python exceptions is caught by the python interpreter automatically importing the `abrt.pth` file installed in `/usr/lib64/python2.7/site-packages/`. If some custom configuration has changed this behavior, Python dumps may not be created.

To get Java runtimes to report unhandled exceptions through abrt, they must be executed with the command line argument `'-agentpath=/usr/lib64/libabrt-java-connector.so'`.

**Note:** Passing exception information to ABRT by using the ABRT library will cause a decrease in the performance of the application.

### **ABRT Config files**

The ability to collect core dumps has been added to `gpfs.snap` using the `'--protocol core'` option.

This attempts to gather core dumps from a number of locations:

- If `core_pattern` is set to dump to a file it will attempt to get dumps from the absolute path or from the root directory (the CWD for all IBM Spectrum Scale processes)
- If `core_pattern` is set to redirect to `abrt` it will try to read the `/etc/abrt/abrt.conf` file and read the `'DumpLocation'` variable. All files and folders under this directory will be gathered.
- If the `'DumpLocation'` value cannot be read then a default of `'/var/tmp/abrt'` is used.
- If `core_pattern` is set to use something other than `abrt` or a file path, core dumps will not be collected for the OS.

Samba can dump to the directory `'/var/adm/ras/cores/'`. Any files in this directory will be gathered.



---

## Chapter 2. Troubleshooting options available in GUI

You can use logs available in the IBM Spectrum Scale GUI to troubleshoot some issues.

### Events

Use **Monitoring > Events** page in the GUI to monitor the events that are reported in the system. The Events page displays events and you can monitor and troubleshoot errors on your system.

There are three options to filter events by their status:

- **Current Issues** displays all unfixed errors and warnings.
- **Unread Messages** displays all unfixed errors and warnings and information messages that are not marked as read.
- **All Events** displays every event, no matter if it is fixed or marked as read.

The status icons help to quickly determine whether the event is informational, a warning, or an error. Click an event and select **Properties** from the **Action** menu to see detailed information on the event. The event table displays the most recent events first.

### Marking events as Read

You can mark certain events as read to change the status of the event in the events view. The status icons become gray in case an error or warning is fixed or if it is marked as read.

There are events on states that start with "MS\*". These events can be errors, warnings, or information messages that cannot be marked as read and these events automatically change the status from current to historic when the problem is resolved or information condition changes. The user must either fix the problem or change the state of some component to make the current event a historical event. There are also message events that start with MM\*. These events never become historic by themselves. The user must use the action **Mark as Read** on those events to make them historical because the system cannot detect itself even if the problem or information is not valid anymore.

### Running fix procedure

Some issues can be resolved by running a fix procedure. Use action **Run Fix Procedure** to do so. The Events page provides a recommendation for which fix procedure to run next.

### Logs

IBM Support might ask you to collect trace files and dump files from the system to help them resolve a problem. Typically, you perform this task from the management GUI. Use **Settings > Download Logs** page to download logs through GUI.

You can download the following two types of log files:

- GUI log files
- GUI and full IBM Spectrum Scale log files

The GUI log files contain the issues that are related to GUI and it is smaller in size as well. The full log files give details of all kinds of IBM Spectrum Scale issues. The GUI log consists of the following types of information:

- Traces from the GUI that contains the information about errors occurred inside GUI code

- Several configuration files of GUI and PostgreSQL
- Dump of PostgreSQL database that contains IBM Spectrum Scale configuration data and events
- Output of most mmls\* commands
- Logs from the performance collector

The full GUI and IBM Spectrum Scale log file help to analyze all kinds of IBM Spectrum Scale issues. These files are large (gigabytes) and might take an hour to download. You need to select the number of days for which you need to download the log files. These logs files are collected from each individual node. In a cluster with hundreds of nodes, downloading these log files might take a long time and the downloaded file can be large in size. It is recommended to limit the number of days so that it reduces the size of the log file. It is always better to reduce the size of the log file as you might need to send it to IBM Support to help fix the issues.

The issues that are reported in the GUI logs are enough to understand the problem in most of the cases. So, it is recommended to try out the GUI log files first before you download the full log files.

---

## Chapter 3. GPFS cluster state information

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

The information is organized as follows:

- “The **mmafmctl Device getstate** command”
- “The **mmdiag** command”
- “The **mmgetstate** command”
- “The **mmlscluster** command” on page 44
- “The **mmlsconfig** command” on page 45
- “The **mmrefresh** command” on page 45
- “The **mmsdrrestore** command” on page 46
- “The **mmexpelnode** command” on page 46

---

### The **mmafmctl Device getstate** command

The **mmafmctl Device getstate** command displays the status of active file management cache filesets and gateway nodes.

When this command displays a NeedsResync target/fileset state, inconsistencies between home and cache are being fixed automatically; however, unmount and mount operations are required to return the state to Active.

The **mmafmctl Device getstate** command is fully described in the *GPFS Commands* chapter in the *IBM Spectrum Scale: Administration and Programming Reference*.

---

### The **mmdiag** command

The **mmdiag** command displays diagnostic information about the internal GPFS state on the current node.

Use the **mmdiag** command to query various aspects of the GPFS internal state for troubleshooting and tuning purposes. The **mmdiag** command displays information about the state of GPFS on the node where it is executed. The command obtains the required information by querying the GPFS daemon process (**mmfsd**), and thus will only function when the GPFS daemon is running.

The **mmdiag** command is fully described in the *GPFS Commands* chapter in *IBM Spectrum Scale: Administration and Programming Reference*.

---

### The **mmgetstate** command

The **mmgetstate** command displays the state of the GPFS daemon on one or more nodes.

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 or is recovering from an internal error.

**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	k1561nx02	1*	3	7	active	
11	k155n09	1*	3	7	active	

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

The **mmgetstate** command is fully described in the *GPFS Commands* chapter in the *IBM Spectrum Scale: Administration and Programming Reference*.

## The mmlscluster command

The **mmlscluster** command displays GPFS cluster configuration information.

The syntax of the **mmlscluster** 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
Repository type:       server-based
```



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 **mmlscluster** command is fully described in the *GPFS Commands* chapter in the *IBM Spectrum Scale: Administration and Programming Reference*.

---

## The mmlsconfig command

The **mmlsconfig** command displays 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.

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 cl1.cluster:  
-----
```

```
clusterName cl1.cluster  
clusterId 680752107138921233  
autoload no  
minReleaseLevel 4.1.0.0  
pagepool 1G  
maxblocksize 4m  
[c5n97g]  
pagepool 3500m  
[common]  
cipherList EXP-RC4-MD5
```

```
File systems in cluster cl1.cluster:  
-----
```

```
/dev/fs2
```

The **mmlsconfig** command is fully described in the *GPFS Commands* chapter in the *IBM Spectrum Scale: Administration and Programming Reference*.

---

## 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 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 *IBM Spectrum Scale: Administration and Programming Reference*.

For more information, see **mmsdrrestore command** in *IBM Spectrum Scale: Administration and Programming Reference*.

---

## The mmexpelnode command

The **mmexpelnode** command instructs 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 syntax of the **mmexpelnode** command is:

```
mmexpelnode [-o | --once] [-f | --is-fenced] [-w | --wait] -N Node[,Node...]
```

Or,

```
mmexpelnode {-l | --list}
```

Or,

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

The flags used by this command are:

**-o | --once**

Specifies that the nodes should not be prevented from rejoining. After the recovery protocol completes, expelled nodes will be allowed to rejoin the cluster immediately, without the need to first invoke **mmexpelnode --reset**.

**-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). Using this flag allows GPFS to start log recovery immediately, skipping the normal 35-second wait.

**-w | --wait**

Instructs the **mmexpelnode** command to wait until GPFS recovery for the failed node has completed before it runs.

**-l | --list**

Lists all currently expelled nodes.

**-r | --reset**

Allows the specified nodes to rejoin the cluster (that is, resets the status of the nodes). To unexpel all of the expelled nodes, issue: **mmexpelnode -r -N all**.

**-N {all | Node[,Node...]}**

Specifies a list of host names or IP addresses that represent the nodes to be expelled or unexpelled. Specify the daemon interface host names or IP addresses as shown by the **mmcluster** command. The **mmexpelnode** command does not support administration node names or node classes.

**Note:** **-N all** can *only* be used to unexpel nodes.

## 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 4. 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”
- “Read-only mode mount”
- “The lsof command” on page 50
- “The mmlsmount command” on page 50
- “The mmapplypolicy -L command” on page 51
- “The mmcheckquota command” on page 57
- “The mmlsnsd command” on page 57
- “The mmwindisk command” on page 58
- “The mmfileid command” on page 59
- “The SHA digest” on page 61

---

### 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. Follow the procedures in “Information to be collected before contacting the IBM Support Center” on page 167, and then contact 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. (See “Why does the offline mmfsck command fail with “Error creating internal storage?”” on page 147.)
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 **mmmout** 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

The **lsof** (list open files) command returns 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 in Linux distributions or by 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

The **mmlsmount** command lists the nodes that have a given GPFS file system mounted.

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.

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
```

The system displays output similar to:

File system fs2 is mounted on 7 nodes:

192.168.3.53	c25m3n12	c34.cluster	
192.168.110.73	c34f2n01	c34.cluster	
192.168.110.74	c34f2n02	c34.cluster	
192.168.148.77	c12c4apv7	c34.cluster	
192.168.132.123	c20m2n03	c34.cluster	(internal mount)
192.168.115.28	js21n92	c34.cluster	(internal mount)
192.168.3.124	c3m3n14	c3.cluster	

File system fs3 is not mounted.

File system fs3 (c3.cluster:fs3) is mounted on 7 nodes:

192.168.2.11	c2m3n01	c3.cluster	
192.168.2.12	c2m3n02	c3.cluster	
192.168.2.13	c2m3n03	c3.cluster	
192.168.3.123	c3m3n13	c3.cluster	

192.168.3.124	c3m3n14	c3.cluster
192.168.110.74	c34f2n02	c34.cluster
192.168.80.20	c21f1n10	c21.cluster

The **mmismount** command is fully described in the *GPFS Commands* chapter in the *IBM Spectrum Scale: Administration and Programming Reference*.

---

## 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 action.
3	Displays the information for each of the preceding trace levels, plus each candidate file and the applicable rule.
4	Displays the information for each of the preceding trace levels, plus each explicitly excluded file, and the applicable rule.
5	Displays the information for each of the preceding trace levels, plus the attributes of candidate and excluded files.
6	Displays the information for each of the preceding trace levels, plus files that are not candidate files, and their attributes.

These terms are used:

### candidate file

A file that matches a policy rule.

### chosen file

A candidate file that has been scheduled for an action.

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
/* List rule */
RULE EXTERNAL LIST 'tmpfiles' EXEC '/tmp/exec.list'
RULE 'all' LIST 'tmpfiles' where name like '%tmp%'
```

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

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

The **mmapplypolicy** command is fully described in the *GPFS Commands* chapter in the *IBM Spectrum Scale: Administration and Programming Reference*.

## 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:

```
[E:-1] Error while loading policy rules.
PCSQLERR: Unexpected SQL identifier token - 'EXCULDE'.
PCSQLCTX: at line 8 of 8: RULE 'exclude 2' {{{EXCULDE}}}
mmapplypolicy: Command failed. Examine previous error messages to determine cause.
```

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. This option also displays total numbers for file migration and deletion.

This command:

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

produces output similar to this:

```
[I] GPFS Current Data Pool Utilization in KB and %
sp1 5120 19531264 0.026214%
system 102400 19531264 0.524288%
[I] Loaded policy rules from policyfile.
Evaluating MIGRATE/DELETE/EXCLUDE rules with CURRENT_TIMESTAMP = 2009-03-04@02:40:12 UTC
parsed 0 Placement Rules, 0 Restore Rules, 3 Migrate/Delete/Exclude Rules,
1 List Rules, 1 External Pool/List 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%'
/* List rule */
RULE EXTERNAL LIST 'tmpfiles' EXEC '/tmp/exec.list'
RULE 'all' LIST 'tmpfiles' where name like '%tmp%'
[I] Directories scan: 10 files, 1 directories, 0 other objects, 0 'skipped' files and/or errors.
[I] Inodes scan: 10 files, 1 directories, 0 other objects, 0 'skipped' files and/or errors.
[I] Summary of Rule Applicability and File Choices:
Rule# Hit_Cnt KB_Hit Chosen KB_Chosen KB_Ill Rule
0 2 32 0 0 0 RULE 'exclude *.save files' EXCLUDE WHERE(.)
1 2 16 2 16 0 RULE 'delete' DELETE FROM POOL 'sp1' WHERE(.)
2 2 32 2 32 0 RULE 'migration to system pool' MIGRATE FROM POOL \
'sp1' TO POOL 'system' WHERE(.)
3 2 16 2 16 0 RULE 'all' LIST 'tmpfiles' WHERE(.)
```



[I] Files with no applicable rules: 5.

[I] GPFS Policy Decisions and File Choice Totals:  
Chose to migrate 32KB: 2 of 2 candidates;  
Chose to premigrate 0KB: 0 candidates;  
Already co-managed 0KB: 0 candidates;  
Chose to delete 16KB: 2 of 2 candidates;  
Chose to list 16KB: 2 of 2 candidates;  
0KB of chosen data is illplaced or illreplicated;  
Predicted Data Pool Utilization in KB and %:  
sp1 5072 19531264 0.025969%  
system 102432 19531264 0.524451%

## 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:

```
[I] GPFS Current Data Pool Utilization in KB and %  
sp1 5120 19531264 0.026214%  
system 102400 19531264 0.524288%  
[I] Loaded policy rules from policyfile.  
Evaluating MIGRATE/DELETE/EXCLUDE rules with CURRENT_TIMESTAMP = 2009-03-04@02:43:10 UTC  
parsed 0 Placement Rules, 0 Restore Rules, 3 Migrate/Delete/Exclude Rules,  
1 List Rules, 1 External Pool/List 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%'  
/* List rule */  
RULE EXTERNAL LIST 'tmpfiles' EXEC '/tmp/exec.list'  
RULE 'all' LIST 'tmpfiles' where name like '%tmp%'  
[I] Directories scan: 10 files, 1 directories, 0 other objects, 0 'skipped' files and/or errors.  
[I] Inodes scan: 10 files, 1 directories, 0 other objects, 0 'skipped' files and/or errors.  
WEIGHT(INF) LIST 'tmpfiles' /fs1/file.tmp1 SHOW()  
WEIGHT(INF) LIST 'tmpfiles' /fs1/file.tmp0 SHOW()  
WEIGHT(INF) DELETE /fs1/file.tmp1 SHOW()  
WEIGHT(INF) DELETE /fs1/file.tmp0 SHOW()  
WEIGHT(INF) MIGRATE /fs1/file1 TO POOL system SHOW()  
WEIGHT(INF) MIGRATE /fs1/file0 TO POOL system SHOW()  
[I] Summary of Rule Applicability and File Choices:  
Rule# Hit_Cnt KB_Hit Chosen KB_Chosen KB_Ill Rule  
0 2 32 0 0 0 RULE 'exclude *.save files' EXCLUDE WHERE(.)  
1 2 16 2 16 0 RULE 'delete' DELETE FROM POOL 'sp1' WHERE(.)  
2 2 32 2 32 0 RULE 'migration to system pool' MIGRATE FROM POOL \  
'sp1' TO POOL 'system' WHERE(.)  
3 2 16 2 16 0 RULE 'all' LIST 'tmpfiles' WHERE(.)
```

[I] Files with no applicable rules: 5.

[I] GPFS Policy Decisions and File Choice Totals:  
Chose to migrate 32KB: 2 of 2 candidates;  
Chose to premigrate 0KB: 0 candidates;  
Already co-managed 0KB: 0 candidates;  
Chose to delete 16KB: 2 of 2 candidates;  
Chose to list 16KB: 2 of 2 candidates;

```

0KB of chosen data is illplaced or illreplicated;
Predicted Data Pool Utilization in KB and %:
sp1      5072    19531264      0.025969%
system  102432  19531264      0.524451%

```

where the lines:

```

WEIGHT(INF) LIST 'tmpfiles' /fs1/file.tmp1 SHOW()
WEIGHT(INF) LIST 'tmpfiles' /fs1/file.tmp0 SHOW()
WEIGHT(INF) DELETE /fs1/file.tmp1 SHOW()
WEIGHT(INF) DELETE /fs1/file.tmp0 SHOW()
WEIGHT(INF) MIGRATE /fs1/file1 TO POOL system SHOW()
WEIGHT(INF) MIGRATE /fs1/file0 TO POOL system SHOW()

```

show the chosen files and the scheduled 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:

```

[I] GPFS Current Data Pool Utilization in KB and %
sp1      5120    19531264      0.026214%
system  102400  19531264      0.524288%
[I] Loaded policy rules from policyfile.
Evaluating MIGRATE/DELETE/EXCLUDE rules with CURRENT_TIMESTAMP = 2009-03-04@02:32:16 UTC
parsed 0 Placement Rules, 0 Restore Rules, 3 Migrate/Delete/Exclude Rules,
      1 List Rules, 1 External Pool/List 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%'
/* List rule */
RULE EXTERNAL LIST 'tmpfiles' EXEC '/tmp/exec.list'
RULE 'all' LIST 'tmpfiles' where name like '%tmp%'
[I] Directories scan: 10 files, 1 directories, 0 other objects, 0 'skipped' files and/or errors.
/fs1/file.tmp1  RULE 'delete' DELETE FROM POOL 'sp1' WEIGHT(INF)
/fs1/file.tmp1  RULE 'all' LIST 'tmpfiles' WEIGHT(INF)
/fs1/file.tmp0  RULE 'delete' DELETE FROM POOL 'sp1' WEIGHT(INF)
/fs1/file.tmp0  RULE 'all' LIST 'tmpfiles' WEIGHT(INF)
/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)
[I] Inodes scan: 10 files, 1 directories, 0 other objects, 0 'skipped' files and/or errors.
WEIGHT(INF) LIST 'tmpfiles' /fs1/file.tmp1 SHOW()
WEIGHT(INF) LIST 'tmpfiles' /fs1/file.tmp0 SHOW()
WEIGHT(INF) DELETE /fs1/file.tmp1 SHOW()
WEIGHT(INF) DELETE /fs1/file.tmp0 SHOW()
WEIGHT(INF) MIGRATE /fs1/file1 TO POOL system SHOW()
WEIGHT(INF) MIGRATE /fs1/file0 TO POOL system SHOW()
[I] Summary of Rule Applicability and File Choices:
Rule#  Hit_Cnt  KB_Hit  Chosen  KB_Chosen  KB_Ill  Rule
  0     2       32     0       0         0     RULE 'exclude *.save files' EXCLUDE WHERE(.)
  1     2       16     2       16        0     RULE 'delete' DELETE FROM POOL 'sp1' WHERE(.)
  2     2       32     2       32        0     RULE 'migration to system pool' MIGRATE FROM POOL \
    'sp1' TO POOL 'system' WHERE(.)
  3     2       16     2       16        0     RULE 'all' LIST 'tmpfiles' WHERE(.)

[I] Files with no applicable rules: 5.

```

```
[I] GPFS Policy Decisions and File Choice Totals:
Chose to migrate 32KB: 2 of 2 candidates;
Chose to premigrate 0KB: 0 candidates;
Already co-managed 0KB: 0 candidates;
Chose to delete 16KB: 2 of 2 candidates;
Chose to list 16KB: 2 of 2 candidates;
0KB of chosen data is illplaced or illreplicated;
Predicted Data Pool Utilization in KB and %:
sp1      5072   19531264      0.025969%
system 102432 19531264      0.524451%
```

where the lines:

```
/fs1/file.tmp1  RULE 'delete' DELETE FROM POOL 'sp1' WEIGHT(INF)
/fs1/file.tmp1  RULE 'all' LIST 'tmpfiles' WEIGHT(INF)
/fs1/file.tmp0  RULE 'delete' DELETE FROM POOL 'sp1' WEIGHT(INF)
/fs1/file.tmp0  RULE 'all' LIST 'tmpfiles' WEIGHT(INF)
/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)
```

show the 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 the following additional information:

```
[I] Directories scan: 10 files, 1 directories, 0 other objects, 0 'skipped' files and/or errors.
/fs1/file1.save  RULE 'exclude *.save files' EXCLUDE
/fs1/file2.save  RULE 'exclude *.save files' EXCLUDE
/fs1/file.tmp1  RULE 'delete' DELETE FROM POOL 'sp1' WEIGHT(INF)
/fs1/file.tmp1  RULE 'all' LIST 'tmpfiles' WEIGHT(INF)
/fs1/file.tmp0  RULE 'delete' DELETE FROM POOL 'sp1' WEIGHT(INF)
/fs1/file.tmp0  RULE 'all' LIST 'tmpfiles' WEIGHT(INF)
/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)
```

where 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 the following additional information:

```
[I] Directories scan: 10 files, 1 directories, 0 other objects, 0 'skipped' files and/or errors.
/fs1/file1.save [2009-03-03@21:19:57 0 0 16384 sp1 2009-03-04@02:09:38 16 root] RULE 'exclude \
*.save files' EXCLUDE
/fs1/file2.save [2009-03-03@21:19:57 0 0 16384 sp1 2009-03-03@21:19:57 16 root] RULE 'exclude \
*.save files' EXCLUDE
/fs1/file.tmp1 [2009-03-04@02:09:31 0 0 0 sp1 2009-03-04@02:09:31 0 root] RULE 'delete' DELETE \
FROM POOL 'sp1' WEIGHT(INF)
/fs1/file.tmp1 [2009-03-04@02:09:31 0 0 0 sp1 2009-03-04@02:09:31 0 root] RULE 'all' LIST \
'tmpfiles' WEIGHT(INF)
/fs1/file.tmp0 [2009-03-04@02:09:38 0 0 16384 sp1 2009-03-04@02:09:38 16 root] RULE 'delete' \
DELETE FROM POOL 'sp1' WEIGHT(INF)
/fs1/file.tmp0 [2009-03-04@02:09:38 0 0 16384 sp1 2009-03-04@02:09:38 16 root] RULE 'all' \
LIST 'tmpfiles' WEIGHT(INF)
/fs1/file1 [2009-03-03@21:32:41 0 0 16384 sp1 2009-03-03@21:32:41 16 root] RULE 'migration \
to system pool' MIGRATE FROM POOL 'sp1' TO POOL 'system' WEIGHT(INF)
/fs1/file0 [2009-03-03@21:21:11 0 0 16384 sp1 2009-03-03@21:32:41 16 root] RULE 'migration \
to system pool' MIGRATE FROM POOL 'sp1' TO POOL 'system' WEIGHT(INF)
```

where the lines:

```
/fs1/file1.save [2009-03-03@21:19:57 0 0 16384 sp1 2009-03-04@02:09:38 16 root] RULE 'exclude \
*.save files' EXCLUDE
/fs1/file2.save [2009-03-03@21:19:57 0 0 16384 sp1 2009-03-03@21:19:57 16 root] RULE 'exclude \
*.save files' EXCLUDE
```

show the attributes of 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 the following additional information:

```
[I] Directories scan: 10 files, 1 directories, 0 other objects, 0 'skipped' files and/or errors.
/fs1/. [2009-03-04@02:10:43 0 0 8192 system 2009-03-04@02:17:43 8 root] NO RULE APPLIES
/fs1/file1.save [2009-03-03@21:19:57 0 0 16384 sp1 2009-03-04@02:09:38 16 root] RULE \
'exclude *.save files' EXCLUDE
```

```

/fs1/file2.save [2009-03-03@21:19:57 0 0 16384 sp1 2009-03-03@21:19:57 16 root] RULE \
'exclude *.save files' EXCLUDE
/fs1/file.tmp1 [2009-03-04@02:09:31 0 0 0 sp1 2009-03-04@02:09:31 0 root] RULE 'delete' \
DELETE FROM POOL 'sp1' WEIGHT(INF)
/fs1/file.tmp1 [2009-03-04@02:09:31 0 0 0 sp1 2009-03-04@02:09:31 0 root] RULE 'all' LIST \
'tmpfiles' WEIGHT(INF)
/fs1/data1 [2009-03-03@21:20:23 0 0 0 sp1 2009-03-04@02:09:31 0 root] NO RULE APPLIES
/fs1/file.tmp0 [2009-03-04@02:09:38 0 0 16384 sp1 2009-03-04@02:09:38 16 root] RULE 'delete' \
DELETE FROM POOL 'sp1' WEIGHT(INF)
/fs1/file.tmp0 [2009-03-04@02:09:38 0 0 16384 sp1 2009-03-04@02:09:38 16 root] RULE 'all' LIST \
'tmpfiles' WEIGHT(INF)
/fs1/file1 [2009-03-03@21:32:41 0 0 16384 sp1 2009-03-03@21:32:41 16 root] RULE 'migration \
to system pool' MIGRATE FROM POOL 'sp1' TO POOL 'system' WEIGHT(INF)
/fs1/file0 [2009-03-03@21:21:11 0 0 16384 sp1 2009-03-03@21:32:41 16 root] RULE 'migration \
to system pool' MIGRATE FROM POOL 'sp1' TO POOL 'system' WEIGHT(INF)

```

where the line:

```
/fs1/data1 [2009-03-03@21:20:23 0 0 0 sp1 2009-03-04@02:09:31 0 root] NO RULE APPLIES
```

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

---

## The mmcheckquota command

The **mmcheckquota** command counts inode and space usage for a file system and writes the collected data into quota files.

Indications leading you to the conclusion that 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 **mmcheckquota** command is fully described in the *GPFS Commands* chapter in the *IBM Spectrum Scale: Administration and Programming Reference*.

---

## The mmlsnsd command

The **mmlsnsd** command displays information about the currently defined disks in the cluster.

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      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      sdnnsd      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
(fsree 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
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
hd7n97	0972846145C8E934	/dev/hd7n97	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 **mmlsnsd** command is fully described in the *GPFS Commands* chapter in the *IBM Spectrum Scale: Administration and Programming Reference*.

## The mmwindisk command

On Windows nodes, use the **mmwindisk** command to view all disks known to the operating system along with partitioning information relevant to GPFS.

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

Disk	Avail	Type	Status	Size	GPFS Partition ID
0		BASIC	ONLINE	137 GiB	
1		GPFS	ONLINE	55 GiB	362DD84E-3D2E-4A59-B96B-BDE64E31ACCF
2		GPFS	ONLINE	200 GiB	BD5E64E4-32C8-44CE-8687-B14982848AD2
3		GPFS	ONLINE	55 GiB	B3EC846C-9C41-4EFD-940D-1AFA6E2D08FB
4		GPFS	ONLINE	55 GiB	6023455C-353D-40D1-BC EB-FF8E73BF6C0F
5		GPFS	ONLINE	55 GiB	2886391A-BB2D-4BDF-BE59-F33860441262
6		GPFS	ONLINE	55 GiB	00845DCC-058B-4DEB-BD0A-17BAD5A54530
7		GPFS	ONLINE	55 GiB	260BCAEB-6E8A-4504-874D-7E07E02E1817
8		GPFS	ONLINE	55 GiB	863B6D80-2E15-457E-B2D5-FEA0BC41A5AC
9	YES	UNALLOC	OFFLINE	55 GiB	
10	YES	UNALLOC	OFFLINE	200 GiB	

Where:

**Disk**

is the Windows disk number as shown in the Disk Management console and the DISKPART command-line utility.

**Avail**

shows the value **YES** when the disk is available and in a state suitable for creating an NSD.

**GPFS Partition ID**

is the unique ID for the GPFS partition on the disk.

The **mmwindisk** command does not provide the NSD volume ID. You can use **mmlsnsd -m** to find the relationship between NSDs and devices, which are disk numbers on Windows.

---

## The mmfileid command

The **mmfileid** command identifies files that are on areas of a disk that are damaged or suspect.

**Attention:** Use this command only when the IBM Support Center directs you to do so.

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

The command syntax is as follows:

```
mmfileid Device
{-d DiskDesc | -F DescFile}
[-o OutputFile] [-f NumThreads] [-t Directory]
[-N {Node[,Node...]} | NodeFile | NodeClass] [--qos QOSClass]
```

The input parameters are as follows:

*Device*

The device name for the file system.

**-d DiskDesc**

A descriptor that identifies the disk to be scanned. *DiskDesc* has the following format:

```
NodeName:DiskName[:PhysAddr1[-PhysAddr2]]
```

It has the following alternative format:

```
:{NsdName|DiskNum|BROKEN}[:PhysAddr1[-PhysAddr2]]
```

*NodeName*

Specifies a node in the GPFS cluster that has access to the disk to scan. You must specify this value if the disk is identified with its physical volume name. Do not specify this value if the disk is identified with its NSD name or its GPFS disk ID number, or if the keyword **BROKEN** is used.

*DiskName*

Specifies the physical volume name of the disk to scan as known on node *NodeName*.

*NsdName*

Specifies the GPFS NSD name of the disk to scan.

*DiskNum*

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

**BROKEN**

Tells the command to scan all the disks in the file system for files with broken addresses that result in lost data.

### *PhysAddr1* [-*PhysAddr2*]

Specifies the range of physical disk addresses to scan. The default value for *PhysAddr1* is zero. The default value for *PhysAddr2* is the value for *PhysAddr1*.

If both *PhysAddr1* and *PhysAddr2* are zero, the command searches the entire disk.

The following lines are examples of valid disk descriptors:

```
k148n07:hdisk9:2206310-2206810
:gpfs1008nsd:
:10:27645856
:BROKEN
```

### **-F** *DescFile*

Specifies a file that contains a list of disk descriptors, one per line.

### **-f** *NumThreads*

Specifies the number of worker threads to create. The default value is 16. The minimum value is 1. The maximum value is the maximum number allowed by the operating system function **pthread\_create** for a single process. A suggested value is twice the number of disks in the file system.

### **-N** {*Node* [,*Node*...] | *NodeFile* | *NodeClass*}

Specifies the list of nodes that participate in determining the disk addresses. This command supports all defined node classes. The default is **all** or the current value of the **defaultHelperNodes** configuration parameter of the **mmchconfig** command.

For general information on how to specify node names, see the topic “Specifying nodes as input to GPFS commands” in the *IBM Spectrum Scale: Administration and Programming Reference*.

### **-o** *OutputFile*

The path name of a file to which the result from the **mmfileid** command is to be written. If not specified, the result is sent to standard output.

### **-t** *Directory*

Specifies the directory to use for temporary storage during **mmfileid** command processing. The default directory is **/tmp**.

### **--qos** *QoSClass*

Specifies the Quality of Service for I/O operations (QoS) class to which the instance of the command is assigned. If you do not specify this parameter, the instance of the command is assigned by default to the **maintenance** QoS class. This parameter has no effect unless the QoS service is enabled. For more information, see the help topic on the **mmchqos** command in the *IBM Spectrum Scale: Administration and Programming Reference*. Specify one of the following QoS classes:

#### **maintenance**

This QoS class is typically configured to have a smaller share of file system IOPS. Use this class for I/O-intensive, potentially long-running GPFS commands, so that they contribute less to reducing overall file system performance.

#### **other**

This QoS class is typically configured to have a larger share of file system IOPS. Use this class for administration commands that are not I/O-intensive.

For more information, see the help topic on setting the Quality of Service for I/O operations (QoS) in the *IBM Spectrum Scale: Administration and Programming Reference*.

You can redirect the output to a file with the **-o** flag and sort the output on the inode number with the **sort** command.

The **mmfileid** command output contains one line for each inode found to be on a corrupted disk sector. Each line of the command output has this format:

```
InodeNumber LogicalDiskAddress SnapshotId Filename
```



### *InodeNumber*

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

### *LogicalDiskAddress*

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

### *SnapshotId*

Indicates the snapshot identifier for the file. A *SnapshotId* of 0 means that the file is not a snapshot file.

### *Filename*

Indicates the name of the file identified by **mmfileid**. File names are relative to the root of the file system in which they reside.

Assume that a disk analysis tool reports that disks `hdisk6`, `hdisk7`, `hdisk8`, and `hdisk9` contain bad sectors, and that the file `addr.in` has the following contents:

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

You run the following command:

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

The command output might be similar to the following example:

```
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 that begin with the word `Address` represent GPFS system metadata files or reserved disk areas. If your output contains any lines like these, do not attempt to replace or repair the indicated files. If you suspect that any of the special files are damaged, call the IBM Support Center for assistance.

The following line of output indicates that inode number 14336, disk address 1072256 contains file **/gpfsB/tesDir/testFile.out**. The 0 to the left of the name indicates that the file does not belong to a snapshot. This file is on a potentially bad disk sector area:

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

The following line of output indicates that inode number 14344, disk address 2922528 contains file **/gpfsB/x.img**. The 1 to the left of the name indicates that the file belongs to snapshot number 1. This file is on a potentially bad disk sector area:

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

---

## 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 **mmremotecenter** 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 **mmremotecoluster** 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 **mmremotecoluster 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 the output produced by the **mmauth show all** 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 5. Resolving deadlocks

- | IBM Spectrum Scale provides functions for deadlock detection, deadlock data collection, deadlock breakup, and cluster overload protection.

The distributed nature of GPFS, the complexity of the locking infrastructure, the dependency on the proper operation of disks and networks, and the overall complexity of operating in a clustered environment all contribute to increasing the probability of a deadlock.

Deadlocks can be disruptive in certain situations, more so than other type of failure. A deadlock effectively represents a single point of failure that can render the entire cluster inoperable. When a deadlock is encountered on a production system, it can take a long time to debug. The typical approach to recovering from a deadlock involves rebooting all of the nodes in the cluster. Thus, deadlocks can lead to prolonged and complete outages of clusters.

To troubleshoot deadlocks, you must have specific types of debug data that must be collected while the deadlock is in progress. Data collection commands must be run manually before the deadlock is broken. Otherwise, determining the root cause of the deadlock after that is difficult. Also, deadlock detection requires some form of external action, for example, a complaint from a user. Waiting for a user complaint means that detecting a deadlock in progress might take many hours.

- | In GPFS V4.1 and later, automated deadlock detection, automated deadlock data collection, deadlock breakup options, and cluster overload detection are provided to make it easier to handle a deadlock situation.
- | • “Automated deadlock detection”
- | • “Automated deadlock data collection” on page 65
- | • “Automated deadlock breakup” on page 66
- | • “Deadlock breakup on demand” on page 67
- | • “Cluster overload detection” on page 68

---

### Automated deadlock detection

Many deadlocks involve long waiters; for example, **mmfsd** threads that have been waiting for some event for a considerable duration of time. With some exceptions, long waiters typically indicate that something in the system is not healthy. There may be a deadlock in progress, some disk may be failing, or the entire system may be overloaded.

All waiters can be broadly divided into four categories:

- Waiters that can occur under normal operating conditions and can be ignored by automated deadlock detection.
- Waiters that correspond to complex operations and can legitimately grow to moderate lengths.
- Waiters that should never be long. For example, most mutexes should only be held briefly.
- Waiters that can be used as an indicator of cluster overload. For example, waiters waiting for I/O completions or network availability.

Automated deadlock detection monitors waiters. Deadlock detection relies on a configurable threshold to determine if a deadlock is in progress. When a deadlock is detected, an alert is issued in the **mmfs.log**, the operating system log, and the **deadlockDetected** callback is triggered.

Automated deadlock detection is enabled by default and controlled with the **mmchconfig** attribute **deadlockDetectionThreshold**. A potential deadlock is detected when a waiter waits longer than **deadlockDetectionThreshold**. To view the current threshold for deadlock detection, enter the following command:

```
mmfsconfig deadlockDetectionThreshold
```

The system displays output similar to the following:

```
deadlockDetectionThreshold 300
```

To disable automated deadlock detection, specify a value of 0 for the **deadlockDetectionThreshold** attribute.

To simplify the process of monitoring for deadlocks, a user callback program can be registered with **mmaddcallback** for the **deadlockDetected** event. This program can be used for recording and notification purposes. When a suspected deadlock is detected, the **deadlockDetected** event is triggered, and the user callback program is run. See the `/usr/lpp/mmfs/samples/deadlockdetected.sample` file for an example of using the **deadlockDetected** event.

The following messages, related to deadlock detection, might be found in the **mmfs.log** files:

```
Enabled automated deadlock detection.
```

```
[A] Deadlock detected: 2015-03-04 02:06:21: waiting 301.291 seconds on node c937f3n04-40g:  
PutACLHandlerThread 2449: on MsgRecordCondvar, reason 'RPC wait' for tmMsgTellAcquire1
```

When a Deadlock detected message is found, it means that a long waiter exceeded the deadlock detection threshold and is suspected to be a deadlock. It takes time to know with certainty if a long waiter is an actual deadlock or not. A real deadlock will not disappear after waiting for a longer period, but a false-positive deadlock can disappear. When selecting a **deadlockDetectionThreshold** value, there is a trade-off between waiting too long and not having timely detection of deadlocks and not waiting long enough causing false-positive deadlock detection. If a false-positive deadlock is detected, a message similar to the following might be found in the **mmfs.log** files:

```
Wed Mar 4 02:11:53.220 2015: [N] Long waiters have disappeared.
```

In addition to the messages found in **mmfs.log** files, the **mmdiag --deadlock** command can be used to query the suspected deadlock waiters currently on a node. Only the longest waiters that are suspected deadlocks are shown. Legitimately long waiters that are ignored by deadlock detection are not shown, but those waiters are shown in the **mmdiag --waiters** section. Other waiters, which are much shorter than the longest deadlock waiters, are not shown because they are typically not relevant (even if their waiter length exceeds the deadlock detection threshold).

The `/var/log/messages` files on Linux and the error report on AIX also have information relevant for deadlock detection, but most details are only shown in the **mmfs.log** files.

While **deadlockDetectionThreshold** is for medium length waiters that can grow to moderate lengths, **deadlockDetectionThresholdForShortWaiters** is for short waiters that should never be long. Waiters that can be legitimately long under normal operating conditions are ignored by automated deadlock detection, for example:

```
TSDELDISKCmdThread: on ThCond 0x1127916B8 (0x1127916B8) (InodeScanStatCondvar),  
reason 'Waiting for PIT worker threads to finish'
```

```
0x3FFDC00ADE8 waiting 4418.790093653 seconds, FsckStaticThread: on ThCond 0x3FFB0011AB8 (0x3FFB0011AB8)  
(FsckStaticThreadCondvar), reason 'Waiting for static fsck work'
```

When many false-positive deadlocks are detected in a cluster (and the long waiters disappear soon after detection), the cluster needs to be checked for a hardware, network, or workload issue. If these issues are not found, then the deadlock detection threshold can be adjusted to avoid routinely detecting false-positive deadlocks.

When you adjust the deadlock detection threshold, you can disable automated deadlock data collection to avoid collecting debug data unnecessarily. Run the workload for a while to determine the longest waiter length detected as a false-positive deadlock. Use that length to determine a better value for **deadlockDetectionThreshold**. You can also try increasing the **deadlockDetectionThreshold** a few times until no more false-positive deadlocks are detected. If you disabled automated deadlock data collection while you were adjusting the threshold, enable it again after the adjustments are complete.

Deadlock amelioration functions should only be used on a stable GPFS cluster to avoid extraneous messages in the **mmfs.log** files and unnecessary debug data collection. If a cluster is not stable, deadlock detection should be disabled.

All deadlock amelioration functions, not just deadlock detection, are disabled by specifying 0 for **deadlockDetectionThreshold**. A positive value must be specified for **deadlockDetectionThreshold** to enable any part of the deadlock amelioration functions.

Deadlock amelioration functions are supported in a multi-cluster environment. When a deadlock is detected, debug data is collected on all local nodes and all non-local nodes that joined the cluster by mounting a local file system. The cluster overload notification applies to such non-local nodes as well. For more information about cluster overload, see “Cluster overload detection” on page 68.

---

## Automated deadlock data collection

In order to effectively troubleshoot a typical deadlock, it is imperative that the following debug data is collected:

- A full internal dump (**mmfsadm dump all**)
- A dump of kthreads (**mmfsadm dump kthreads**)
- Trace data (10-30 seconds of trace data)

Automated deadlock data collection can be used to help gather this crucial debug data on detection of a potential deadlock.

Automated deadlock data collection is enabled by default and controlled with the **mmchconfig** attribute **deadlockDataCollectionDailyLimit**. The **deadlockDataCollectionDailyLimit** attribute specifies the maximum number of times debug data can be collected in a 24-hour period. To view the current data collection interval, enter the following command:

```
mmfsconfig deadlockDataCollectionDailyLimit
```

The system displays output similar to the following:

```
deadlockDataCollectionDailyLimit 10
```

To disable automated deadlock data collection, specify a value of 0 for **deadlockDataCollectionDailyLimit**.

**Note:** The 24-hour period for **deadlockDataCollectionDailyLimit** is enforced passively. When there is a need to collect debug data, the **deadlockDataCollectionDailyLimit** is examined to determine whether 24 hours passed since the beginning of this period and whether a new period for **deadlockDataCollectionDailyLimit** needs to be started or not. If the number of debug data collections exceeds the **deadlockDataCollectionDailyLimit** value before the period reaches 24 hours, then no debug data will be collected until the next period starts. Sometimes exceptions are made to help capture the

most relevant debug data. There should be enough disk space available for debug data collection, and old debug data needs to be moved intermittently to make space for new debug data.

Another **mmchconfig** attribute, **deadlockDataCollectionMinInterval**, can be used to control the amount of time between consecutive debug data collections. The default is 300 seconds because debug data collected within 5 minutes already covered the start of a newly detected deadlock that is 5 minutes or longer.

The following messages, related to deadlock data collection, might be found in the **mmfs.log** files:

```
[I] Enabled automated deadlock debug data collection.
```

```
[N] sdrServ: Received deadlock notification from 192.168.116.56  
[N] GPFS will attempt to collect debug data on this node.
```

```
[N] Debug data is not collected. deadlockDataCollectionDailyLimit 10 has been exceeded.
```

```
[N] Debug data has not been collected. It was collected recently at 2014-01-29 12:58:00.
```

Trace data is part of the debug data that is collected when a suspected deadlock is detected. However, on a typical customer system, GPFS tracing is not routinely turned on. In this case, the automated debug data collection turns on tracing, waits for 20 seconds, collects the trace, and turns off tracing. The 20 seconds of trace will not cover the formation of the deadlock, but it might still provide some helpful debug data.

---

## Automated deadlock breakup

Automated deadlock breakup helps resolve a deadlock situation without human intervention. To break up a deadlock, less disruptive actions are tried first; for example, causing a file system panic. If necessary, more disruptive actions are then taken; for example, shutting down a GPFS **mmfsd** daemon.

If a system administrator prefers to control the deadlock breakup process, the **deadlockDetected** callback can be used to notify system administrators that a potential deadlock was detected. The information from the **mmdiag --deadlock** section can then be used to help determine what steps to take to resolve the deadlock.

Automated deadlock breakup is disabled by default and controlled with the **mmchconfig** attribute **deadlockBreakupDelay**. The **deadlockBreakupDelay** attribute specifies how long to wait after a deadlock is detected before attempting to break up the deadlock. Enough time must be provided to allow the debug data collection to complete. To view the current breakup delay, enter the following command:

```
mmfsconfig deadlockBreakupDelay
```

The system displays output similar to the following:

```
deadlockBreakupDelay 0
```

The value of 0 shows that automated deadlock breakup is disabled. To enable automated deadlock breakup, specify a positive value for **deadlockBreakupDelay**. If automated deadlock breakup is to be enabled, a delay of 300 seconds or longer is recommended.

Automated deadlock breakup is done on a node-by-node basis. If automated deadlock breakup is enabled, the breakup process is started when the suspected deadlock waiter is detected on a node. The process first waits for the **deadlockBreakupDelay**, and then goes through various phases until the deadlock waiters disappear. There is no central coordination on the deadlock breakup, so the time to take deadlock breakup actions may be different on each node. Breaking up a deadlock waiter on one node can cause some deadlock waiters on other nodes to disappear, so no breakup actions need to be taken on those other nodes.

If a suspected deadlock waiter disappears while waiting for the **deadlockBreakupDelay**, the automated deadlock breakup process stops immediately without taking any further action. To lessen the number of breakup actions that are taken in response to detecting a false-positive deadlock, increase the **deadlockBreakupDelay**. If you decide to increase the **deadlockBreakupDelay**, a deadlock can potentially exist for a longer period.

If your goal is to break up a deadlock as soon as possible, and your workload can afford an interruption at any time, then enable automated deadlock breakup from the beginning. Otherwise, keep automated deadlock breakup disabled to avoid unexpected interruptions to your workload. In this case, you can choose to break the deadlock manually, or use the function that is described in the “Deadlock breakup on demand” topic.

| Due to the complexity of the GPFS code, asserts or segmentation faults might happen during a deadlock  
| breakup action. That might cause unwanted disruptions to a customer workload still running normally  
| on the cluster. A good reason to use deadlock breakup on demand is to not disturb a partially working  
| cluster until it is safe to do so. Try not to break up a suspected deadlock prematurely to avoid  
| unnecessary disruptions. If automated deadlock breakup is enabled all of the time, it is good to set  
| **deadlockBreakupDelay** to a large value such as 3600 seconds. If using **mmcommon breakDeadlock**, it is  
| better to wait until the longest deadlock waiter is an hour or longer. Much shorter times can be used if a  
| customer prefers fast action in breaking a deadlock over assurance that a deadlock is real.

The following messages, related to deadlock breakup, might be found in the **mmfs.log** files:

```
[I] Enabled automated deadlock breakup.  
[N] Deadlock breakup: starting in 300 seconds  
[N] Deadlock breakup: aborting RPC on 1 pending nodes.  
[N] Deadlock breakup: panicking fs fs1  
[N] Deadlock breakup: shutting down this node.  
[N] Deadlock breakup: the process has ended.
```

---

## Deadlock breakup on demand

Deadlocks can be broken up on demand, which allows a system administrator to choose the appropriate time to start the breakup actions.

A deadlock can be localized, for example, it might involve only one of many file systems in a cluster. The other file systems in the cluster can still be used, and a mission critical workload might need to continue uninterrupted. In these cases, the best time to break up the deadlock is after the mission critical workload ends.

The **mmcommon** command can be used to break up an existing deadlock in a cluster when the deadlock was previously detected by deadlock amelioration. To start the breakup on demand, use the following syntax:

```
mmcommon breakDeadlock [-N {Node[,Node...] | NodeFile | NodeClass}]
```

If the **mmcommon breakDeadlock** command is issued without the **-N** parameter, then every node in the cluster receives a request to take action on any long waiter that is a suspected deadlock.

If the **mmcommon breakDeadlock** command is issued with the **-N** parameter, then only the nodes that are specified receive a request to take action on any long waiter that is a suspected deadlock. For example, assume that there are two nodes, called **node3** and **node6**, that require a deadlock breakup. To send the breakup request to just these nodes, issue the following command:

```
mmcommon breakDeadlock -N node3,node6
```

Shortly after running the **mmcommon breakDeadlock** command, issue the following command:

```
mmdsh -N all /usr/lpp/mmfs/bin/mmdiag --deadlock
```

The output of the **mmdsh** command can be used to determine if any deadlock waiters still exist and if any additional actions are needed.

The effect of the **mmcommon breakDeadlock** command only persists on a node until the longest deadlock waiter that was detected disappears. All actions that are taken by **mmcommon breakDeadlock** are recorded in the **mmfs.log** file. When **mmcommon breakDeadlock** is issued for a node that did not have a deadlock, no action is taken except for recording the following message in the **mmfs.log** file:

```
[N] Received deadlock breakup request from 192.168.40.72: No deadlock to break up.
```

The **mmcommon breakDeadlock** command provides more control over breaking up deadlocks, but multiple breakup requests might be required to achieve satisfactory results. All waiters that exceeded the **deadlockDetectionThreshold** might not disappear when **mmcommon breakDeadlock** completes on a node. In complicated deadlock scenarios, some long waiters can persist after the longest waiters disappear. Waiter length can grow to exceed the **deadlockDetectionThreshold** at any point, and waiters can disappear at any point as well. Examine the waiter situation after **mmcommon breakDeadlock** completes to determine whether the command must be repeated to break up the deadlock.

Another way to break up a deadlock on demand is to enable automated deadlock breakup by changing **deadlockBreakupDelay** to a positive value. By enabling automated deadlock breakup, breakup actions are initiated on existing deadlock waiters. The breakup actions repeat automatically if deadlock waiters are detected. Change **deadlockBreakupDelay** back to 0 when the results are satisfactory, or when you want to control the timing of deadlock breakup actions again. If automated deadlock breakup remains enabled, breakup actions start on any newly detected deadlocks without any intervention.

---

## Cluster overload detection

A healthy workload is a workload that is operating within the resource capacity of a cluster. Overload is a condition where the cluster does not have enough available resources for a workload. An overloaded cluster can cause slow response times and render a cluster unusable in severe cases. A GPFS cluster that is working correctly should not be overloaded. An overload condition must be avoided to keep a cluster productive.

A cluster overload condition does not affect how GPFS works outside of the deadlock amelioration functions. However, cluster overload detection and notification can be used for monitoring hardware, network, or workload conditions to help maintain a healthy production cluster.

Cluster overload detection is enabled by default and controlled with the **mmchconfig** attribute **deadlockOverloadThreshold**. The **deadlockOverloadThreshold** attribute can be adjusted to ensure that overload conditions are detected according to the criteria you set, instead of reporting overload conditions that you can tolerate. To view the current threshold for cluster overload detection, enter the following command:

```
mmfsconfig deadlockOverloadThreshold
```

The system displays output similar to the following:

```
deadlockOverloadThreshold 5
```

To disable cluster overload detection, specify a value of 0 for the **deadlockOverloadThreshold** attribute.

To simplify the process of monitoring for a cluster overload condition, a user callback program can be registered with **mmaddcallback** for the **deadlockOverload** event. This program can be used for recording and notification purposes. Whenever a node detects an overload condition, the **deadlockOverload** event is triggered, and the user callback program is run.



Deadlock amelioration uses certain I/O completion and network communication waiters, heuristically, to indicate a cluster overload condition. These types of waiters are used for overload evaluation because, even in the range of a few seconds, they might cause arbitrarily long waiters of many other kinds. Cluster overload is not detected unnecessarily, for example, when a single I/O completion waiter is long for a short period.

When a node detects an overload condition, it notifies all nodes in the cluster that the cluster is now overloaded. The notification process uses the cluster manager and the **gpfsNotifyOverload** event. Overload is a cluster-wide condition because all the nodes in a cluster work together, and long waiters on one node can affect other nodes in the cluster. To reduce network traffic, each node checks whether the overload condition should be cleared or not. After a node does not detect an overload condition and is not informed that the cluster is still overloaded, each node will mark the cluster as no longer overloaded after a short period.

The following messages, related to cluster overload, might be found in the **mmfs.log** files:

```
[W] Warning: cluster myCluster may become overloaded soon.
[W] Warning: cluster myCluster is now overloaded.
[I] Forwarding 'overloaded' status to cluster manager myClusterMgr of cluster myCluster
[I] This node is the cluster manager of Cluster myCluster, sending 'overloaded' status to the entire cluster
[N] Received cluster overload notification from 192.168.148.18
[N] Cluster myCluster is no longer overloaded.
```

When a cluster is overloaded, the **mmchconfig** attribute **deadlockDetectionThresholdIfOverloaded** is used to detect suspected deadlocks instead of **deadlockDetectionThreshold**. The default value for the **deadlockDetectionThresholdIfOverloaded** attribute is 1800 seconds because all waiters might become much longer in an overloaded cluster. To view the current threshold for deadlock detection in an overloaded cluster, enter the following command:

```
mmfsconfig deadlockDetectionThresholdIfOverloaded
```

The system displays output similar to the following:

```
deadlockDetectionThresholdIfOverloaded 1800
```

If automated deadlock breakup is enabled, it is disabled temporarily until the overload condition is cleared. This process avoids unnecessary breakup actions when a false-positive deadlock is detected.



---

## Chapter 6. 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 AIX in IBM Knowledge Center ([www.ibm.com/support/knowledgecenter/ssw\\_aix/welcome](http://www.ibm.com/support/knowledgecenter/ssw_aix/welcome)) and search for the appropriate kernel debugging documentation 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 the Microsoft website ([www.microsoft.com](http://www.microsoft.com)), might be useful in troubleshooting:

- Debugging Tools for Windows
- Process Monitor
- Process Explorer
- Microsoft Windows Driver Kit
- Microsoft Windows Software Development 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 fully described in the *GPFS Commands* chapter in the *IBM Spectrum Scale: Administration and Programming Reference*. For details on the **mmpmon** command, see the *Monitoring GPFS I/O performance with the mmpmon command* topic in the *IBM Spectrum Scale: Administration and Programming Reference*.



---

## Chapter 7. Installation and configuration issues

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

An IBM Spectrum Scale installation problem should be suspected when GPFS modules are not loaded successfully, 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 IBM Spectrum Scale 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”
- “GPFS modules cannot be loaded on Linux” on page 79
- “GPFS daemon will not come up” on page 79
- “GPFS daemon went down” on page 83
- “IBM Spectrum Scale failures due to a network failure” on page 84
- “Kernel panics with a 'GPFS dead man switch timer has expired, and there's still outstanding I/O requests' message” on page 85
- “Quorum loss” on page 85
- “Delays and deadlocks” on page 86
- “Node cannot be added to the GPFS cluster” on page 87
- “Remote node expelled after remote file system successfully mounted” on page 87
- “Disaster recovery issues” on page 88
- “GPFS commands are unsuccessful” on page 89
- “Application program errors” on page 91
- “Troubleshooting Windows problems” on page 92
- “OpenSSH connection delays” on page 93

---

### Installation and configuration problems

This topic describes about the issues that you might encounter while installing or configuring IBM Spectrum Scale.

The *IBM Spectrum Scale: Concepts, Planning, and Installation Guide* provides the step-by-step procedure for installing and migrating IBM Spectrum Scale, 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 version. Did you reboot your IBM Spectrum Scale nodes before and after the installation/upgrade of IBM Spectrum Scale? If you did, see “GPFS daemon will not come up” on page 79. If not, reboot. For more information, see the *Initialization of the GPFS daemon* topic in the *IBM Spectrum Scale: Concepts, Planning, and Installation Guide*.

- Not being able to access a file system. See “File system will not mount” on page 95.
- New GPFS functions do not operate. See “GPFS commands are unsuccessful” on page 89.

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

This topic describes about the steps that you need to perform when a cluster crashes after IBM Spectrum Scale reinstallation.

After reinstalling IBM Spectrum Scale 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 -p primaryServer
```

where *primaryServer* is the name of the primary GPFS cluster configuration server.

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 users exist 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 77.

## Problems with the `/etc/hosts` file

This topic describes about the issues relating to the `/etc/hosts` file that you might come across while installing or configuring IBM Spectrum Scale.

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

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

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

## Linux configuration considerations

This topic describes about the Linux configuration that you need to consider while installing or configuring IBM Spectrum Scale on your cluster.

**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. 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.

## 3. What are the performance tuning suggestions?

For an up-to-date list of tuning suggestions, see the IBM Spectrum Scale FAQ in IBM Knowledge Center ([www.ibm.com/support/knowledgecenter/STXKQY/gpfsclustersfaq.html](http://www.ibm.com/support/knowledgecenter/STXKQY/gpfsclustersfaq.html)).

For Linux on z Systems, see also the Device Drivers, Features, and Commands ([www.ibm.com/support/knowledgecenter/api/content/linuxonibm/liaaf/lnz\\_r\\_dd.html](http://www.ibm.com/support/knowledgecenter/api/content/linuxonibm/liaaf/lnz_r_dd.html)) topic in the Linux on z Systems library overview.

## Protocol authentication problem determination

You can use a set of GPFS commands to identify and rectify issues that are related to authentication configurations.

To do basic authentication problem determination, perform the following steps:

1. Issue the **mmces state show auth** command to view the current state of authentication.
2. Issue the **mmces events active auth** command to see whether events are currently contributing to make the state of the authentication component unhealthy.
3. Issue the **mmuserauth service list** command to view the details of the current authentication configuration.
4. Issue the **mmuserauth service check -N cesNodes --server-reachability** command to verify the state of the authentication configuration across the cluster.
5. Issue the **mmuserauth service check -N cesNodes --rectify** command to rectify the authentication configuration.

**Note:** Server reachability cannot be rectified by using the **--rectify** parameter.

## Problems with running commands on other nodes

This topic describes about the problems that you might encounter relating to running remote commands during installing and configuring IBM Spectrum Scale.

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/ssh** and **/usr/bin/scp**. You also have the option of specifying your own remote shell and remote file copy commands to be used instead of the default **ssh** and **scp**. The remote shell and copy commands must adhere to the same syntax forms as **ssh** and **scp** 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

This topic describes about issues with running remote commands due to authorization problems in IBM Spectrum Scale.

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

| **Note:** Use the **ssh** and **scp** commands that are shipped with the OpenSSH package supported by GPFS.  
| Refer to the IBM Spectrum Scale FAQ in IBM Knowledge Center ([www.ibm.com/support/knowledgecenter/STXKQY/gpfsclustersfaq.html](http://www.ibm.com/support/knowledgecenter/STXKQY/gpfsclustersfaq.html)) for the latest OpenSSH information.

For the **ssh** and **scp** 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 **ssh** and **scp** 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

This topic describes about the issues with running GPFS commands on remote nodes due to connectivity problems.

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

```
mmde1node -N k145n04
Verifying GPFS is stopped on all affected nodes ...
mmdsh: 6027-1617 There are no available nodes on which to run the command.
mmde1node: 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

This topic describes about the error messages that are displayed for rsh issues in IBM Spectrum Scale.

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

This topic describes about the issues relating to IBM Spectrum Scale cluster configuration data.



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/gen/mmLockDir` 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 *requiredCondition*.

To release the lock:

1. Determine the PID and the system that owns the lock by issuing:

```
mmcommon showLocks
```

The **mmcommon showLocks** command displays information about the lock server, lock name, lock holder, PID, and extended information. If a GPFS administration command is not responding, stopping the command will free the lock. If another process has this 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.

2. If any locks are held and you want to release them manually, from any node in the GPFS cluster issue the command:

```
mmcommon freeLocks <lockName>
```

## GPFS error messages for cluster configuration data file problems

This topic describes about the error messages relating to the cluster configuration data file issues in IBM Spectrum Scale.

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 configuration 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

This topic describes about the procedure for recovering the cluster configuration data file in IBM Spectrum Scale.

A copy of the IBM Spectrum Scale 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 IBM Spectrum Scale 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 IBM Spectrum Scale 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 IBM Spectrum Scale cluster configuration server, restore the file by issuing this command from the node on which it is missing:

```
mmsdrrestore -p primaryServer
```

where *primaryServer* is the name of the primary GPFS cluster configuration server.

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

This topic describes about the procedure for automatically backing up the cluster data in IBM Spectrum Scale.

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 *GPFS user exits* topic in the *IBM Spectrum Scale: Administration and Programming Reference*.

## Error numbers specific to GPFS applications calls

This topic describes about the error numbers specific to GPFS application 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 IBM Spectrum Scale FAQ in IBM Knowledge Center ([www.ibm.com/support/knowledgecenter/STXKQY/gpfsclustersfaq.html](http://www.ibm.com/support/knowledgecenter/STXKQY/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 9, “Disk issues,” on page 127 for unavailable disks.

**6027-341 [D]**

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

**6027-342 [E]**

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

**6027-343 [E]**

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

**6027-344 [E]**

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

---

## GPFS modules cannot be loaded on Linux

You must build the GPFS portability layer binaries based on the kernel configuration of your system. For more information, see *The GPFS open source portability layer* topic in the *IBM Spectrum Scale: Concepts, Planning, and Installation Guide*. During **mmstartup** processing, GPFS loads the **mmfslinux** kernel module.

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

1. If the portability layer is not built, you may see messages similar to:

```
Mon Mar 26 20:56:30 EDT 2012: runmmfs starting
Removing old /var/adm/ras/mmfs.log.* files:
Unloading modules from /lib/modules/2.6.32.12-0.6-ppc64/extra
runmmfs: The /lib/modules/2.6.32.12-0.6-ppc64/extra/mmfslinux.ko kernel extension does not exist.
runmmfs: Unable to verify kernel/module configuration.
Loading modules from /lib/modules/2.6.32.12-0.6-ppc64/extra
runmmfs: The /lib/modules/2.6.32.12-0.6-ppc64/extra/mmfslinux.ko kernel extension does not exist.
runmmfs: Unable to verify kernel/module configuration.
Mon Mar 26 20:56:30 EDT 2012 runmmfs: error in loading or unloading the mmfs kernel extension
Mon Mar 26 20:56:30 EDT 2012 runmmfs: stopping GPFS
```

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.
3. 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.

For more information about the **mmfslinux** module, see the *Building the GPFS portability layer* topic in the *IBM Spectrum Scale: Concepts, Planning, and Installation Guide*.

---

## 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 GPFS log does not contain the message:

**6027-300 [N]**

mmfsd ready

- The GPFS 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 **mmismount** command can be used to ensure that all file systems are unmounted.

## Steps to follow if the GPFS daemon does not come up

This topic describes about the steps that you need to follow if the GPFS daemon does not come up after installation of IBM Spectrum Scale.

1. See “GPFS modules cannot be loaded on Linux” on page 79 if your node is running Linux, to verify that you have built the portability layer.
2. Verify that the GPFS daemon is active by issuing:

```
ps -e | grep mmfsd
```

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 by issuing:

```
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 GPFS log. If you see the message:

**6027-1592**

Unable to retrieve GPFS cluster files from node *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.

- Issue the **mmiscluster** command to list the cluster configuration. This will also update the GPFS configuration data on the node. Correct any reported errors before continuing.
- 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
```

If any of these commands produce unexpected results, this may be an indication of corrupted GPFS cluster configuration data file information. Follow the procedures in “Information to be collected before contacting the IBM Support Center” on page 167, and then contact the IBM Support Center.

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 quorum node indicates that it is part of a quorum by writing an **mmfsd ready** message to the GPFS 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 the **mmshutdown** command.
  - b. Remove the shared segment in an AIX environment:
    - 1) Issue the **mmshutdown** command.
    - 2) Issue the **mmfsadm cleanup** command.
  - c. 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, follow the procedures in “Additional information to collect for GPFS daemon crashes” on page 168, and then contact the IBM Support Center.

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

This topic describes about the steps that you need to perform if you are unable to start GPFS after installing a new version of IBM Spectrum Scale.

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

- If message:

**6027-2700 [E]**

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 GPFS 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 GPFS 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 the *Building the GPFS portability layer* topic in the *IBM Spectrum Scale: Concepts, Planning, and Installation Guide*.
  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 73.

## GPFS error messages for shared segment and network problems

This topic describes about the error messages relating to issues in shared segment and network in IBM Spectrum Scale.

For shared segment problems, follow the problem determination and repair actions specified with the following 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.

For network problems, follow the problem determination and repair actions specified with the following message:

**6027-306 [E]**

Could not initialize inter-node communication

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

This topic describes about the application call error numbers 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 IBM Spectrum Scale FAQ in IBM Knowledge Center ([www.ibm.com/support/knowledgecenter/STXKQY/gpfsclustersfaq.html](http://www.ibm.com/support/knowledgecenter/STXKQY/gpfsclustersfaq.html)).

**6027-341 [D]**

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

**6027-342 [E]**

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

**6027-343 [E]**

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

## 6027-344 [E]

Node *nodeName* is incompatible because its version 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 GPFS log contains the message:

### 6027-650 [X]

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 79.

- 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**
- mmpuatacl**
- 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 versions of the commands are:

**tsdf**  
**tsldisk**  
**tslfs**  
**tslsmgr**  
**tslspolicy**  
**tslquota**  
**tslssnapshot**  
**tssnapdir**  
**tssnaplatest**

These are found in the `/usr/lpp/mmfs/bin` directory.

**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.

- For all other errors, follow the procedures in “Additional information to collect for GPFS daemon crashes” on page 168, and then contact the IBM Support Center.

---

## IBM Spectrum Scale failures due to a network failure

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

This dependency is direct because various IBM Spectrum Scale internal messages flow on the network, and may be indirect if the underlying disk technology is dependent on the network. Symptoms included in an indirect failure would be inability to complete I/O or GPFS moving disks to the **down** state.

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 GPFS 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 a sufficient number of nodes fail, GPFS will lose the quorum of nodes, which exhibits itself by messages appearing in the GPFS 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 21.

---

## 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 the *IBM Spectrum Scale: 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** errors. 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:

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

When using subnets ( see the *Using remote access with public and private IP addresses* topic in the *IBM Spectrum Scale: Advanced Administration Guide* ), different interfaces may be in use for intra-cluster and intercluster communication. The presence of a hang or dropped packets 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.

**Note:** A deadlock can occur if user exit scripts that will be called by the **mmaddcallback** facility are placed in a GPFS file system. The scripts should be placed in a local file system so they are accessible even when the networks fail.

To debug a deadlock, do the following:

1. 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 (Maximum disk size allowed is 1.1 TB)
dm2           140095488      1 yes    yes    136434304 ( 97%)    278232 ( 0%)
dm4           140095488      1 yes    yes    136318016 ( 97%)    287442 ( 0%)
dm5           140095488     4000 yes    yes    133382400 ( 95%)    386018 ( 0%)
dm0nsd       140095488     4005 yes    yes    134701696 ( 96%)    456188 ( 0%)
dm1nsd       140095488     4006 yes    yes    133650560 ( 95%)    492698 ( 0%)
dm15         140095488     4006 yes    yes    140093376 (100%)     62 ( 0%)
-----
(pool total)      840572928                                814580352 ( 97%)    1900640 ( 0%)
=====
(total)           840572928                                814580352 ( 97%)    1900640 ( 0%)

```

### Inode Information

```

-----
Number of used inodes:      4244
Number of free inodes:     157036
Number of allocated inodes: 161280
Maximum number of inodes:  512000

```

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. If the file system is approaching these limits, you may notice the following error messages:

#### 6027-533 [W]

Inode space *inodeSpace* in file system *fileSystem* is approaching the limit for the maximum number of inodes.

### 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 automated deadlock detection and deadlock data collection are enabled, look in the latest GPFS log file to determine if the system detected the deadlock and collected the appropriate debug data. Look in `/var/adm/ras/mmfs.log.latest` for messages similar to the following:

```
Thu Feb 13 14:58:09.524 2014: [A] Deadlock detected: 2014-02-13 14:52:59: waiting 309.888 seconds on node  
p7fbn12: SyncHandlerThread 65327: on LkObjCondvar, reason 'waiting for R0 lock'  
Thu Feb 13 14:58:09.525 2014: [I] Forwarding debug data collection request to cluster manager p7fbn11 of  
cluster cluster1.gpfs.net  
Thu Feb 13 14:58:09.524 2014: [I] Calling User Exit Script gpfsDebugDataCollection: event deadlockDebugData,  
Async command /usr/lpp/mmfs/bin/mmcommon.  
Thu Feb 13 14:58:10.625 2014: [N] sdrServ: Received deadlock notification from 192.168.117.21  
Thu Feb 13 14:58:10.626 2014: [N] GPFS will attempt to collect debug data on this node.  
mmtrace: move /tmp/mmfs/lxtrace.trc.p7fbn12.recycle.cpu0  
/tmp/mmfs/trcfile.140213.14.58.10.deadlock.p7fbn12.recycle.cpu0  
mmtrace: formatting /tmp/mmfs/trcfile.140213.14.58.10.deadlock.p7fbn12.recycle to  
/tmp/mmfs/trcprt.140213.14.58.10.deadlock.p7fbn12.gz
```

This example shows that deadlock debug data was automatically collected in `/tmp/mmfs`. If deadlock debug data was not automatically collected, it would need to be manually collected.

To determine which nodes have the longest waiting threads, issue this command on each node:

```
/usr/lpp/mmfs/bin/mmdiag --waiters waitTimeInSeconds
```

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

```
mmfsadm dump all
```

#### Notes:

- a. Each node can potentially dump more than 200 MB of data.
  - b. Run the `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.
3. If the deadlock situation cannot be corrected, follow the instructions in “Additional information to collect for delays and deadlocks” on page 168, then contact the IBM Support Center.

---

## 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 nodeName 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 `mmiscluster` 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 issues

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 *Establishing disaster recovery for your GPFS cluster* topic in the *IBM Spectrum Scale: Advanced Administration Guide* for detailed information about GPFS disaster recovery
- “Creating a master GPFS log file” on page 2
- “Information to be collected before contacting the IBM Support Center” on page 167

The following two messages might appear in the GPFS log for active/active disaster recovery scenarios with GPFS replication. The purpose of these messages is to record quorum override decisions that are made after the loss of most of the disks:

6027-435 [N]

The file system descriptor quorum has been overridden.

6027-490 [N]

The descriptor replica on disk *diskName* has been excluded.

A message similar to these appear in the log on the file system manager, node every time it reads the file system descriptor with an overridden quorum:

```
...
6027-435 [N] The file system descriptor quorum has been overridden.
6027-490 [N] The descriptor replica on disk gpfs23nsd has been excluded.
6027-490 [N] The descriptor replica on disk gpfs24nsd has been excluded.
...
```

For more information on quorum override, see the *IBM Spectrum Scale: Concepts, Planning, and Installation Guide* and search for *quorum*.

For PPRC and FlashCopy-based configurations, more problem determination information can be collected from the ESS log file. This information and the appropriate ESS documentation must be referred while working 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 more information, see the following topics:

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

## Disaster recovery setup problems

The following setup problems might impact disaster recovery implementation:

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 124.
2. In certain cases, it might 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 tolerates and recovers 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 might still be able to perform a restricted mount and attempt to recover parts of their data from the damaged file system. For more information on restricted mounts, see “Restricted mode mount” on page 49.

3. When you issue **mmfsctl syncFSconfig**, you might get an error similar to the following:

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

In such scenarios, 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 by using ssh or rsh without the use of a password.

## Other problems with disaster recovery

You might encounter the following issues that are related to disaster recovery in IBM Spectrum Scale:

1. Currently, users are advised to always specify the **all** option when you issue the **mmfsctl syncFSconfig** command, rather than the device name of one specific file system. Issuing this command enables GPFS to detect and correctly resolve the configuration discrepancies that might occur as a result of the manual administrative action in the target GPFS cluster to which the configuration is imported.
2. The optional **SpecFile** parameter to the **mmfsctl syncFSconfig** that is specified with the **-S** flag must be a fully qualified path name that defines 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 issued. A copy of this file must be available at the provided path name on all peer contact nodes that are 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.latest** and **/var/adm/ras/mmfs.log.previous** on the local node and on the file system manager node. These files enumerate 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 76.
3. The **ssh** command is not functioning correctly. See “Authorization problems” on page 75.  
If **ssh** 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
sshd: 0826-813 Permission is denied.
mmdsh: 6027-1615 k145n02 remote shell process had return code 1.
mmlscluster: 6027-1591 Attention: Unable to retrieve GPFS cluster files from node k145n02
sshd: 0826-813 Permission is denied.
mmdsh: 6027-1615 k145n01 remote shell process had return code 1.
mmlscluster: 6027-1592 Unable to retrieve GPFS cluster files from node k145n01
```

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

If you encounter this type of failure, determine why **ssh** 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 127.
  - 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.

In addition to **mmchfs -V**, you may need to run the **mmmigratefs** command. See the *File system format changes between versions of GPFS* topic in the *IBM Spectrum Scale: Administration and Programming Reference*.

**Note:** Before issuing the **-V** option (with **full** or **compat**), see the *Migration, coexistence and compatibility* topic in the *IBM Spectrum Scale: 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 *IBM Spectrum Scale: Administration and Programming Reference*.

## GPFS error messages for unsuccessful GPFS commands

This topic describes about the 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 83.

#### 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 79.

#### 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:

#### 6027-1627

The following nodes are not aware of the configuration server change: *nodeList*. Do not start GPFS on the preceding 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 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 all -L** 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.

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. Go to the directory just obtained, and use this **mmlsattr -L** command to obtain the fileset name:

```
mmlsattr -L . | grep "fileset name:"
```

The system produces output similar to:

```
fileset name: myFileset
```

- c. Use the **mmlsquota -j** command to check the quota limit of the fileset. For example, using the fileset name found in the previous step, issue this command:

```
mmlsquota -j myFileset -e
```

The system produces output similar to:

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

The **mmquota** output is similar when checking the user and group quota. 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 57.

**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 **mmquota** 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

This topic describes about the error messages that IBM Spectrum Scale displays 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 [E]

File system is read-only.

---

## Troubleshooting Windows problems

The topics that follow apply to Windows Server 2008.

### Home and .ssh directory ownership and permissions

This topic describes about the issues related to .ssh directory ownership and permissions.

Make sure users own their home directories, 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 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

You might have 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.



## GPFS Windows and SMB2 protocol (CIFS serving)

SMB2 is a version of the Server Message Block (SMB) protocol that was introduced with Windows Vista and Windows Server 2008.

Various enhancements include the following (among others):

- reduced “chattiness” of the protocol
- larger buffer sizes
- faster file transfers
- caching of metadata such as directory content and file properties
- better scalability by increasing the support for number of users, shares, and open files per server

The SMB2 protocol is negotiated between a client and the server during the establishment of the SMB connection, and it becomes active only if both the client and the server are SMB2 capable. If either side is not SMB2 capable, the default SMB (version 1) protocol gets used.

The SMB2 protocol does active metadata caching on the client redirector side, and it relies on Directory Change Notification on the server to invalidate and refresh the client cache. However, GPFS on Windows currently does not support Directory Change Notification. As a result, if SMB2 is used for serving out a IBM Spectrum Scale file system, the SMB2 redirector cache on the client will not see any cache-invalidate operations if the actual metadata is changed, either directly on the server or via another CIFS client. In such a case, the SMB2 client will continue to see its cached version of the directory contents until the redirector cache expires. Therefore, the use of SMB2 protocol for CIFS sharing of GPFS file systems can result in the CIFS clients seeing an inconsistent view of the actual GPFS namespace.

A workaround is to disable the SMB2 protocol on the CIFS server (that is, the GPFS compute node). This will ensure that the SMB2 never gets negotiated for file transfer even if any CIFS client is SMB2 capable.

To disable SMB2 on the GPFS compute node, follow the instructions under the “MORE INFORMATION” section at the Microsoft Support website ([support.microsoft.com/kb/974103](http://support.microsoft.com/kb/974103)).

---

## OpenSSH connection delays

OpenSSH can be sensitive to network configuration issues that often do not affect other system components. One common symptom is a substantial delay (20 seconds or more) to establish a connection. When the environment is configured correctly, a command such as `ssh gandalf date` should only take one or two seconds to complete.

If you are using OpenSSH and experiencing an SSH connection delay (and if IPv6 is not supported in your environment), try disabling IPv6 on your Windows nodes and remove or comment out any IPv6 addresses from the `/etc/resolv.conf` file.

---

## File protocol authentication setup issues

When trying to enable Active Directory Authentication for file (smb, nfs), the creation might fail due to a timeout. In some cases, the AD server can return multiple IPs that cannot be queried within the allotted timeout period and/or IPs that belong to networks inaccessible by the IBM Spectrum Scale nodes.

**You can try the following workarounds to resolve this issue:**

- Remove any invalid/unreachable IPs from the AD DNS.  
If you removed any invalid/unreachable IPs, retry the `mmuserauth service create` command that previously failed.
- You can also try to disable any adapters that might not be in use.

For example, on Windows 2008: **Start -> Control Panel -> Network and Sharing Center -> Change adapter settings -> Right-click the adapter that you are trying to disable and click Disable**  
If you disabled any adapters, retry the **mmuserauth service create** command that previously failed.

---

## Chapter 8. File system issues

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 errors encountered with GPFS file systems:

- “File system will not mount”
- “File system will not unmount” on page 104
- “File system forced unmount” on page 105
- “Unable to determine whether a file system is mounted” on page 108
- “Multiple file system manager failures” on page 108
- “Discrepancy between GPFS configuration data and the on-disk data for a file system” on page 109
- “Errors associated with storage pools, filesets and policies” on page 109
- “Failures using the `mmbackup` command” on page 116
- “Snapshot problems” on page 116
- “Failures using the `mmpmon` command” on page 119
- “NFS issues” on page 121
- “Problems working with Samba” on page 123
- “Data integrity” on page 124
- “Messages requeuing in AFM” on page 124

---

### 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 GPFS log.
- Your application cannot access the data it needs. Check the GPFS log for messages.
- Return codes or error messages from the `mmmount` command.
- The `mmismount` 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 has been achieved. Check the GPFS log to see if an `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 127.

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, you can also look locally for error log reports, and follow the problem determination and repair actions specified in your storage system vendor problem determination guide. If the disk has failed, follow the procedures in “NSD and underlying disk subsystem failures” on page 127.

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 83.
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 **mmmout** 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), follow the procedures in “Information to be collected before contacting the IBM Support Center” on page 167, and then 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 *IBM Spectrum Scale: 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 accidentally 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 the *IBM Spectrum Scale: 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 **mmfsfs -A** command to check whether the automatic mount option has been specified. If automatic mount option is expected, check the GPFS 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 21.
13. Verify the **maxblocksize** configuration parameter on all clusters involved. If **maxblocksize** is less than the block size of the local or remote file system you are attempting to mount, you will not be able to mount it.
14. If the file system has encryption rules, see “Mount failure for a file system with encryption rules” on page 143.
15. To mount a file system on a remote cluster, ensure that the cluster that owns and serves the file system and the remote cluster have proper authorization in place. The authorization between clusters is set up with the **mmauth** command.

Authorization errors on AIX are similar to the following:

```
c13c1apv6.gpfs.net: Failed to open remotefs.
c13c1apv6.gpfs.net: Permission denied
c13c1apv6.gpfs.net: Cannot mount /dev/remotefs on /gpfs/remotefs: Permission denied
```

Authorization errors on Linux are similar to the following:

```
mount: /dev/remotefs is write-protected, mounting read-only
mount: cannot mount /dev/remotefs read-only
mmmount: 6027-1639 Command failed. Examine previous error messages to determine cause.
```

For more information about mounting a file system that is owned and served by another GPFS cluster, see the *IBM Spectrum Scale: Advanced Administration Guide*.

## GPFS error messages for file system mount problems

6027-419

Failed to read a file system descriptor.

6027-482 [E]

Remount failed for device *name: errnoDescription*

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 -A** 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/automountdir** 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 Red Hat Enterprise Linux 5, 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:

```
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 Red Hat Enterprise Linux 5, 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.

**Note:** If **automountdir** is mounted (as in step 2) and the **mmcommon startAutomounter** command is not able to bring up the automount daemon, manually **umount** the **automountdir** before issuing the **mmcommon startAutomounter** again.

4. Verify that the mount command was issued to GPFS by examining the GPFS 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) 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 95.

## 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 GPFS 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 95.
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

When a remote file system does not mount, the problem might be with how the file system was defined to both the local and remote nodes, or the communication paths between them. Review the *Mounting a file system owned and served by another GPFS cluster* topic in the *IBM Spectrum Scale: Advanced Administration Guide* to ensure that your setup is correct.

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

- “Remote file system I/O fails with the “Function not implemented” error message when UID mapping is enabled”
- “Remote file system will not mount due to differing GPFS cluster security configurations” on page 101
- “Cannot resolve contact node address” on page 101
- “The remote cluster name does not match the cluster name supplied by the mmremotecluster command” on page 101
- “Contact nodes down or GPFS down on contact nodes” on page 102
- “GPFS is not running on the local node” on page 102
- “The NSD disk does not have an NSD server specified and the mounting cluster does not have direct access to the disks” on page 102
- “The cipherList option has not been set properly” on page 103
- “Remote mounts fail with the “permission denied” error message” on page 103

## 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 by issuing the **mmlsconfig enableUIDremap** command.
3. That UID mapping helper scripts are working correctly.

For more information about configuring UID mapping, see the IBM white paper entitled *UID Mapping for GPFS in a Multi-cluster Environment* in IBM Knowledge Center ([www.ibm.com/support/knowledgecenter/SSFKCN/com.ibm.cluster.gpfs.doc/gpfs\\_uid/uid\\_gpfs.html](http://www.ibm.com/support/knowledgecenter/SSFKCN/com.ibm.cluster.gpfs.doc/gpfs_uid/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 GPFS 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 gpfs1x2 requires
secure connections. Contact the administrator to obtain the target clusters
key and register the key using "mmremoteccluster 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 GPFS 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 **mmremoteccluster** command.

The SHA digest field of the **mmauth show** and **mmremoteccluster** 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 61.

## 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 mmremoteccluster command

A mount command fails with a message similar to this:

```
Cannot mount gpfs1x2:gpfs66: Network is unreachable
```

and the GPFS log contains message similar to this:

```
Mon Jun 25 12:47:18 2007: Waiting to join remote cluster gpfs1x2
Mon Jun 25 12:47:18 2007: Command: mount gpfs1x2:gpfs66 27226
Mon Jun 25 12:47:18 2007: Failed to join remote cluster gpfs1x2
Mon Jun 25 12:47:18 2007: Command err 719: mount gpfs1x2: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 **gpfs1x2.ibm.com** and not **gpfs1x2**

```
mmlscluster
```

Output is similar to this:

GPFS cluster information

=====

```
GPFS cluster name:      gpfs1x2.ibm.com
GPFS cluster id:       649437685184692490
GPFS UID domain:      gpfs1x2.ibm.com
Remote shell command:  /usr/bin/ssh
Remote file copy command: /usr/bin/scp
Repository type:      server-based
```

GPFS cluster configuration servers:

-----

```
Primary server:  gpfs1x2.ibm.com
Secondary server: (none)
```

Node	Daemon node name	IP address	Admin node name	Designation
1	gpfs1x2	198.117.68.68	gpfs1x2.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.
```

The GPFS log will have entries similar to this:

```
Mon Jun 25 13:11:14 2007: Command: mount gpfs1x22:gpfs22 19004
Mon Jun 25 13:11:14 2007: Waiting to join remote cluster gpfs1x22
Mon Jun 25 13:11:15 2007: Connecting to 199.13.68.4 gpfs1x22
Mon Jun 25 13:16:36 2007: Failed to join remote cluster gpfs1x22
Mon Jun 25 13:16:36 2007: Command err 78: mount gpfs1x22:gpfs22 19004
```

To resolve the problem, use the **mmremoteclass 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 **mmremoteclass 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 an 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 gpfs1x2.ibm.com:gpfs66: Stale NFS file handle
```

The GPFS log will contain information similar to this:

```
Mon Jun 25 14:10:46 2007: Command: mount gpfs1x2.ibm.com:gpfs66 28147
Mon Jun 25 14:10:47 2007: Waiting to join remote cluster gpfs1x2.ibm.com
Mon Jun 25 14:10:47 2007: Connecting to 199.13.68.4 gpfs1x2
Mon Jun 25 14:10:47 2007: Connected to 199.13.68.4 gpfs1x2
Mon Jun 25 14:10:47 2007: Joined remote cluster gpfs1x2.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 gpfs1x2.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 GPFS 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.

### Remote mounts fail with the “permission denied” error message

There are many reasons why remote mounts can fail with a “permission denied” error message.

Follow these steps to resolve permission denied problems:

1. Check with the remote cluster's administrator to make sure that the proper keys are in place. The `mmauth show` command on both clusters will help with this.
2. Check that the grant access for the remote mounts has been given on the remote cluster with the `mmauth grant` command. Use the `mmauth show` command from the remote cluster to verify this.
3. Check that the file system access permission is the same on both clusters using the `mmauth show` command and the `mmremotefs show` command. If a remote cluster is only allowed to do a read-only mount (see the `mmauth show` command), the remote nodes must specify `-o ro` on their mount requests (see the `mmremotefs show` command). If you try to do remote mounts with read/write (`rw`) access for remote mounts that have read-only (`ro`) access, you will get a “permission denied” error.

See the *IBM Spectrum Scale: Administration and Programming Reference* for detailed information about the `mmauth` command and the `mmremotefs` command.

### 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 GPFS 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 sdcnsd.
```

```
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 sdcnsd.
```

```
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 50. 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.

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 132 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” 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 **mmumount -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:

```
mmumount fileSystem -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 GPFS 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 **mmismount -L** 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 different failure groups 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 an **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 **mmdisk** 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 the IBM Support Center.
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 57.

If the file system must be repaired without quotas:

- a. Disable quota management by issuing the command:  
`mmchfs Device -Q no`
  - b. Issue the **mmmout** 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 57.
  - g. Issue the **mmmout** 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. See the *NSD creation considerations* topic in the *IBM Spectrum Scale: Concepts, Planning, and Installation Guide* and the *Establishing disaster recovery for your GPFS cluster* topic in the *IBM Spectrum Scale: 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 [X]**

File System `fileSystem` unmounted by the system with return code `value` reason code `value`

**6027-474 [X]**

Recovery Log I/O failed, unmounting file system `fileSystem`

## 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 19 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 9, "Disk issues," on page 127 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 file system *fileSystem* is mounted) is issued.

If you encounter this message, perform problem determination, 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 file system *fileSystem* 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 127.

The required action here is to address the underlying condition that 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 NSD disk servers are being used, the most common failure is the loss of access through the communications network. 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 105, as well as these additional messages:

- When a forced unmount occurred on all nodes:

**6027-635 [E]**

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. See "NSD and underlying disk subsystem failures" on page 127 for NSD problem determination and repair procedures.

**6027-636 [E]**

Disk 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 new manager for *fileSystem*.

- 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 [E]

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 **mmfsdisk** 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, **mmaddisk**, **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 preceding 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, follow the procedures in “Information to be collected before contacting the IBM Support Center” on page 167, and then contact the IBM Support Center.

---

## Errors associated with storage pools, filesets and policies

When an error is suspected while working with storage pools, policies and filesets, check the relevant section in the *IBM Spectrum Scale: 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” on page 110
- “Negative values occur in the 'predicted pool utilizations', when some files are 'ill-placed'” on page 111

- “Policies - usage errors” on page 111
- “Errors encountered with policies” on page 112
- “Filesets - usage errors” on page 113
- “Errors encountered with filesets” on page 114
- “Storage pools - usage errors” on page 114
- “Errors encountered with storage pools” on page 115

## 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	0s ( 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.

- 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:

- Starts with the 'Current pool utilization' for the **system** storage pool, which is 3 GB.
- Subtracts 5 GB, the size of file **abc**.
- 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:

- 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 51.

2. There is a 1 MB limit on the total size of the policy file installed in GPFS.
3. 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.

4. 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.

5. When a policy file is installed, GPFS verifies that the named 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 storage pools referenced in the policy.

## 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 51.

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 51.

- 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.

## Filesets - usage errors

These are common mistakes and misunderstandings encountered when dealing with filesets:

1. Fileset junctions look very much like ordinary directories, but they cannot be deleted by the usual commands such as **rm -r** or **rmdir**. Using these commands on a fileset junction could result in a Not owner message on an AIX system, or an Operation not permitted message on a Linux system.

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.

2. 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.

3. 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 --deleted** command indicates deleted filesets and shows their names in parentheses.

4. 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 is to finish the deletion by reissuing the fileset delete command:

```
mmdeletfileset fs1 fsname1 -f
```

The **mmlsfileset** command identifies filesets in this state by displaying a status of 'Deleting'.

5. 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.
6. By default, the **mmdeletfileset** 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.

7. When deleting a small dependent fileset, it may be faster to use the **rm -rf** command instead of the **mmdelfileset** command with the **-f** option.

## Errors encountered with filesets

These are errors 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 *Filesets and backup* topic in the *IBM Spectrum Scale: 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.

## Active file management fileset errors

When the **mmafmctl Device getstate** command displays a NeedsResync target/fileset state, inconsistencies exist between the home and cache. To ensure that the cached data is synchronized with the home and the fileset is returned to Active state, either the file system must be unmounted and mounted or the fileset must be unlinked and linked. Once this is done, the next update to fileset data will trigger an automatic synchronization of data from the cache to the home.

## 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 the desired level of 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. 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.

The **mmddf** command can be used to find the maximum disk size allowed for a storage pool.

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, follow the procedures in “Information to be collected before contacting the IBM Support Center” on page 167, and then 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 reassigned 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** or **mmchattr** command did not complete its last migration or that **-I defer** was used.

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 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 9, “Disk issues,” on page 127.

---

## 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 `mmbackup` command was issued to back up a file system owned by a remote cluster.
- 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 errors from `mmbackup` normally indicate the underlying problem.

## GPFS error messages for `mmbackup` errors

6027-1995

Device *deviceName* is not mounted on node *nodeName*.

## TSM error messages

ANS1312E

Server media mount not possible.

---

## Snapshot problems

Use the `mmlssnapshot` command as a general hint for snapshot-related problems, to find out what snapshots exist, and what state they are in. Use the `mmsnapdir` command to find 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

The `mmlssnapshot` and `mmsnapdir` 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 or node failures could cause a snapshot command to fail. The response to these types of errors is to 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 usage restrictions or incorrect snapshot names.

An example of a snapshot restriction error is exceeding the maximum number of snapshots allowed at one time. For simple errors of these types, you can determine the source of the error by reading the error message or by reading the description of the command. You can also run the **mmlssnapshot** command to see the complete list of existing snapshots.

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 an existing snapshot. The rules for naming global and fileset snapshots are designed to minimize conflicts between the file system administrator and the fileset owners. These rules can result in errors when fileset snapshot names are duplicated across different filesets or when the snapshot command **-j** option (specifying a qualifying fileset name) is provided or omitted incorrectly. To resolve name problems review the **mmlssnapshot** output with careful attention to the Fileset column. You can also specify the **-s** or **-j** options of the **mmlssnapshot** command to limit the output. For snapshot deletion, the **-j** option must exactly match the Fileset column.

For more information about snapshot naming conventions, see the **mmcrsnapshot** command in the *IBM Spectrum Scale: Administration and Programming Reference*.

## 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

The following errors might be encountered when restoring from a snapshot:

- The **mmrestorefs** command fails with an **ENOSPC** message. In this case, there are not enough free blocks in the file system to restore the selected snapshot. You can add space to the file system by adding a new disk. As an alternative, you can delete a different snapshot from the file system to free some existing space. You cannot delete the snapshot that is being restored. After there is additional free space, issue the **mmrestorefs** command again.
- The **mmrestorefs** command fails with quota exceeded errors. Try adjusting the quota configuration or disabling quota, and then issue the command again.
- The **mmrestorefs** command is interrupted and some user data is not be restored completely. Try repeating the **mmrestorefs** command in this instance.
- The **mmrestorefs** command fails because of an incorrect file system, fileset, or snapshot name. To fix this error, issue the command again with the correct name.
- The **mmrestorefs -j** command fails with the following error:

6027-953

Failed to get a handle for fileset *filesetName*, snapshot *snapshotName* in file system *fileSystem*.  
*errorMessage*.

In this case, the file system that contains the snapshot to restore should be mounted, and then the fileset of the snapshot should be linked.

If you encounter additional errors that cannot be resolved, contact the IBM Support Center.

## 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 be successful but the snapshot may not be accessed.

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, **mmlsnapshot** 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.

The inode numbers that are used for and within these special **.snapshots** directories are constructed dynamically and do not follow the standard rules. These inode numbers are visible to applications through standard commands, such as **stat**, **readdir**, or **ls**. The inode numbers reported for these directories can also be reported differently on different operating systems. Applications should not expect consistent numbering for such inodes.

---

## Failures using the mmpmon command

The **mmpmon** command manages performance monitoring and displays performance information.

The **mmpmon** command is thoroughly documented in the *Monitoring GPFS I/O performance with the mmpmon command* topic in the *IBM Spectrum Scale: Advanced Administration Guide*, and the *GPFS Commands* chapter in the *IBM Spectrum Scale: Administration and Programming Reference*. Before proceeding with **mmpmon** problem determination, review all of this material 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 *IBM Spectrum Scale: Advanced Administration Guide*. For the limitations regarding concurrent usage of **mmpmon**, see *Running mmpmon concurrently from multiple users* in *IBM Spectrum Scale: Advanced Administration Guide*.
- 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 *IBM Spectrum Scale: Advanced Administration Guide* and search for *Counter sizes and counter wrapping*.
3. See *Monitoring GPFS I/O performance with the mmpmon command* in the *IBM Spectrum Scale: Advanced Administration Guide* and search for *Other information 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. Follow the procedures in “Information to be collected before contacting the IBM Support Center” on page 167, and then 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, follow the procedures in “Information to be collected before contacting the IBM Support Center” on page 167, and then contact the IBM Support Center.

## Tracing the mmpmon command

When the **mmpmon** command 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 the **mmpmon** command, 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 issues

This topic describes some of the possible problems that can be encountered when GPFS interacts with NFS.

For details on how GPFS and NFS interact, see the *NFS and GPFS* topic in the *IBM Spectrum Scale: Administration and Programming Reference*.

These are some of the problems encountered when GPFS interacts with NFS:

- “NFS client with stale inode data”
- “NFS V4 problems”

### 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 *NFS Version 4 Protocol* paper and other information found in the Network File System Version 4 (nfsv4) section of the IETF Datatracker website ([datatracker.ietf.org/wg/nfsv4/documents](http://datatracker.ietf.org/wg/nfsv4/documents)).
2. The *Managing GPFS access control lists and NFS export* topic in the *IBM Spectrum Scale: Administration and Programming Reference*.
3. The *GPFS exceptions and limitations to NFS V4 ACLs* topic in the *IBM Spectrum Scale: 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.

---

## Determining the health of integrated SMB server

The following commands can be used to determine the health of SMB services:

- To check the overall CES cluster state, issue the following command:

```
mmiscluster --ces
```

The system displays output similar to this:

```
GPFS cluster information
```

```
=====
```

```
GPFS cluster name:      boris.nsd001st001
GPFS cluster id:       3992680047366063927
```

```
Cluster Export Services global parameters
```

```
-----
```

```
Shared root directory:  /gpfs/fs0
Enabled Services:      NFS SMB
Log level:             2
Address distribution policy: even-coverage
```

Node	Daemon node name	IP address	CES IP address list
4	prt001st001	172.31.132.1	10.18.24.25 10.18.24.32 10.18.24.34 10.18.24.36 9.11.102.89
5	prt002st001	172.31.132.2	9.11.102.90 10.18.24.19 10.18.24.21 10.18.24.23 10.18.24.30
6	prt003st001	172.31.132.3	10.18.24.38 10.18.24.39 10.18.24.41 10.18.24.42 9.11.102.43
7	prt004st001	172.31.132.4	9.11.102.37 10.18.24.26 10.18.24.28 10.18.24.18 10.18.24.44
8	prt005st001	172.31.132.5	9.11.102.36 10.18.24.17 10.18.24.33 10.18.24.35 10.18.24.37
9	prt006st001	172.31.132.6	9.11.102.41 10.18.24.24 10.18.24.20 10.18.24.22 10.18.24.40
10	prt007st001	172.31.132.7	9.11.102.42 10.18.24.31 10.18.24.27 10.18.24.29 10.18.24.43

This shows at a glance whether nodes are failed or whether they host public IP addresses. For successful SMB operation at least one CES node must be HEALTHY and hosting at least one IP address.

- To show which services are enabled, issue the following command:

```
mmces service list
```

The system displays output similar to this:

```
Enabled services: NFS SMB
NFS is running, SMB is running
```

For successful SMB operation, SMB needs to be enabled and running.

- To determine the overall health state of SMB on all CES nodes, issue the following command:

```
mmces state show smb -a
```

The system displays output similar to this:

```
NODE      SMB
prt001st001 HEALTHY
prt002st001 HEALTHY
prt003st001 HEALTHY
prt004st001 HEALTHY
prt005st001 HEALTHY
prt006st001 HEALTHY
prt007st001 HEALTHY
```

- To show the reason for a currently active (failed) state on all nodes, issue the following command:

```
mmces events active SMB -a
```

The system displays output similar to this:

```
NODE COMPONENT  EVENT NAME SEVERITY  DETAILS
```

In this case nothing is listed because all nodes are healthy and so there are no active events. If a node was unhealthy it would look similar to this:

NODE	COMPONENT	EVENT NAME	SEVERITY	DETAILS
prt001st001	SMB	ctdb_down	ERROR	CTDB process not running
prt001st001	SMB	smbd_down	ERROR	SMBD process not running

- To show the history of events generated by the monitoring framework, issue the following command  

```
mmces events list SMB
```

The system displays output similar to this:

NODE	TIMESTAMP	EVENT NAME	SEVERITY	DETAILS
prt001st001	2015-05-27 14:15:48.540577+07:07MST	smbd_up	INFO	SMBD process now running
prt001st001	2015-05-27 14:16:03.572012+07:07MST	smbport_up	INFO	SMB port 445 is now active
prt001st001	2015-05-27 14:28:19.306654+07:07MST	ctdb_recovery	WARNING	CTDB Recovery detected
prt001st001	2015-05-27 14:28:34.329090+07:07MST	ctdb_recovered	INFO	CTDB Recovery finished
prt001st001	2015-05-27 14:33:06.002599+07:07MST	ctdb_recovery	WARNING	CTDB Recovery detected
prt001st001	2015-05-27 14:33:19.619583+07:07MST	ctdb_recovered	INFO	CTDB Recovery finished
prt001st001	2015-05-27 14:43:50.331985+07:07MST	ctdb_recovery	WARNING	CTDB Recovery detected
prt001st001	2015-05-27 14:44:20.285768+07:07MST	ctdb_recovered	INFO	CTDB Recovery finished
prt001st001	2015-05-27 15:06:07.302641+07:07MST	ctdb_recovery	WARNING	CTDB Recovery detected
prt001st001	2015-05-27 15:06:21.609064+07:07MST	ctdb_recovered	INFO	CTDB Recovery finished
prt001st001	2015-05-27 22:19:31.773404+07:07MST	ctdb_recovery	WARNING	CTDB Recovery detected
prt001st001	2015-05-27 22:19:46.839876+07:07MST	ctdb_recovered	INFO	CTDB Recovery finished
prt001st001	2015-05-27 22:22:47.346001+07:07MST	ctdb_recovery	WARNING	CTDB Recovery detected
prt001st001	2015-05-27 22:23:02.050512+07:07MST	ctdb_recovered	INFO	CTDB Recovery finished

- To retrieve monitoring state from health monitoring component, issue the following command:  

```
mmces state show
```

The system displays output similar to this:

NODE	AUTH	NETWORK	NFS	OBJECT	SMB	CES
prt001st001	DISABLED	HEALTHY	HEALTHY	DISABLED	DISABLED	HEALTHY

- To check the monitor log, issue the following command:

```
grep smb /var/adm/ras/mmcesmonitor.log | head -n 10
```

The system displays output similar to this:

```
2015-05-29T06:42:34.559-07:00 prt003st001 D:15573:MonitorEventScheduler_smb:smb:Trigger monitoring event for
MonitorEventScheduler_smb (interval 15)
2015-05-29T06:42:34.559-07:00 prt003st001 I:15573:Thread-5:smb:Monitor SMB service ...
2015-05-29T06:42:34.560-07:00 prt003st001 D:15573:Thread-5:smb:ProcessMonitor smbd started:
2015-05-29T06:42:34.588-07:00 prt003st001 D:15573:Thread-5:smb:ProcessMonitor smbd succeeded
2015-05-29T06:42:34.589-07:00 prt003st001 D:15573:Thread-5:smb:PortMonitor SMB started:
2015-05-29T06:42:34.594-07:00 prt003st001 D:15573:Thread-5:smb:ProcessMonitor ctdbd started:
2015-05-29T06:42:34.617-07:00 prt003st001 D:15573:Thread-5:smb:ProcessMonitor ctdbd succeeded
2015-05-29T06:42:34.618-07:00 prt003st001 D:15573:Thread-5:smb:CommandMonitor /usr/lpp/mmfs/bin/ctdb status -x +
| /usr/bin/cut -d '+' -f 4,5,6,8,11 | /usr/bin/tee /dev/stderr | /bin/grep 0+0+0+0+Y started:
2015-05-29T06:42:34.633-07:00 prt003st001 D:15573:Thread-5:smb:CommandMonitor /usr/lpp/mmfs/bin/ctdb status -x +
| /usr/bin/cut -d '+' -f 4,5,6,8,11 | /usr/bin/tee /dev/stderr | /bin/grep 0+0+0+0+Y succeeded.
Return code check for 0
2015-05-29T06:42:34.633-07:00 prt003st001 D:15573:Thread-5:smb:CommandMonitor /usr/lpp/mmfs/bin/ctdb status
| /bin/grep 'Recovery mode.*NORMAL' started:
```

- The following logs can also be checked:

```
/var/adm/ras/*
/var/log/messages
```

---

## 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 an **MMFS\_SYSTEM\_UNMOUNT** and a corruption error are generated.

If actual disk data corruption occurs, this error will appear on each node in succession. Before proceeding with the following steps, follow the procedures in “Information to be collected before contacting the IBM Support Center” on page 167, and then contact the IBM Support Center.

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 59 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” 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 the *IBM Spectrum Scale: 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.

---

## Messages queuing in AFM

Sometimes requests in the AFM messages queue on the gateway node get requeued because of errors at home. For example, if there is no space at home to perform a new write, a write message that is queued is not successful and gets requeued. The administrator would see the failed message getting requeued in the queue on the gateway node. The administrator has to resolve the issue by adding more space at home and running the **mmafmctl resumeRequeued** command, so that the requeued messages are executed at home again. If **mmafmctl resumeRequeued** is not run by an administrator, AFM would still execute the message in the regular order of message executions from cache to home.



Running the **mmfsadm dump afm all** command on the gateway node shows the queued messages. Queued messages show in the dumps similar to the following example:

```
c12c4apv13.gpfs.net: Normal Queue: (listed by execution order) (state: Active)
c12c4apv13.gpfs.net: Write [612457.552962] requeued file3 (43 @ 293) chunks 0 bytes 0 0
```



---

## Chapter 9. Disk issues

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. 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”
- “GPFS has declared NSDs built on top of AIX logical volumes as down” on page 136
- “Disk accessing commands fail to complete due to problems with some non-IBM disks” on page 138
- “Persistent Reserve errors” on page 138
- “GPFS is not using the underlying multipath device” on page 141

---

### 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 105.
- The **mmismount** 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 **mmisdisk** command shows that disks are down.

**Note:** If you are reinstalling the operating system on one node and erasing all partitions from the system, GPFS descriptors will be removed from any NSD this node can access locally. The results of this action might require recreating the file system and restoring from backup. If you experience this problem, do not unmount the file system on any node that is currently mounting the file system. Contact the IBM Support Center immediately to see if the problem can be corrected.

### 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 stanzas. For a complete explanation of disk stanzas, see the following *IBM Spectrum Scale: Administration and Programming Reference* topics:

- *Stanza files*
- **mmchdisk** command
- **mmchnsd** command
- **mmcrfs** command
- **mmcrnsd** command

For disks that are SAN-attached to all nodes in the cluster, **device=DiskName** should refer to the disk device name in **/dev** on the node where the **mmcrnsd** command is issued. If a server list is specified, **device=DiskName** 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, one of these messages is displayed:

**6027-2108**

Error found while processing stanza

or

**6027-1636**

Error found while checking disk descriptor *descriptor*

Usually, this message is preceded by one or more messages describing the error more specifically.

Another possible error from **mmcrnsd** is:

**6027-2109**

Failed while processing disk stanza on node *nodeName*.

or

**6027-1661**

Failed while processing disk descriptor *descriptor* on node *nodeName*.

One of these errors 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 file system **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, **nodec26f4gp04**.
- 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 devices 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 "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 preceding 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

NSDs are deleted with the **mmdeInsd** 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 **mmdeInsd** 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 device *deviceName* 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 **mmdeInsd -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  
mmdeInsd -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 19). 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 may result in a similar marking of disks as stopped.
  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.

## 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 19 for hardware configuration error log entries.

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 three 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 the *File system replication parameters* topic in the *IBM Spectrum Scale: Concepts, Planning, and Installation Guide*.

## 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 command scans file system metadata for disk addresses belonging to the disk in question, then replaces them with a special “broken disk address” value, which may take a while.

**CAUTION:**

**Be extremely careful with using the -p option of mmdeldisk, because by design it destroys references to data blocks, making affected blocks unavailable. This is a last-resort tool, to be used when data loss may have already occurred, to salvage the remaining data—which means it cannot take any precautions. If you are not absolutely certain about the state of the file system and the impact of running this command, do not attempt to run it without first contacting the IBM Support Center.**

```
mmdeldisk fs1 gpfs1n12 -p
```

4. Invoke the **mmfileid** command with the operand **:BROKEN**:

```
mmfileid :BROKEN
```

For more information, see “The mmfileid command” on page 59.

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 block-size 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 data or metadata replication is enabled, and the status of an existing disk changes so that the disk is no longer available for block allocation (if strict replication is enforced), you may receive an **errno** of **ENOSPC** when you create or append data to an existing file. A disk becomes unavailable for new block allocation if it is being deleted, replaced, or it has been suspended. 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 -K** command with the original value of the **-K** option (**always** or **whenpossible**) and then run the **mmrestripefs -r** command. To determine if a disk has strict replication enforced, issue the **mmlsfs -K** command.

**Note:** A disk in a **down** state that has not been explicitly suspended is still available for block allocation, and thus a spontaneous disk failure will not result in application I/O requests failing with **ENOSPC**. While new blocks will be allocated on such a disk, nothing will actually be written to the disk until its availability changes to **up** following an **mmchdisk start** command. Missing replica updates that took place while the disk was down will be performed when **mmchdisk start** runs.

### 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 49). 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 block-size 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 cluster configuration data 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:

**6027-418**

Inconsistent file system quorum. readQuorum=*value* writeQuorum=*value* quorumSize=*value*

**6027-482 [E]**

Remount failed for device *name: errnoDescription*

**6027-485**

Perform **mmchdisk** for any disk failures and re-mount.

**6027-636 [E]**

Disk 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:

### **EIO = 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.

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, check the operating system error log on all nodes for indications of errors. Take corrective action by issuing the **mmchdisk** command.

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 GPFS log:

**6027-361 [E]**

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 the *NSD server considerations* topic in the *IBM Spectrum Scale: Concepts, Planning, and Installation Guide*.

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.

**Note:** In general, after fixing the path to a disk, you must run the `mmnsddiscover` command on the server that lost the path to the NSD. (Until the `mmnsddiscover` command is run, the reconnected node will see its local disks and start using them by itself, but it will not act as the NSD server.)

After that, you must run the command on all client nodes that need to access the NSD on that server; or you can achieve the same effect with a single `mmnsddiscover` invocation if you utilize the `-N` option to specify a node list that contains all the NSD servers and clients that need to rediscover paths.

## 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:  
`mmdeledisk 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:

```
mm1nsd -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 **pvids** 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. For more information, refer to AIX in IBM Knowledge Center ([www.ibm.com/support/knowledgecenter/ssw\\_aix/welcome](http://www.ibm.com/support/knowledgecenter/ssw_aix/welcome)) and search for *operating system and device management*.

## 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.

---

## Disk accessing commands fail to complete due to problems with some non-IBM disks

Certain disk commands, such as **mmcrfs**, **mmadddisk**, **mmrpldisk**, **mmmout** 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 vgrname**) 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.

---

## 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

**Note:** While Persistent Reserve (PR) is supported on both AIX and Linux, *reserve\_policy* is applicable only to AIX.

Persistent Reserve 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 (sddisk) 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. `s`
3. 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/tspreadkeys 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
```

## Clearing a leftover Persistent Reserve reservation

Message number **6027-2202** indicates that a specified disk has a SCSI-3 PR reservation, which prevents the **mmcrnsd** command from formatting it. The following example is specific to a Linux environment. Output on AIX is similar but not identical.

Before trying to clear the PR reservation, use the following instructions to verify that the disk is really intended for GPFS use. Note that in this example, the device name is specified without a prefix (**/dev/sdp** is specified as **sdp**).

1. Display all the registration key values on the disk:

```
/usr/lpp/mmfs/bin/tspreadkeys sdp
```

The system displays information similar to:

```
Registration keys for sdp
1. 00006d0000000001
```

If the registered key values all start with 0x00006d, which indicates that the PR registration was issued by GPFS, proceed to the next step to verify the SCSI-3 PR reservation type. Otherwise, contact your system administrator for information about clearing the disk state.

2. Display the reservation type on the disk:

```
/usr/lpp/mmfs/bin/tspreadres sdp
```

The system displays information similar to:

```
yes:LU_SCOPE:WriteExclusive-AllRegistrants:0000000000000000
```

If the output indicates a PR reservation with type **WriteExclusive-AllRegistrants**, proceed to the following instructions for clearing the SCSI-3 PR reservation on the disk.

If the output does not indicate a PR reservation with this type, contact your system administrator for information about clearing the disk state.

To clear the SCSI-3 PR reservation on the disk, follow these steps:

1. Choose a hex value (*HexValue*); for example, **0x111abc** that is not in the output of the **tspreadkeys** command run previously. Register the local node to the disk by entering the following command with the chosen *HexValue*:

```
/usr/lpp/mmfs/bin/tspregister sdp 0x111abc
```

2. Verify that the specified *HexValue* has been registered to the disk:

```
/usr/lpp/mmfs/bin/tspreadkeys sdp
```

The system displays information similar to:

```
Registration keys for sdp
1. 00006d0000000001
2. 0000000000111abc
```

3. Clear the SCSI-3 PR reservation on the disk:

```
/usr/lpp/mmfs/bin/tspclear sdp 0x111abc
```

4. Verify that the PR registration has been cleared:

```
/usr/lpp/mmfs/bin/tspreadkeys sdp
```

The system displays information similar to:

```
Registration keys for sdp
```

5. Verify that the reservation has been cleared:

```
/usr/lpp/mmfs/bin/tspreadres sdp
```

The system displays information similar to:

```
no:::
```

The disk is now ready to use for creating an NSD.

## Manually enabling or disabling Persistent Reserve

**Attention:** Manually enabling or disabling Persistent Reserve should only be done under the supervision of the IBM Support Center with GPFS stopped on the node.



The IBM Support Center 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.

---

## GPFS is not using the underlying multipath device

You can view the underlying disk device where I/O is performed on an NSD disk by using the **mmlsdisk** command with the **-M** option.

The **mmlsdisk** command output might show unexpected results for multipath I/O devices. For example if you issue this command:

```
mmlsdisk dmfs2 -M
```

The system displays information similar to:

Disk name	I/O performed on node	Device	Availability
m0001	localhost	/dev/sdb	up

The following command is available on Linux only.

```
# multipath -ll
mpathae (36005076304ffc0e50000000000000001) dm-30 IBM,2107900
[size=10G][features=1 queue_if_no_path][hwhandler=0]
\_ round-robin 0 [prio=8][active]
  \_ 1:0:5:1 sdhr 134:16 [active][ready]
  \_ 1:0:4:1 sdgl 132:16 [active][ready]
  \_ 1:0:1:1 sdff 130:16 [active][ready]
  \_ 1:0:0:1 sddz 128:16 [active][ready]
  \_ 0:0:7:1 sdct 70:16 [active][ready]
  \_ 0:0:6:1 sdbn 68:16 [active][ready]
  \_ 0:0:5:1 sdah 66:16 [active][ready]
  \_ 0:0:4:1 sdb 8:16 [active][ready]
```

The **mmlsdisk** output shows that I/O for NSD **m0001** is being performed on disk **/dev/sdb**, but it should show that I/O is being performed on the device-mapper multipath (DMM) **/dev/dm-30**. Disk **/dev/sdb** is one of eight paths of the DMM **/dev/dm-30** as shown from the **multipath** command.

This problem could occur for the following reasons:

- The previously installed user exit **/var/mmfs/etc/nsddevices** is missing. To correct this, restore user exit **/var/mmfs/etc/nsddevices** and restart GPFS.
- The multipath device type does not match the GPFS known device type. For a list of known device types, see **/usr/lpp/mmfs/bin/mmdevdiscover**. After you have determined the device type for your multipath device, use the **mmchconfig** command to change the NSD disk to a known device type and then restart GPFS.

The following output shows that device type **dm-30** is **dmm**:

```
/usr/lpp/mmfs/bin/mmdevdiscover | grep dm-30
dm-30 dmm
```

To change the NSD device type to a known device type, create a file that contains the NSD name and device type pair (one per line) and issue this command:

```
mmchconfig updateNsdType=/tmp/filename
```

where the contents of **/tmp/filename** are:

```
m0001 dmm
```

The system displays information similar to:

```
mmchconfig: Command successfully completed  
mmchconfig: Propagating the cluster configuration data to all  
affected nodes. This is an asynchronous process.
```

---

## Chapter 10. Encryption issues

The topics that follow provide solutions for problems that may be encountered while setting up or using encryption.

---

### Unable to add encryption policies

If the `mmchpolicy` command fails when you are trying to add encryption policies, perform the following diagnostic steps:

1. Confirm that the `gpfs.crypto` and `gpfs.gskit` packages are installed.
2. Confirm that the file system is at GPFS 4.1 or later and the fast external attributes (`--fastea`) option is enabled.
3. Examine the error messages that are logged in the `mmfs.log.latest` file, which is located at: `/var/adm/ras/mmfs.log.latest`.

---

### Receiving “Permission denied” message

If you experience a “Permission denied” failure while creating, opening, reading, or writing to a file, perform the following diagnostic steps:

1. Confirm that the key server is operational and correctly set up and can be accessed through the network.
2. Confirm that the `/var/mmfs/etc/RKM.conf` file is present on all nodes from which the file is supposed to be accessed. The `/var/mmfs/etc/RKM.conf` file must contain entries for all the RKMs needed to access the file.
3. Verify that the master keys needed by the file and the keys that are specified in the encryption policies are present on the key server.
4. Examine the error messages in the `/var/adm/ras/mmfs.log.latest` file.

---

### “Value too large” failure when creating a file

If you experience a “Value too large to be stored in data type” failure when creating a file, follow these diagnostic steps.

1. Examine error messages in `/var/adm/ras/mmfs.log.latest` to confirm that the problem is related to the extended attributes being too large for the inode. The size of the encryption extended attribute is a function of the number of keys used to encrypt a file. If you encounter this issue, update the encryption policy to reduce the number of keys needed to access any given file.
2. If the previous step does not solve the problem, create a new file system with a larger inode size.

---

### Mount failure for a file system with encryption rules

If you experience a mount failure for a file system with encryption rules, follow these diagnostic steps.

1. Confirm that the `gpfs.crypto` and `gpfs.gskit` packages are installed.
2. Confirm that the `/var/mmfs/etc/RKM.conf` file is present on the node and that the content in `/var/mmfs/etc/RKM.conf` is correct.
3. Examine the error messages in `/var/adm/ras/mmfs.log.latest`.

---

### “Permission denied” failure of key rewrap

If you experience a “Permission denied” failure of a key rewrap, follow these diagnostic steps.

When **mmapplypolicy** is invoked to perform a key rewrap, the command may issue messages like the following:

```
[E] Error on gpfs_enc_file_rewrap_key(/fs1m/s1s/test4,KEY-d7bd45d8-9d8d-4b85-a803-e9b794ec0af2:hs21n56_new,KEY-40a0b68b-c86d-4519-9e48-3714d3b71e20:js21n92)
Permission denied(13)
```

If you receive a message similar to this, follow these steps:

1. Check for syntax errors in the migration policy syntax.
2. Ensure that the new key is not already being used for the file.
3. Ensure that both the original and the new keys are retrievable.
4. Examine the error messages in `/var/adm/ras/mmfs.log.latest` for additional details.

---

## Chapter 11. Other problem determination hints and tips

These hints and tips might be helpful when investigating problems related to logical volumes, quorum nodes, or system performance that can be encountered while using GPFS.

See these topics for more information:

- “Which physical disk is associated with a logical volume?”
- “Which nodes in my cluster are quorum nodes?”
- “What is stored in the /tmp/mmfs directory and why does it sometimes disappear?” on page 146
- “Why does my system load increase significantly during the night?” on page 146
- “What do I do if I receive message 6027-648?” on page 147
- “Why can't I see my newly mounted Windows file system?” on page 147
- “Why is the file system mounted on the wrong drive letter?” on page 147
- “Why does the offline mmfsck command fail with "Error creating internal storage"?” on page 147
- “Questions related to active file management” on page 148
- “Questions related to File Placement Optimizer (FPO)” on page 148

---

### 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%         108:107:107:108
```

---

### 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
Repository type:       server-based
```

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.

---

## 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 the IBM Support Center.

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 directories 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
```

**Note:** This state information (possibly large amounts of data in the form of GPFS dumps and traces) can be dumped automatically as part the first failure data capture mechanisms of GPFS, and can accumulate in the (default **/tmp/mmfs**) directory that is defined by the **dataStructureDump** configuration parameter. It is recommended that a cron job (such as **/etc/cron.daily/tmpwatch**) be used to remove **dataStructureDump** directory data that is older than two weeks, and that such data is collected (for example, via **gpfs.snap**) within two weeks of encountering any problem that requires investigation.

---

## 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.**

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.

---

## Why is the file system mounted on the wrong drive letter?

Before mounting a GPFS file system, you must be certain that the drive letter required for GPFS is freely available and is not being used by a local disk or a network-mounted file system on *all* computation nodes where the GPFS file system will be mounted.

---

## Why does the offline `mmfsck` command fail with "Error creating internal storage"?

The `mmfsck` command requires some temporary space on the file system manager for storing internal data during a file system scan. The internal data will be placed in the directory specified by the `mmfsck -t` command line parameter (`/tmp` by default). The amount of temporary space that is needed is proportional to the number of inodes (used and unused) in the file system that is being scanned. If GPFS is unable to create a temporary file of the required size, the `mmfsck` command will fail with the following error message:

```
Error creating internal storage
```

This failure could be caused by:

- The lack of sufficient disk space in the temporary directory on the file system manager
- The lack of sufficient pagepool on the file system manager as shown in `mmlsconfig pagepool` output
- Insufficiently high `filesize` limit set for the `root` user by the operating system
- The lack of support for large files in the file system that is being used for temporary storage. Some file systems limit the maximum file size because of architectural constraints. For example, JFS on AIX does not support files larger than 2 GB, unless the **Large file support** option has been specified when the file system was created. Check local operating system documentation for maximum file size limitations.

---

## Why do I get timeout executing function error message?

If any of the commands fails due to timeout while executing `mmccr`, rerun the command to fix the issue. This timeout issue is likely related to an increased workload of the system.

---

## Questions related to active file management

The following questions are related to active file management (AFM).

### How can I change the mode of a fileset?

The mode of an AFM client cache fileset cannot be changed from local-update mode to any other mode; however, it can be changed from read-only to single-writer (and vice versa), and from either read-only or single-writer to local-update.

To change the mode, do the following:

1. Ensure that fileset status is active and that the gateway is available.
2. Umount the file system.
3. Unlink the fileset.
4. Run the **mmchfileset** command to change the mode.
5. Mount the file system again.
6. Link the fileset again.

### Why are setuid/setgid bits in a single-writer cache reset at home after data is appended?

The setuid/setgid bits in a single-writer cache are reset at home after data is appended to files on which those bits were previously set and synced. This is because over NFS, a write operation to a setuid file resets the setuid bit.

### How can I traverse a directory that has not been cached?

On a fileset whose metadata in all subdirectories is not cached, any application that optimizes by assuming that directories contain two fewer subdirectories than their hard link count will not traverse the last subdirectory. One such example is **find**; on Linux, a workaround for this is to use **find -noleaf** to correctly traverse a directory that has not been cached.

### What extended attribute size is supported?

For an operating system in the gateway whose Linux kernel version is below 2.6.32, the NFS max rsize is 32K, so AFM would not support an extended attribute size of more than 32K on that gateway.

### What should I do when my file system or fileset is getting full?

The **.ptrash** directory is present in cache and home. In some cases, where there is a conflict that AFM cannot resolve automatically, the file is moved to **.ptrash** at cache or home. In cache the **.ptrash** gets cleaned up when eviction is triggered. At home, it is not cleared automatically. When the administrator is looking to clear some space, the **.ptrash** should be cleaned up first.

---

## Questions related to File Placement Optimizer (FPO)

The following questions are related to File Placement Optimizer (FPO).

### Why is my data not read from the network locally when I have an FPO pool (write-affinity enabled storage pool) created?

When you create a storage pool that is to contain files that make use of FPO features, you must specify **allowWriteAffinity=yes** in the storage pool stanza.



To enable the policy to read replicas from local disks, you must also issue the following command:

```
mmchconfig readReplicaPolicy=local
```

### **How can I change a failure group for a disk in an FPO environment?**

To change the failure group in a write-affinity-enabled storage pool, you must use the **mmdeldisk** and **mmaddisk** commands; you cannot use **mmchdisk** to change it directly.

### **Why does Hadoop receive a fixed value for the block group factor instead of the GPFS default value?**

When a customer does not define the **dfs.block.size** property in the configuration file, the GPFS connector will use a fixed block size to initialize Hadoop. The reason for this is that Hadoop has only one block size per file system, whereas GPFS allows different chunk sizes (block-group-factor × data block size) for different data pools because block size is a per-pool property. To avoid a mismatch when using Hadoop with FPO, define **dfs.block.size** and **dfs.replication** in the configuration file.

### **How can I retain the original data placement when I restore data from a TSM server?**

When data in an FPO pool is backed up in a TSM server and then restored, the original placement map will be broken unless you set the write affinity failure group for each file before backup.

### **How is an FPO pool file placed at AFM home and cache?**

For AFM home or cache, an FPO pool file written on the local side will be placed according to the write affinity depth and write affinity failure group definitions of the local side. When a file is synced from home to cache, it follows the same FPO placement rule as when written from the gateway node in the cache cluster. When a file is synced from cache to home, it follows the same FPO data placement rule as when written from the NFS server in the home cluster.

To retain the same file placement at both home and cache, ensure that each has the same cluster configuration, and set the write affinity failure group for each file.



## Chapter 12. Reliability, Availability, and Serviceability (RAS) events

The following tables list the RAS events that are applicable to various components of the IBM Spectrum Scale system.

**Note:** The recorded events are stored in local database on each node. The user can get a list of recorded events using `mmces events list` command. The recorded events can also be displayed through GUI.

*Table 7. Events for the AUTH component*

Event	EventType	Severity	Message	Description	Cause	User Action
ads_down	STATE_CHANGE	ERROR	External ADS server is unresponsive	External ADS server is unresponsive	The local node is unable to connect to any ADS server.	Local node is unable to connect to any Active Directory Service server. Verify network connection and check that Active Directory Service server(s) are operational.
ads_failed	STATE_CHANGE	ERROR	local winbindd is unresponsive	local winbindd is unresponsive.	The local winbindd does not respond to ping requests. It is needed for Active Directory Service.	Local winbindd does not respond to ping requests. Try to restart winbindd, and if not successful, perform winbindd troubleshooting.
ads_up	STATE_CHANGE	INFO	external ADS server is up	External ADS server is up.		External Active Directory Service server is operational, no user action required.
ads_warn	INFO	WARNING	external ADS server monitoring returned unknown result	External ADS server monitoring returned unknown result.	An internal error occurred while monitoring the external ADS server.	An internal error occurred while monitoring external Active Directory Service server. Perform trouble check.
ldap_down	STATE_CHANGE	ERROR	external LDAP server {0} is unresponsive	External LDAP server <LDAP server> is unresponsive.	The local node is unable to connect to the LDAP server.	Local node is unable to connect to LDAP server. Verify network connection and check that LDAP server is operational.
ldap_up	STATE_CHANGE	INFO	external LDAP server {0} is up	The external LDAP server is operational.		NA

Table 7. Events for the AUTH component (continued)

Event	EventType	Severity	Message	Description	Cause	User Action
nis_down	STATE_CHANGE	ERROR	external NIS server {0} is unresponsive	External NIS server <NIS server> is unresponsive.	The local node is unable to connect to any NIS server.	Local node is unable to connect to any Network Information Server server. Verify network connection and check that Network Information Service server(s) are operational.
nis_failed	STATE_CHANGE	ERROR	ypbind is unresponsive	ypbind is unresponsive.	The local ypbind daemon does not respond.	Local ypbind daemon does not respond. Try to restart ypbind, and if not successful, perform ypbind troubleshooting.
nis_up	STATE_CHANGE	INFO	external NIS server {0} is up	External Network Information Service (NIS) server is operational.		NA
nis_warn	INFO	WARNING	external NIS monitoring returned unknown result	The external NIS server monitoring returned an unknown result.	An internal error occurred while monitoring external Network Information Service server.	Perform trouble check.
sssd_down	STATE_CHANGE	ERROR	SSSD process not running	The SSSD process not running.	The SSSD authentication service is not running.	Perform trouble check.
sssd_restart	INFO	INFO	SSSD process was not running. Trying to start it	Attempt to start the SSSD authentication process.	The SSSD process was not running.	NA
sssd_up	STATE_CHANGE	INFO	SSSD process is now running	SSSD process is now running.	The SSSD authentication process is running.	NA
sssd_warn	INFO	WARNING	SSSD process monitoring returned unknown result	SSSD authentication process monitoring returned an unknown result.	An internal error occurred while monitoring the SSSD.	Perform trouble check.
wnbd_down	STATE_CHANGE	ERROR	WINBINDD process not running	The WINBINDD authentication process not running .	The WINBINDD authentication service is not running.	Verify the configuration and Active Directory connection.
wnbd_restart	INFO	INFO	WINBINDD process was not running. Trying to start it	Attempt to start the WINBINDD authentication process.	The WINBINDD process was not running.	NA
wnbd_up	STATE_CHANGE	INFO	WINBINDD process is now running	The WINBINDD authentication service is operational.		NA

Table 7. Events for the AUTH component (continued)

Event	EventType	Severity	Message	Description	Cause	User Action
wnbd_warn	INFO	WARNING	WINBINDD process monitoring returned unknown result	WINBINDD process monitoring returned an unknown result.	An internal error occurred during the monitoring of WINBINDD.	Perform trouble check.
yp_down	STATE_CHANGE	ERROR	YPBIND process not running	The YPBIND process not running .	The YPBIND authentication service is not running .	Perform trouble check.
yp_restart	INFO	INFO	YPBIND process was not running. Trying to start it	Attempt to start the YPBIND process.	The YPBIND process is not running.	NA
yp_up	STATE_CHANGE	INFOmight	YPBIND process is now running	The YPBIND service is operational.		NA
yp_warn	INFO	WARNING	YPBIND process monitoring returned unknown result	The YPBIND process monitoring returned an unknown result.	An internal error occurred while monitoring YPBIND.	Perform trouble check.

Table 8. Events for the GPFS component

Event	EventType	Severity	Message	Description	Cause	User Action
cesnodestatechange_info	INFO	INFO	Message: A CES node state change: Node {0} {1} {2} flag	Informational. Shows the modified node state, like the node turned to suspended mode, network down, or others.	A node state change was detected.	Actions might depend on the new node state.
cesquorumloss	STATE_CHANGE	ERROR	CES quorum loss	The cluster got in an inconsistent state.	The number of required quorum nodes does not match the minimum requirements. Reasons might be network or hardware issues.	Recover from the underlying issue. Ensure that the cluster nodes are up and running.
gpfs_down	STATE_CHANGE	ERROR	GPFS process not running	Check the state of the file system daemon.	The file system daemon is not running, but expected to run.	Check for the root cause of this failure in the logs.
gpfs_up	STATE_CHANGE	INFO	GPFS process now running	Check the state of the file system daemon.	The file system daemon is running.	NA
gpfs_warn	INFO	WARNING	GPFS process monitoring returned unknown result.	Check the state of the file system daemon.	The file system daemon state might not be checked due to a problem.	Find potential issues for this kind of failure in the logs.

Table 8. Events for the GPFS component (continued)

Event	EventType	Severity	Message	Description	Cause	User Action
shared_root_bad	STATE_CHANGE	ERROR	Shared root is unavailable	The shared root file system is bad or not available. This file system is required to run the cluster because it stores the cluster-wide information. This problem will trigger a failover.	The CES framework detects the shared_root file system to be unavailable on the node.	Resolve the underlying issue. Check that the shared root file system is mounted.
shared_root_ok	STATE_CHANGE	INFO	Shared root is available	The shared root file system is available. This file system is required to run the cluster, because it stores the cluster-wide information.	The CES framework detects the shared_root file system to be ok.	NA

Table 9. Events for the KEYSTONE component

Event	EventType	Severity	Message	Description	Cause	User action
ks_failed	STATE_CHANGE	ERROR	keystone (httpd) process should be {0} but is {1}	The keystone (httpd) process is in an unexpected mode.	If the object authentication is local, AD or LDAP the process failed unexpectedly. If the object authentication is none or user-defined the process is expected to be stopped but it was not.	Make sure that the process is in the expected state.
ks_ok	STATE_CHANGE	INFO	keystone(httpd) process as expected, state is {0}	The keystone (httpd) process is in the expected state.	If the object authentication is local, AD or LDAP the process is running. If the object authentication is none or user-defined the process stopped as expected.	NA
ks_restart	INFO	WARNING	The {0} service failed. Trying to recover			
ks_url_exfail	STATE_CHANGE	WARNING	Keystone request failed using {0}			
ks_url_failed	STATE_CHANGE	ERROR	Keystone request failed using {0}	A keystone URL request failed.	An HTTP request to keystone failed.	Check that httpd / keystone is running on the expected server and is accessible with the defined ports.

Table 9. Events for the KEYSTONE component (continued)

Event	EventType	Severity	Message	Description	Cause	User action
ks_url_ok	STATE_CHANGE	INFO	Keystone request successfully using {0}	A keystone URL request was successful.	A HTTP request to keystone returned successfully.	NA
ks_url_warn	INFO	WARNING	Keystone request on {0} returned unknown result	A keystone URL request returned an unknown result.	A simple HTTP request to keystone returned with an unexpected error.	Check that httpd / keystone is running on the expected server and is accessible with the defined ports.
ks_warn	INFO	WARNING	keystone(httpd) process monitoring returned unknown result	The keystone (httpd) monitoring returned an unknown result.	A status query for httpd returned an unexpected error.	Check service script and settings of httpd.
postgresql_failed	STATE_CHANGE	ERROR	postgresql-obj process should be {0} but is {1}	The postgresql-obj process is in an unexpected mode.	The database back-end for object authentication is supposed to run on a single node. Either the DB is not running on the designated node or it is running on a different node.	Check that postgresql-obj is running on the expected server.
postgresql_ok	STATE_CHANGE	INFO	postgresql-obj process as expected, state is {0}	The postgresql-obj process is in the expected mode.	The database back-end for object authentication is supposed to run on the right node while being stopped on others.	NA
postgresql_warn	INFO	WARNING	postgresql-obj process monitoring returned unknown result	The postgresql-obj process monitoring returned an unknown result	A status query for postgresql-obj returned with an unexpected error.	Check postgres database engine.

Table 10. Events for the NFS component

Event	EventType	Severity	Message	Description	Cause	User Action
dbus_error	STATE_CHANGE	WARNING	DBus availability check failed	The NFS service is registered to DBus, and DBus is used to send export related information to this server.	The DBus was detected as down. This might cause several issues on the local node.	Stop the NFS service, restart the DBus, and start the NFS service again.
disable_nfs_service	INFO	INFO	Ganesha NFS service was disabled	The NFS service was disabled on this node. Disabling a Service means, that also all configuration files are removed. This is different from stopping a running service.	The user has executed 'mmces service disable nfs'.	NA

Table 10. Events for the NFS component (continued)

Event	EventType	Severity	Message	Description	Cause	User Action
enable_nfs_service	INFO	INFO	Ganesha NFS service was enabled	The NFS service was enabled on this node. Enabling a protocol service means, that also all required configuration files are automatically installed with the current valid configuration settings.	The user has executed 'mmces service enable nfs'	NA
ganeshaexit	INFO	INFO	Ganesha NFS was stopped	An NFS server instance has terminated.	A NFS instance was terminated or was killed somehow.	Restart the NFS service when the root cause for this issue is solved.
ganeshaGrace	INFO	INFO	Ganesha NFS is set to grace	The NFS server is sent to grace mode for a limited time. This gives previously connected clients time to recover their file locks.	The grace period is always cluster wide. NFS export configurations might have changed, and one or more NFS servers were restarted.	NA
nfs3_down	INFO	WARNING	NFS v3 check returned down	The NFS v3 NULL check failed when expected. The NFS v3 NULL check is done to see if the NFS server reacts on NFS v3 requests. The v3 protocol must be enabled for this check. If this down state is detected, further checks are done to figure out if the NFS server is working. If the NFS server seems to be not working, then a failover is triggered. If NFS v3 and NFS v4 protocols are configured, then only the v3 NULL test is executed.	The NFS server might hang or is under high load so that the request might not be processed.	Check the health state of the NFS server and restart, if necessary.
nfs3_up	INFO	INFO	NFS v4 check returned up	The NFS v4 NULL check was successful.	The NFS v4 NULL check works as expected.	NA



Table 10. Events for the NFS component (continued)

Event	EventType	Severity	Message	Description	Cause	User Action
nfs4_down	INFO	WARNING	NFS v4 check returned down	The NFS v4 NULL check failed when expected. The NFS v4 NULL check is done to see whether the NFS server reacts on NFS v4 requests. The v4 protocol must be enabled for this check. If this down state is detected, further checks are done to figure out if the NFS server is working. If the NFS server seems to be not working, then a failover is triggered.	The NFS server might hang or is under high load so that the request might not be processed.	Check the health state of the NFS server and restart, if necessary.
nfs4_up	INFO	INFO	NFS v4 check returned up	The NFS v4 NULL check was successful.	The NFS v4 NULL check works as expected.	NA
nfs_active	STATE_CHANGE	INFO	NFS is now active	The NFS service must be up and running, and in a healthy state to provide the configured file exports.	The NFS server was detected as alive (again).	NA
nfs_dbus_error	STATE_CHANGE	WARNING	NFS check via DBus failed	The NFS service must be registered on DBus to be fully working.	The NFS service is registered on DBus, but there was a problem accessing it.	Check the health state of the NFS service, restart the NFS service. Check the log files for reported issues.
nfs_dbus_failed	STATE_CHANGE	WARNING	NFS check via DBus did not return expected message	NFS service configuration settings (log configuration settings) are queried via DBus. The result is checked for expected keywords.	The NFS service is registered on DBus, but the check via DBus did not return the expected result.	Stop the NFS service and start it again. Check the log configuration of the NFS service.
nfs_dbus_ok	STATE_CHANGE	INFO	NFS check via DBus successful	Check that the NFS service is registered on DBus and working.	The NFS service is registered on DBus and working.	NA
nfs_in_grace	STATE_CHANGE	WARNING	NFS in grace mode	The monitor detected that Ganesha is in grace mode. During this time, the ganesha state is shown as degraded.	The NFS service was started or restarted.	NA
nfs_not_active	STATE_CHANGE	ERROR	NFS is not active	A check showed that a running Ganesha instance shows no activity at all.	Process might have hung.	Restart Ganesha.

Table 10. Events for the NFS component (continued)

Event	EventType	Severity	Message	Description	Cause	User Action
nfs_not_dbus	STATE_CHANGE	WARNING	NFS service not available as Dbus service. Consider restart of NFS server.	The NFS service is currently not registered on Dbus. In this mode, the NFS service is not fully working. Exports cannot be added or removed, and not set in grace mode, which is important for data consistency.	The NFS service might have been started while the Dbus was down.	Stop the NFS service, restart the Dbus, and start the NFS service again.
nfsd_down	STATE_CHANGE	ERROR	NFSD process not running	Checks for an NFS service process.	The NFS server process was not detected.	Check the health state of the NFS server and restart, if necessary. The process might hang or is in failed state.
nfsd_up	STATE_CHANGE	INFO	NFSD process now running	Checks for a NFS service process.	The NFS server process was detected. Some further checks are done then.	NA
nfsd_warn	INFO	WARNING	NFSD process monitoring returned unknown result	Checks for a NFS service process.	The NFS server process state might not be determined due to a problem.	Check the health state of the NFS server and restart, if necessary.
portmapper_down	STATE_CHANGE	ERROR	Portmapper port 111 is not active	The portmapper is needed to provide the NFS services to clients.	The portmapper is not running on port 111.	NA
portmapper_up	STATE_CHANGE	INFO	Portmapper port is now active	The portmapper is needed to provide the NFS services to clients.	The portmapper is running on port 111.	NA
portmapper_warn	INFO	WARNING	Portmapper port monitoring (111) returned unknown result	The portmapper is needed to provide the NFS services to clients.	The portmapper status might not be determined due to a problem.	Restart the portmapper, if necessary.
postIpChange_info	INFO	INFO	IP addresses modified (post change)	Information that IP addresses are moved around the cluster nodes.	CES IP addresses were moved or added to the node, and activated.	NA
rquotad_down	INFO	INFO	rpc.rquotad not running	Currently not in use. Future.	NA	NA
rquotad_up	INFO	INFO	rpc.rquotad is running	Currently not in use. Future.		
start_nfs_service	INFO	INFO	Ganesha NFS service was started	Information about a NFS service start.	The NFS service was started (like 'mmces service start nfs').	NA
statd_down	STATE_CHANGE	ERROR	rpc.statd is not running	The statd process is used by NFS v3 to handle file locks.	The statd process is not running.	Stop and start the NFS service. This attempts to start the statd process also.
statd_up	STATE_CHANGE	INFO	rpc.statd is running	The statd process is used by NFS v3 to handle file locks.	The statd process is running.	NA
stop_nfs_service	INFO	INFO	Ganesha NFS service was stopped	Information about a NFS service stop.	The NFS service was stopped (like 'mmces service stop nfs').	NA

Table 11. Events for the Network component

Event	EventType	Severity	Message	Description	Cause	User Action
bond_degraded	STATE_CHANGE	INFO	Some slaves of the bond {0} went down	Some of the bond parts are malfunctioning		Check the network configuration and cabling of the relevant network adapters
bond_down	STATE_CHANGE	ERROR	All slaves of the bond {0} went down	All slaves of a network bond went down.	There could be hard- and software related issues.	Check the network configuration and cabling of the relevant network adapters.
bond_up	STATE_CHANGE	INFO	All slaves of the bond {0} are working as expected	The bond is functioning properly		NA
ces_disable_node network	INFO	INFO	Network was disabled	Clean up after a 'mmchnode --ces-disable' command. The network configuration is modified accordingly.	Clean up after a 'mmchnode --ces-disable' command.	NA
ces_enable_node network	INFO	INFO	Network was disabled	Called to handle any network- sepcific issues involved after a 'mmchnode --ces-enable' command. The network configuration is modified accordingly.	Called after a 'mmchnode --ces-enable' command.	NA
ces_startup_network	INFO	INFO	CES network service was started	Information that the CES network has started.	CES network IPs are started.	NA
handle_network_problem_info	INFO	INFO	Handle network problem - Problem: {0}, Argument: {1}	Information about network- related reconfigurations. This can be enable or disable IPs, assign or unassign IPs for example.	A change in the network configuration. Details are part of the information message.	NA
many_tx_errors	STATE_CHANGE	ERROR	NIC {0} had many TX errors since the last monitoring cycle	The network adapter had many TX errors since the last monitoring cycle.	The cabling is most likely damage.	Check cable contacts or try a different cable.
move_cesip_from	INFO	INFO	Address {0} was moved from this node to node {1}	Information that a CES IP address was moved from the current node to another node.	Rebalancing of CES IP addresses.	NA
move_cesip_to	INFO	INFO	Address {0} was moved from node {1} to this node	Information that a CES IP address was moved from another node to the current node.	Rebalancing of CES IP addresses.	NA
move_cesips_infos	INFO	INFO	A move request for ip addresses was executed	CES IP addresses can be moved in case of node failovers from one node to one or more other nodes. This message is logged on a node observing this, not necessarily on any affected node.	A CES IP movement was detected.	NA

Table 11. Events for the Network component (continued)

Event	EventType	Severity	Message	Description	Cause	User Action
network_connectivity_down	STATE_CHANGE	ERROR	NIC {0} can not connect to the gateway	The network adapter can not connect to the gateway.	There could be hard- and software related issues; gateway may be down.	Check the network configuration of the network adapter, path to the gateway and gateway itself.
network_connectivity_up	STATE_CHANGE	INFO	NIC {0} can connect to the gateway	The network adapter can connect to the gateway.		NA
network_down	STATE_CHANGE	ERROR	Network is down	The network is down.	There might be hardware and software-related issues	Check for network-related issues, network cards, bonds, cabling, configurations, and so on.
network_found	INFO	INFO	NIC {0} was found	A new network adapter was found.		NA
network_ips_down	STATE_CHANGE	ERROR	No relevant NICs detected	No relevant network adapters detected		Check the IPs, which are relevant for IBM Spectrum Scale"
network_ips_up	STATE_CHANGE	INFO	Relevant IPs are served by found NICs	Relevant IPs are served by network adapters		NA
network_link_down	STATE_CHANGE	ERROR	Physical link of the NIC {0} is down	The physical link of the adapter is down.	There could be hard- and software related issues.	Check the network configuration and cabling of the network adapter.
network_link_up	STATE_CHANGE	INFO	Physical link of the NIC {0} is up	The physical link of the adapter is up.		NA
network_up	STATE_CHANGE	INFO	Network is up	The Network is running.		NA
network_vanished	INFO	INFO	NIC {0} has vanished	One of network adapters can not be detected anymore.	Network configuration changes	NA
no_tx_errors	STATE_CHANGE	INFO	NIC {0} had no or an insignificant number of TX errors	The NIC had no or an insignificant number of TX errors.		NA

Table 12. Events for the Object component

Event	EventType	Severity	Message	Description	Cause	User Action
account-auditor_failed	STATE_CHANGE	ERROR	account-auditor process should be {0} but is {1}	The account-auditor process is not in the expected state.	The account-auditor process is expected to be running on the singleton node only.	Check the status of openstack-swift-account-auditor process and object singleton flag.
account-auditor_ok	STATE_CHANGE	INFO	account-auditor process as expected, state is {0}	The account-auditor process is in the expected state.	The account-auditor process is expected to be running on the singleton node only.	NA
account-auditor_warn	INFO	WARNING	account-auditor process monitoring returned unknown result	The account-auditor check returned an unknown result.	A status query for openstack-swift-account-auditor returned with an unexpected error.	Check service script and settings.

Table 12. Events for the Object component (continued)

Event	EventType	Severity	Message	Description	Cause	User Action
account-reaper_failed	STATE_CHANGE	ERROR	account-reaper process should be {0} but is {1}	The account-reaper process is not running.	The account-reaper process is not running.	Check the status of openstack-swift-account-reaper process.
account-reaper_ok	STATE_CHANGE	INFO	account-reaper process as expected, state is {0}	The account-reaper process is running.	The account-reaper process is running.	NA
account-reaper_warn	INFO	WARNING	account-reaper process monitoring returned unknown result	The account-reaper check returned an unknown result.	A status query for openstack-swift-account-reaper returned with an unexpected error.	Check service script and settings.
account-replicator_failed	STATE_CHANGE	ERROR	account-replicator process should be {0} but is {1}	The account-replicator process is not running.	The account-replicator process is not running.	Check the status of openstack-swift-account-replicator process.
account-replicator_ok	STATE_CHANGE	INFO	account-replicator process as expected, state is {0}	The account-replicator process is running.	The account-replicator process is running.	NA
account-replicator_warn	INFO	WARNING	account-replicator process monitoring returned unknown result	The account-replicator check returned an unknown result.	A status query for openstack-swift-account-replicator returned with an unexpected error.	Check the service script and settings.
account-server_failed	STATE_CHANGE	ERROR	account process should be {0} but is {1}	The account-server process is not running.	The account-server process is not running.	Check the status of openstack-swift-account process.
account-server_ok	STATE_CHANGE	INFO	account process as expected, state is {0}	The account-server process is running.	The account-server process is running.	NA
account-server_warn	INFO	WARNING	account process monitoring returned unknown result	The account-server check returned an unknown result.	A status query for openstack-swift-account returned with an unexpected error.	Check the service script and settings.
container-auditor_failed	STATE_CHANGE	ERROR	container-auditor process should be {0} but is {1}	The container-auditor process is not in the expected state.	The container-auditor process is expected to be running on the singleton node only.	Check the status of openstack-swift-container-auditor process and object singleton flag.
container-auditor_ok	STATE_CHANGE	INFO	container-auditor process as expected, state is {0}	The container-auditor process is in the expected state.	The container-auditor process is expected to be running on the singleton node only.	NA
container-auditor_warn	INFO	WARNING	container-auditor process monitoring returned unknown result	The container-auditor check returned an unknown result.	A status query for openstack-swift-container-auditor returned with an unexpected error.	Check service script and settings.
container-replicator_failed	STATE_CHANGE	ERROR	container-replicator process should be {0} but is {1}	The container-replicator process is not running.	The container-replicator process is not running.	Check the status of openstack-swift-container-replicator process.

Table 12. Events for the Object component (continued)

Event	EventType	Severity	Message	Description	Cause	User Action
container-replicator_ok	STATE_CHANGE	INFO	container-replicator process as expected, state is {0}	The container-replicator process is running.	The container-replicator process is running.	NA
container-replicator_warn	INFO	WARNING	container-replicator process monitoring returned unknown result	The container-replicator check returned an unknown result.	A status query for openstack-swift-container-replicator returned with an unexpected error.	Check service script and settings.
container-server_failed	STATE_CHANGE	ERROR	container process should be {0} but is {1}	The container-server process is not running.	The container-server process is not running.	Check the status of openstack-swift-container process.
container-server_ok	STATE_CHANGE	INFO	container process as expected, state is {0}	The container-server process is running.	The container-server process is running.	NA
container-server_warn	INFO	WARNING	container process monitoring returned unknown result	The container-server check returned an unknown result.	A status query for openstack-swift-container returned with an unexpected error.	Check the service script and settings.
container-updater_failed	STATE_CHANGE	ERROR	container-updater process should be {0} but is {1}	The container-updater process is not in the expected state.	The container-updater process is expected to be running on the singleton node only.	Check the status of openstack-swift-container-updater process and object singleton flag.
container-updater_ok	STATE_CHANGE	INFO	container-updater process as expected, state is {0}	The container-updater process is in the expected state.	The container-updater process is expected to be running on the singleton node only.	NA
container-updater_warn	INFO	WARNING	container-updater process monitoring returned unknown result	The container-updater check returned an unknown result.	A status query for openstack-swift-container-updater returned with an unexpected error.	Check the service script and settings.
disable_Address_database_node	INFO	INFO	Disable Address Database Node	Event to signal that the database flag was removed from this node.	A CES IP with a singleton/database flag linked to it was removed/moved from/to this node.	NA
disable_Address_singleton_node	INFO	INFO	Disable Address Singleton Node	Event to signal that the singleton flag was removed from this.	A CES IP with a singleton/database flag linked to it was removed/moved from/to this node.	NA
enable_Address_database_node	INFO	INFO	Enable Address Database Node	Event to signal that the database flag was moved to this node.	A CES IP with a singleton/database flag linked to it was removed/moved from/to this node.	NA

Table 12. Events for the Object component (continued)

Event	EventType	Severity	Message	Description	Cause	User Action
enable_Address_singleton_node	INFO	INFO	Enable Address Singleton Node	Event to signal that the singleton flag was moved to this node.	A CES IP with a singleton/database flag linked to it was removed/moved from/to this node.	NA
ibmobjectizer_failed	STATE_CHANGE	ERROR	ibmobjectizer process should be {0} but is {1}	The ibmobjectizer process is in the expected state.	The ibmobjectizer process is expected to be running on the singleton node only.	Check the status of the ibmobjectizer process and object singleton flag.
ibmobjectizer_ok	STATE_CHANGE	INFO	ibmobjectizer process as expected, state is {0}	The ibmobjectizer process is not in the expected state.	The ibmobjectizer process is expected to be running on the singleton node only.	NA
ibmobjectizer_warn	INFO	WARNING	ibmobjectizer process monitoring returned unknown result	The ibmobjectizer check returned an unknown result.	A status query for ibmobjectizer returned with an unexpected error.	Check the service script and settings.
memcached_failed	STATE_CHANGE	ERROR	memcached process should be {0} but is {1}	The memcached process is not running.	The memcached process is not running.	Check the status of memcached process.
memcached_ok	STATE_CHANGE	INFO	memcached process as expected, state is {0}	The memcached process is running.	The memcached process is running.	NA
memcached_warn	INFO	WARNING	memcached process monitoring returned unknown result	The memcached check returned an unknown result.	A status query for memcached returned with an unexpected error.	Check the service script and settings.
obj_restart	INFO	WARNING	The {0} service failed. Trying to recover			
object-expirer_failed	STATE_CHANGE	ERROR	object-expirer process should be {0} but is {1}	The object-expirer process is not in the expected state.	The object-expirer process is expected to be running on the singleton node only.	Check the status of openstack-swift-object-expirer process and object singleton flag.
object-expirer_ok	STATE_CHANGE	INFO	object-expirer process as expected, state is {0}	The object-expirer process is in the expected state.	The object-expirer process is expected to be running on the singleton node only.	NA
object-expirer_warn	INFO	WARNING	object-expirer process monitoring returned unknown result	The object-expirer check returned an unknown result.	A status query for openstack-swift-object-expirer returned with an unexpected error.	Check the service script and settings.
object-replicator_failed	STATE_CHANGE	ERROR	object-replicator process should be {0} but is {1}	The object-replicator process is not running.	The object-replicator process is not running.	Check the status of openstack-swift-object-replicator process.
object-replicator_ok	STATE_CHANGE	INFO	object-replicator process as expected, state is {0}	The object-replicator process is running.	The object-replicator process is running.	NA

Table 12. Events for the Object component (continued)

Event	EventType	Severity	Message	Description	Cause	User Action
object-replicator_warn	INFO	WARNING	object-replicator process monitoring returned unknown result	The object-replicator check returned an unknown result.	A status query for openstack-swift-object-replicator returned with an unexpected error.	Check the service script and settings.
object-server_failed	STATE_CHANGE	ERROR	object process should be {0} but is {1}	The object-server process is not running.	The object-server process is not running.	Check the status of the openstack-swift-object process.
object-server_ok	STATE_CHANGE	INFO	object process as expected, state is {0}	The object-server process is running.	The object-server process is running.	NA
object-server_warn	INFO	WARNING	object process monitoring returned unknown result	The object-server check returned an unknown result.	A status query for openstack-swift-object-server returned with an unexpected error.	Check the service script and settings.
object-updater_failed	STATE_CHANGE	ERROR	object-updater process should be {0} but is {1}	The object-updater process is not in the expected state.	The object-updater process is expected to be running on the singleton node only.	Check the status of the openstack-swift-object-updater process and object singleton flag.
object-updater_ok	STATE_CHANGE	INFO	object-updater process as expected, state is {0}	The object-updater process is in the expected state.	The object-updater process is expected to be running on the singleton node only.	NA
object-updater_warn	INFO	WARNING	object-updater process monitoring returned unknown result	The object-updater check returned an unknown result	A status query for openstack-swift-object-updater returned with an unexpected error.	Check the service script and settings.
openstack-object-sof_failed	STATE_CHANGE	ERROR	object-sof process should be {0} but is {1}	The swift-on-file process is not in the expected state.	The swift-on-file process is expected to be running then the capability is enabled and stopped when disabled.	Check the status of the openstack-swift-object-sof process and capabilities flag in spectrum-scale-object.conf.
openstack-object-sof_ok	STATE_CHANGE	INFO	object-sof process as expected, state is {0}	The swift-on-file process is in the expected state.	The swift-on-file process is expected to be running then the capability is enabled and stopped when disabled.	NA
openstack-object-sof_warn	INFO	INFO	object-sof process monitoring returned unknown result	The openstack-swift-object-sof check returned an unknown result.	A status query for openstack-swift-object-sof returned with an unexpected error.	Check the service script and settings.
postIpcChange_info	INFO	INFO	IP addresses modified {0}	CES IP addresses have been moved and activated.		NA
proxy-server_failed	STATE_CHANGE	ERROR	proxy process should be {0} but is {1}	The proxy-server process is not running.	The proxy-server process is not running.	Check the status of the openstack-swift-proxy process.
proxy-server_ok	STATE_CHANGE	INFO	proxy process as expected, state is {0}	The proxy-server process is running.	The proxy-server process is running.	NA



Table 12. Events for the Object component (continued)

Event	EventType	Severity	Message	Description	Cause	User Action
proxy-server_warn	INFO	WARNING	proxy process monitoring returned unknown result	The proxy-server process monitoring returned an unknown result.	A status query for openstack-swift-proxy-server returned with an unexpected error.	Check the service script and settings.
ring_checksum_failed	STATE_CHANGE	ERROR	Checksum of ring file {0} does not match the one in CCR	Files for object rings have been modified unexpectedly.	Checksum of file did not match the stored value.	Check the ring files.
ring_checksum_ok	STATE_CHANGE	INFO	Checksum of ring file {0} is OK	Files for object rings were successfully checked.	Checksum of file found unchanged.	NA
ring_checksum_warn	INFO	WARNING	Issue while checking checksum of ring file {0}	Checksum generation process failed.	The ring_checksum check returned an unknown result.	Check the ring files and the md5sum executable.

Table 13. Events for the SMB component

Event	EventType	Severity	Message	Description	Cause	User Action
ctdb_down	STATE_CHANGE	ERROR	CTDB process not running	The CTDB process is not running.		Perform trouble check.
ctdb_recovered	STATE_CHANGE	INFO	STATE_CHANGE	CTDB completed database recovery.		NA
ctdb_recovery	STATE_CHANGE	WARNING	CTDB Recovery detected	CTDB is performing a database recovery.		NA
ctdb_state_down	STATE_CHANGE	ERROR	CTDB state is {0}	The CTDB state is unhealthy.		Perform trouble check.
ctdb_state_up	STATE_CHANGE	INFO	CTDB state is healthy	The CTDB state is healthy.		NA
ctdb_up	STATE_CHANGE	INFO	CTDB process now running	The CTDB process is running.		NA
ctdb_warn	INFO	WARNING	CTDB monitoring returned unknown result	The CTDB check returned unknown result.		Perform trouble check.
smb_restart	INFO	WARNING	The SMB service failed. Trying to recover	Attempt to start the SMBD process.	The SMBD process was not running.	NA
smbd_down	STATE_CHANGE	ERROR	SMBD process not running	The SMBD process is not running.		Perform trouble check.
smbd_up	STATE_CHANGE	INFO	SMBD process now running	The SMBD process is running.		NA
smbd_warn	INFO	WARNING	SMBD process monitoring returned unknown result	The SMBD process monitoring returned an unknown result.		Perform trouble check.
smbport_down	STATE_CHANGE	ERROR	SMB port {0} is not active	SMBD is not listening on a TCP protocol port.		Perform trouble check.
smbport_up	STATE_CHANGE	INFO	SMB port {0} is now active	An SMB port was activated.		NA
smbport_warn	INFO	WARNING	SMB port monitoring {0} returned unknown result	An internal error occurred while monitoring SMB TCP protocol.		Perform trouble check.



---

## Chapter 13. Contacting IBM support center

Specific information about a problem such as: symptoms, traces, error logs, GPFS 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 be collected before contacting the IBM Support Center”
- “How to contact the IBM Support Center” on page 169.

---

### Information to be collected 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.

#### Information to be collected for all problems related to GPFS

Regardless of the problem encountered with GPFS, the following 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. A tar file generated by the **gpfs.snap** command that contains data from the nodes in the cluster. In large clusters, the **gpfs.snap** command can collect data from certain nodes (for example, the affected nodes, NSD servers, or manager nodes) using the **-N** option.

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

- a. Any error log entries relating to the event:
  - On an AIX node, issue this command:  

```
errpt -a
```
  - On a Linux node, create a tar file of all the entries in the **/var/log/messages** file from all nodes in the cluster or the nodes that experienced the failure. For example, issue the following command to create a tar file that includes all nodes in the cluster:  

```
mmdash -v -N all "cat /var/log/messages" > all.messages
```
  - On a Windows node, use the **Export List...** dialog in the Event Viewer to save the event log to a file.
- b. A master GPFS log file that is merged and chronologically sorted for the date of the failure (see “Creating a master GPFS log file” on page 2).
- c. 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**.
- d. On a failing Linux node, gather the installed software packages and the versions of each package by issuing this command:  

```
rpm -qa
```
- e. On a failing AIX node, gather the name, most recent level, state, and description of all installed software packages by issuing this command:  

```
lslpp -l
```
- f. File system attributes for all of the failing file systems, issue:  

```
mmfsfs Device
```

- g. The current configuration and state of the disks for all of the failing file systems, issue:
 

```
mmlsdisk Device
```
  - h. A copy of file `/var/mmfs/gen/mmsdrfs` from the primary cluster configuration server.
4. For Linux on z Systems, collect the data of the operating system as described in the *Linux on z Systems Troubleshooting Guide* ([www.ibm.com/support/knowledgecenter/linuxonibm/liaaf/lnz\\_r\\_sv.html](http://www.ibm.com/support/knowledgecenter/linuxonibm/liaaf/lnz_r_sv.html)).
  5. If you are experiencing one of the following problems, see the appropriate section before contacting the IBM Support Center:
    - For delay and deadlock issues, see “Additional information to collect for delays and deadlocks.”
    - For file system corruption or MMFS\_FSSTRUCT errors, see “Additional information to collect for file system corruption or MMFS\_FSSTRUCT errors.”
    - For GPFS daemon crashes, see “Additional information to collect for GPFS daemon crashes.”

## Additional information to collect for delays and deadlocks

When a delay or deadlock situation is suspected, the IBM Support Center will need additional information to assist with problem diagnosis. If you have not done so already, ensure you have the following information available before contacting the IBM Support Center:

1. Everything that is listed in “Information to be collected for all problems related to GPFS” on page 167.
2. The deadlock debug data collected automatically.
3. If the cluster size is relatively small and the `maxFilesToCache` setting is not high (less than 10,000), issue the following command:

```
gpfs.snap --deadlock
```

If the cluster size is large or the `maxFilesToCache` setting is high (greater than 1M), issue the following command:

```
gpfs.snap --deadlock --quick
```

## Additional information to collect for file system corruption or MMFS\_FSSTRUCT errors

When file system corruption or MMFS\_FSSTRUCT errors are encountered, the IBM Support Center will need additional information to assist with problem diagnosis. If you have not done so already, ensure you have the following information available before contacting the IBM Support Center:

1. Everything that is listed in “Information to be collected for all problems related to GPFS” on page 167.
2. Unmount the file system everywhere, then run `mmfsck -n` in offline mode and redirect it to an output file.

The IBM Support Center will determine when and if you should run the `mmfsck -y` command.

## Additional information to collect for GPFS daemon crashes

When the GPFS daemon is repeatedly crashing, the IBM Support Center will need additional information to assist with problem diagnosis. If you have not done so already, ensure you have the following information available before contacting the IBM Support Center:

1. Everything that is listed in “Information to be collected for all problems related to GPFS” on page 167.
2. Ensure the `/tmp/mmfs` directory exists on all nodes. If this directory does not exist, the GPFS daemon will not generate internal dumps.
3. Set the traces on this cluster and *all* clusters that mount any file system from this cluster:

```
mmtracectl --set --trace=def --trace-recycle=global
```

4. Start the trace facility by issuing:

```
mmtracectl --start
```

5. Recreate the problem if possible or wait for the assert to be triggered again.
6. Once the assert is encountered on the node, turn off the trace facility by issuing:
 

```
mmtracectl --off
```

 If traces were started on multiple clusters, **mmtracectl --off** should be issued immediately on all clusters.
7. Collect **gpfs.snap** output:
 

```
gpfs.snap
```

---

## How to contact the IBM Support Center

The IBM Support Center is available for various types of IBM hardware and software problems that GPFS customers may encounter.

These problems include the following:

- IBM hardware failure
- Node halt or crash not related to a hardware failure
- Node hang or response problems
- Failure in other software supplied by IBM

### If you have an IBM Software Maintenance service contract

If you have an IBM Software Maintenance service contract, contact the IBM Support Center, as follows:

Your location	Method of contacting the IBM Support Center
In the United States	Call <b>1-800-IBM-SERV</b> for support.
Outside the United States	Contact your local IBM Support Center or see the Directory of worldwide contacts ( <a href="http://www.ibm.com/planetwide">www.ibm.com/planetwide</a> ).

When you contact the IBM Support Center, the following will occur:

1. You will be asked for the information you collected in "Information to be collected before contacting the IBM Support Center" on page 167.
2. You will be given a time period during which an IBM representative will return your call. Be sure that the person you identified as your contact can be reached at the phone number you provided in the PMR.
3. An online Problem Management Record (PMR) will be created to track the problem you are reporting, and you will be advised to record the PMR number for future reference.
4. You may be requested to send data related to the problem you are reporting, using the PMR number to identify it.
5. Should you need to make subsequent calls to discuss the problem, you will also use the PMR number to identify the problem.

### If you do not have an IBM Software Maintenance service contract

If you do not have an IBM Software Maintenance service contract, contact your IBM sales representative to find out how to proceed. Be prepared to provide the information you collected in "Information to be collected before contacting the IBM Support Center" on page 167.

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



---

## Chapter 14. Message severity tags

GPFS has adopted a message severity tagging convention. This convention applies to some newer messages and to some messages that are being updated and adapted to be more usable by scripts or semi-automated management programs.

A severity tag is a one-character alphabetic code (**A** through **Z**), optionally followed by a colon (**:**) and a number, and surrounded by an opening and closing bracket (**[ ]**). For example:

**[E]** or **[E:nnn]**

If more than one substring within a message matches this pattern (for example, **[A]** or **[A:nnn]**), the severity tag is the first such matching string.

When the severity tag includes a numeric code (*nnn*), this is an error code associated with the message. If this were the only problem encountered by the command, the command return code would be *nnn*.

If a message does not have a severity tag, the message does not conform to this specification. You can determine the message severity by examining the text or any supplemental information provided in the message catalog, or by contacting the IBM Support Center.

Each message severity tag has an assigned priority that can be used to filter the messages that are sent to the error log on Linux. Filtering is controlled with the **mmchconfig** attribute **systemLogLevel**. The default for **systemLogLevel** is **error**, which means GPFS will send all error **[E]**, critical **[X]**, and alert **[A]** messages to the error log. The values allowed for **systemLogLevel** are: **alert**, **critical**, **error**, **warning**, **notice**, **configuration**, **informational**, **detail**, or **debug**. Additionally, the value **none** can be specified so no messages are sent to the error log.

Alert **[A]** messages have the highest priority, and debug **[B]** messages have the lowest priority. If the **systemLogLevel** default of **error** is changed, only messages with the specified severity and all those with a higher priority are sent to the error log. The following table lists the message severity tags in order of priority:

Table 14. Message severity tags ordered by priority

Severity tag	Type of message (systemLogLevel attribute)	Meaning
<b>A</b>	<b>alert</b>	Indicates a problem where action must be taken immediately. Notify the appropriate person to correct the problem.
<b>X</b>	<b>critical</b>	Indicates a critical condition that should be corrected immediately. The system discovered an internal inconsistency of some kind. Command execution might be halted or the system might attempt to continue despite the inconsistency. Report these errors to the IBM Support Center.
<b>E</b>	<b>error</b>	Indicates an error condition. Command execution might or might not continue, but this error was likely caused by a persistent condition and will remain until corrected by some other program or administrative action. For example, a command operating on a single file or other GPFS object might terminate upon encountering any condition of severity <b>E</b> . As another example, a command operating on a list of files, finding that one of the files has permission bits set that disallow the operation, might continue to operate on all other files within the specified list of files.

Table 14. Message severity tags ordered by priority (continued)

Severity tag	Type of message (systemLogLevel attribute)	Meaning
<b>W</b>	<b>warning</b>	Indicates a problem, but command execution continues. The problem can be a transient inconsistency. It can be that the command has skipped some operations on some objects, or is reporting an irregularity that could be of interest. For example, if a multipass command operating on many files discovers during its second pass that a file that was present during the first pass is no longer present, the file might have been removed by another command or program.
<b>N</b>	<b>notice</b>	Indicates a normal but significant condition. These events are unusual but not error conditions, and might be summarized in an email to developers or administrators for spotting potential problems. No immediate action is required.
<b>C</b>	<b>configuration</b>	Indicates a configuration change; such as, creating a file system or removing a node from the cluster.
<b>I</b>	<b>informational</b>	Indicates normal operation. This message by itself indicates that nothing is wrong; no action is required.
<b>D</b>	<b>detail</b>	Indicates verbose operational messages; no is action required.
<b>B</b>	<b>debug</b>	Indicates debug-level messages that are useful to application developers for debugging purposes. This information is not useful during operations.



---

## Chapter 15. Messages

This topic contains explanations for GPFS error messages.

Messages for GPFS Native RAID in the ranges 6027-1850 – 6027-1899 and 6027-3000 – 6027-3099 are documented in *IBM Spectrum Scale RAID: Administration*.

---

**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 [N]**   **mmfsd ready**

**Explanation:** The **mmfsd** server is up and running.

**User response:** None. Informational message only.

---

**6027-301**      **File *fileName* could not be run with error *errno*.**

**Explanation:** The named shell script could not be executed. This message is followed by the error string that is returned by the **exec**.

**User response:** Check file existence and access permissions.

---

**6027-302 [E]**   **Could not execute *script***

**Explanation:** The **verifyGpfsReady=yes** configuration attribute is set, but the **/var/mmfs/etc/gpfsready** script could not be executed.

**User response:** Make sure **/var/mmfs/etc/gpfsready** exists and is executable, or disable the **verifyGpfsReady** option via **mmchconfig verifyGpfsReady=no**.

---

**6027-303 [N]**   ***script* killed by signal *signal***

**Explanation:** The **verifyGpfsReady=yes** configuration attribute is set and **/var/mmfs/etc/gpfsready** script did not complete successfully.

**User response:** Make sure **/var/mmfs/etc/gpfsready** completes and returns a zero exit status, or disable the **verifyGpfsReady** option via **mmchconfig verifyGpfsReady=no**.

---

**6027-304 [W]**   ***script* ended abnormally**

**Explanation:** The **verifyGpfsReady=yes** configuration attribute is set and **/var/mmfs/etc/gpfsready** script did not complete successfully.

**User response:** Make sure **/var/mmfs/etc/gpfsready** completes and returns a zero exit status, or disable the **verifyGpfsReady** option via **mmchconfig verifyGpfsReady=no**.

---

**6027-305 [N]**   ***script* failed with exit code *code***

**Explanation:** The **verifyGpfsReady=yes** configuration attribute is set and **/var/mmfs/etc/gpfsready** script did not complete successfully.

**User response:** Make sure **/var/mmfs/etc/gpfsready** completes and returns a zero exit status, or disable the **verifyGpfsReady** option via **mmchconfig verifyGpfsReady=no**.

---

**6027-306 [E]**   **Could not initialize inter-node 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 [I]**   ***command* initializing. {*Version*  
*versionName*: **Built** *date time*}**

**Explanation:** The **mmfsd** server has started execution.

**User response:** None. Informational message only.

---

**6027-311 [N]**   ***programName* is shutting down.**

**Explanation:** The stated program is about to terminate.

**User response:** None. Informational message only.

---

**6027-312 [E]**   **Unknown trace class '*traceClass*'.**

**Explanation:** The trace class is not recognized.

**User response:** Specify a valid trace class.

---

**6027-313 [X] 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 [E] 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 [E] 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 [E] 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 [A] 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 [E] 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 [A] 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 [W] 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/gen/mmfs.cfg` file.

**User response:** Check the `/var/mmfs/gen/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 [E]**   **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 [E]**   **Configuration error: check *fileName*.**

**Explanation:** A configuration error is found.

**User response:** Check the `mmfs.cfg` file and other error messages.

---

**6027-336 [E]**   **Value '*value*' for configuration parameter '*parameter*' is not valid. Check *fileName*.**

**Explanation:** A configuration error was found.

**User response:** Check the `mmfs.cfg` file.

---

**6027-337 [N]**   **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 [N]**   **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 [E]**   **Nonnumeric trace value '*value*' after class '*class*'.**

**Explanation:** The specified trace value is not recognized.

**User response:** Specify a valid trace integer value.

---

**6027-340**     **Child process *file* failed to start due to error *rc*: *errStr*.**

**Explanation:** A failure occurred when GPFS attempted to start a program.

**User response:** If the program was a user exit script, verify the script file exists and has appropriate permissions assigned. If the program was not a user exit script, then this is an internal GPFS error or the GPFS installation was altered.

---

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

**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 use by the IBM Support Center.

**User response:** Verify your GPFS daemon version.

---

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

**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 use by the IBM Support Center.

**User response:** Verify your GPFS daemon version.

---

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

**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 use by the IBM Support Center.

**User response:** Verify your GPFS daemon version.

---

**6027-344 [E] Node *nodeName* is incompatible because its version 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 use by the IBM Support Center.

**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 [E] 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 [E] 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 [E] 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 [E] 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 [E] 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.**

**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 [E] 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 `mmaddisk` 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 `mmaddisk` command.

---

**6027-364**      **Error processing disks.**

**Explanation:** An error occurred when the `mmaddisk` command was reading disks in the file system.

**User response:** Determine the reason why the disks are inaccessible for reading, then reissue the `mmaddisk` command.

---

**6027-365 [I]**    **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:** `mmaddisk` 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 `mmaddisk` 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]** I/O still in progress near sector number on disk *diskName*.

**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.

---

**6027-382** 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 of *blockSize* 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.

---

**6027-426** 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.

---



---

**6027-435 [N] The file system descriptor quorum has been overridden.**

**Explanation:** The `mmfsctl exclude` command was previously issued to override the file system descriptor quorum after a disaster.

**User response:** None. Informational message only.

---

**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*' on node *nodeName*.**

**Explanation:** A disk name that is not valid was specified in a GPFS disk command.

**User response:** Correct the parameters of the executing GPFS disk command.

---

**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/sgname` 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/sgname` 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 file system *filesystem*. Created on *IPandTime*.**

**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 sure that no other file system has this disk configured because you may cause data corruption in both file systems if this is not the case.

---

**6027-471**     **Disk *diskName*: 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 reconfigured 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 operative.

---

**6027-472 [E]**   **File system format version *versionString* is not supported.**

**Explanation:** The current file system format version is not supported.

**User response:** Verify:

1. The disks are correctly defined on all nodes.
2. The paths to the disks are correctly defined and operative.

---

**6027-473 [X]**   **File System *filesystem* unmounted by the system with return code *value* 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 [X]**   **Recovery Log I/O failed, unmounting file system *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 '*--inode-limit*' is not enabled. Use option '*-V*' to enable most recent features.**

**Explanation:** *mmchfs --inode-limit* is not enabled under the current file system format version.

**User response:** Run *mmchfs -V*, this will change the file system format to the latest format supported.

---

**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*.

---

---

**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 [N]**    **Mount of *fsName* was blocked by *fileName***

**Explanation:** The internal or external mount of the file system was blocked by the existence of the specified file.

**User response:** If the file system needs to be mounted, remove the specified file.

---

**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 [E]**    **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 [E]**    **Remount failed for device *name*: *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 [N]**    **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 `mmchdisk` 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. This is allowed, but the files will not be replicated and will therefore be at risk.**

**Explanation:** You 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 [N]**    **The descriptor replica on disk *diskName* 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 the following *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 not 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.

---

**6027-499 [X]**    **An unexpected Device Mapper path *dmDevice (nsdId)* has been detected. The new path does not have a Persistent Reserve set up. File system *fileSystem* will be internally unmounted.**

**Explanation:** A new device mapper path is detected or a previously failed path is activated after the local

device discovery has finished. This path lacks a Persistent Reserve, and can not be used. All device paths must be active at mount time.

**User response:** Check the paths to all disks making up the file system. Repair any paths to disks which have failed. Remount the file system.

---

**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:** `mmfsmnhelp` 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 [E]**    **The quota record *recordNumber* 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 [W]**    **Inode space *inodeSpace* in file system *fileSystem* is approaching the limit for the maximum number of inodes.**

**Explanation:** The number of files created is approaching the file system limit.

**User response:** Use the **mmchfileset** command to increase the maximum number of files 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 [N]   Fileset in file system  
*fileSystem: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 [X]   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:** Progress information.

**User response:** None. Informational message only.

---

**6027-566**     **Error processing user file metadata.**

**Explanation:** Error encountered while processing user file metadata.

**User response:** None. Informational message only.

---

**6027-567**     **Waiting for pending file system scan to finish ...**

**Explanation:** Progress information.

**User response:** None. Informational message only.

---



---

**6027-568**     **Waiting for *number* pending file system scans to finish ...**

**Explanation:** Progress information.

**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 the 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*. Failed with error *errorCode***

**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

**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.

---

**6027-584**      **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:** Progress information.

**User response:** None. Informational message only.

---

**6027-590 [W]**      **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 size of 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 [E]**      **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 size MB.

**Explanation:** The specified disk is too large compared to the disks that were initially used to create the storage pool.

**User response:** Specify a smaller disk or add the disk to a new storage pool.

---

**6027-595 [E]** While creating quota files, file *fileName*, with no valid quota information was found in the root directory. Remove files with reserved quota file names (for example, `user.quota`) without valid quota information from the root directory by: - mounting the file system without quotas, - removing the files, and - 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 [I]** While creating quota files, file *fileName* containing quota information was found in the root directory. This file will be used as *quotaType* 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 [E]** The quota command was requested to process quotas for a type (`user`, `group`, or `fileset`), which is not enabled.

**Explanation:** A quota command was 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 [E]** 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 [E]** 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: *valueK* = *valueM*, max block size: *valueK* = *valueM*.**

**Explanation:** Displays the current pool size.

**User response:** None. Informational message only.

---

**6027-604 [E]**   **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 [N]**   **File system has been renamed.**

**Explanation:** Self-explanatory.

**User response:** None. Informational message only.

---

**6027-606 [E]**   **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 [E]**   **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 [I]**   **Recovery: *fileSystem*, 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 [N]**   **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 [X]** 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 [I]** 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 [X]** 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 [E]** Failed to join remote cluster *clusterName*

**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** File system manager takeover 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 [N]**   **Node *nodeName* resigned as manager for *fileSystem*.**

**Explanation:** Progress report produced by the **mmchmgr** command.

**User response:** None. Informational message only.

---

**6027-630 [N]**   **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 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**      **The best choice node *nodeName* is already the 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 [E]**   **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 [E]**   **Disk 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 [E]**   **RVSD is not active.**

**Explanation:** The RVSD subsystem needs to be activated.

**User response:** See the appropriate IBM Reliable Scalable Cluster Technology (RSCT) document ([www.ibm.com/support/knowledgecenter/SGVKBA/welcome](http://www.ibm.com/support/knowledgecenter/SGVKBA/welcome)) and search on *diagnosing IBM Virtual Shared Disk problems*.

---

**6027-638 [E]**   **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 [E]**   **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 [E]**   **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 [E] 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 [N] File system has been deleted.**

**Explanation:** Self-explanatory.

**User response:** None. Informational message only.

---

**6027-643 [I] Node *nodeName* completed take over 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 95. A common cause for such errors is lack of space in `/var`.

---

**6027-646 [E] 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 [E] File *fileName* could not be run with error *errno*.**

**Explanation:** The specified shell script could not be run. This message is followed by the error string that is returned by the `exec`.

**User response:** Check file existence and access permissions.

---

**6027-648 EDITOR environment variable must be full 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 [X] 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.

---

**6027-680 [E]**    **Disk failure. Volume *name*. rc = *value*.  
Physical volume *name*.**

**Explanation:** An I/O request to a disk or a request to fence a disk has failed in such a manner that GPFS can no longer use the disk.

**User response:** Check the disk hardware and the software subsystems in the path to the disk.

---

**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 [E]**    **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 [E]**    **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 [E]**    **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 [E]**    **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 [E]**    **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 [E]**    **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 [X]**    **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 [X]**    **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 [X] 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 [E] 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 [E] 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 [E]** 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 [E]** 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 [E]** 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 [E]** 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 [E]** 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 [E]** 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 [E]**    **Quota initialization/start-up 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 [W]**    **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 [E] 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 [E]** 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 [E]** Error code *value*

**Explanation:** Provides additional information about an error.

**User response:** See accompanying error messages.

---

**6027-752 [E]** 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 [E]** 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 [I]** Waiting for challenge *challengeValue* (node *nodeNumber*, sequence *sequenceNumber*) 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 [E]** 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.

---

**6027-760 [E]** Unable to perform i/o to the disk. This node is either fenced from accessing the disk or this node's disk lease has expired.

**Explanation:** A read or write to the disk failed due to either being fenced from the disk or no longer having a disk lease.

**User response:** Verify disk hardware fencing setup is correct if being used. Ensure network connectivity between this node and other nodes is operational.

---

**6027-761 [W]** Attention: excessive timer drift between *node* and *node* (*number over number sec*).

**Explanation:** GPFS has detected an unusually large difference in the rate of clock ticks (as returned by the **times()** system call) between two nodes. Another node's TOD clock and tick rate changed dramatically relative to this node's TOD clock and tick rate.

**User response:** Check error log for hardware or device driver problems that might cause timer interrupts to be lost or a recent large adjustment made to the TOD clock.

---

---

**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 [W]** 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 [N]** 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 [N]** 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 [I] 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 [I] Recovering nodes in cluster** *cluster:*  
*nodeList*

**Explanation:** Recovery for one or more nodes in the cited cluster has begun.

**User response:** No response is needed if this message is followed by 'recovered nodes' entries on the cited cluster 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 [E] Device not ready.**

**Explanation:** A device is not ready for operation.

**User response:** Check previous messages for further information.

---

**6027-785 [E] Cannot establish connection.**

**Explanation:** This node cannot establish a connection to another node.

**User response:** Check previous messages for further information.

---

**6027-786 [E] 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 [E] 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 [E] 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.

---

**6027-861 Set afm ctl failed**

**Explanation:** The `tsfattr` call failed.

**User response:** Check for additional error messages. Resolve the problems before reattempting the failing operation.

---

**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 [E] 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 [E] 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 [E] 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 [E] *pathName* is not within a mounted GPFS file system.**

**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). Mount the GPFS file system. Correct the

command invocation and reissue the command.

---

**6027-873 [W] 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.

---

**6027-874 [E] 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 [E] 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 [E] 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 [E] 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 [E] 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 [E] 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(*inodeNumber*): *errorString***

**Explanation:** An error occurred during a `gpfs_seek_inode` operation.

**User response:** Reissue the command. If the problem persists, contact the contact the IBM Support Center

---

**6027-881 [E] 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 [E] Error on gpfs\_ireaddir(*rootPath*/*pathName*): *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\_next\_inode(*maxInodeNumber*): *errorString***

**Explanation:** An error occurred during a `gpfs_next_inode` operation.

**User response:** Reissue the command. If the problem persists, contact the IBM Support Center.

---

**6027-884 [E:nnn] Error during directory scan**

**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 [E:nnn] Error during inode scan: *errorString***

**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 [E:nnn] Error during policy decisions scan**

**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 [W] Error on gpfs\_igetstoragepool(*dataPoolId*):**  
*errorString*

**Explanation:** An error occurred during a `gpfs_igetstoragepool` operation. Possible inode corruption.

**User response:** Use `mmfsck` command. If the problem persists, contact the IBM Support Center.

---

**6027-888 [W] Error on gpfs\_igetfilesetName(*filesetId*):**  
*errorString*

**Explanation:** An error occurred during a `gpfs_igetfilesetName` operation. Possible inode corruption.

**User response:** Use `mmfsck` command. If the problem persists, contact the IBM Support Center.

---

**6027-889 [E] 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 [E] 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 [X] 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 [E] 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 [X] 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 [X] 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 [X] 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 [X] 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 [X] 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 [X] 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 [X] 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 [E] 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 [E] 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 [E] Error parsing work file** *fileName*. **Service index:** *number*

**Explanation:** An error occurred while attempting to read the specified work file.

**User response:** Investigate the file and possibly reissue the command. Make sure that there is enough free space in the file system. If the error persists, contact the IBM Support Center.

---

**6027-903 [E:nnn] Error while loading policy rules.**

**Explanation:** An error occurred while attempting to read or parse the policy file, which may contain syntax errors. Subsequent messages include more information about the error.

**User response:** Read all of the related error messages and try to correct the problem.

---

**6027-904 [E] 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 [E] Error: Out of memory. Service index:**  
*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 [E:nnn] Error on system(command)**

**Explanation:** An error occurred during the system call with the specified argument string.

**User response:** Read and investigate related error messages.

---

**6027-907 [E:nnn] Error 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 [W] 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 [E] 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 [E:nnn] Error during policy execution**

**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 [E] Error on changeSpecification change for**  
**pathName. errorString**

**Explanation:** This message provides more details about a **gpfs\_fcntl()** error.

**User response:** Use the **mmlsattr** and **mmchattr** commands to examine the file, and then reissue the change command.

---

**6027-912 [E] Error on restriping of *pathName*.  
*errorString***

**Explanation:** This provides more details on a `gdfs_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 [E] 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 [E] 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 [X] 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 [W] 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

## 6027-924 [X] • 6027-934 [W]

command software is installed on all nodes in the clusters and then reissue the command.

---

**6027-924 [X] 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 [E] 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 `mmlsnapshot` command to list the snapshot names for the file system.

---

**6027-929 [W] Attention: In RULE '*ruleName*' (*ruleNumber*), both pools '*poolName*' and '*poolName*' 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 [W] 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 [E] 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 [E] 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 [W] 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 [W] 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 [W] 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 [W] 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 [E] Error creating shared temporary sub-directory *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 [E] 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.

---

**6027-939 [E] Error on `gpfs_quotactl(pathName,commandCode,resourceId): errorString`**

**Explanation:** An error occurred while attempting `gpfs_quotactl()`.

**User response:** Correct the policy rules and/or enable GPFS quota tracking. If problem persists 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.

---

**6027-946 Not a regular file or directory.**

**Explanation:** An `mmlsatr` or `mmchattr` command error occurred.

**User response:** Correct the problem and reissue the command.

---

**6027-947 Stat failed: A file or directory in the path name does not exist.**

**Explanation:** A file or directory in the path name does not exist.

**User response:** Correct the problem and reissue the command.

---

**6027-948 [E:nnn] *fileName: get clone attributes failed: errorString***

**Explanation:** The `tsfattr` call failed.

**User response:** Check for additional error messages.

Resolve the problems before reattempting the failing operation.

---

**6027-949 [E]** *fileName*: **invalid clone attributes.**

**Explanation:** Self explanatory.

**User response:** Check for additional error messages. Resolve the problems before reattempting the failing operation.

---

**6027-950 [E:nnn]** File cloning requires the 'fastea' feature to be enabled.

**Explanation:** The file system **fastea** feature is not enabled.

**User response:** Enable the **fastea** feature by issuing the **mmchfs -V** and **mmmigratefs --fastea** commands.

---

**6027-951 [E]** Error on *operationName* to work file  
*fileName*: *errorString*

**Explanation:** An error occurred while attempting to do a (write-like) operation on 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.

---

**6027-953** Failed to get a handle for fileset  
*filesetName*, *snapshot snapshotName* in file  
system *fileSystem*. *errorMessage*.

**Explanation:** Failed to get a handle for a specific fileset snapshot in the file system.

**User response:** Correct the command line and reissue the command. If the problem persists, contact the IBM Support Center.

---

**6027-954** Failed to get the maximum inode  
number in the active file system.  
*errorMessage*.

**Explanation:** Failed to get the maximum inode number in the current active file system.

**User response:** Correct the command line and reissue the command. If the problem persists, contact the IBM Support Center.

---

**6027-955** Failed to set the maximum allowed  
memory for the specified *fileSystem*  
command.

**Explanation:** Failed to set the maximum allowed memory for the specified command.

**User response:** Correct the command line and reissue the command. If the problem persists, contact the IBM Support Center.

---

**6027-956** Cannot allocate enough buffer to record  
different items.

**Explanation:** Cannot allocate enough buffer to record different items which are used in the next phase.

**User response:** Correct the command line and reissue the command. If the problem persists, contact the system administrator.

---

**6027-957** Failed to get the root directory inode of  
fileset *filesetName*

**Explanation:** Failed to get the root directory inode of a fileset.

**User response:** Correct the command line and reissue the command. If the problem persists, contact the IBM Support Center.

---

**6027-959** '*fileName*' is not a regular file.

**Explanation:** Only regular files are allowed to be clone parents.

**User response:** This file is not a valid target for **mmclone** operations.

---

**6027-960** cannot access '*fileName*': *errorString*.

**Explanation:** This message provides more details about a **stat()** error.

**User response:** Correct the problem and reissue the command.

---

**6027-961** Cannot execute *command*.

**Explanation:** The **mmeditac** command cannot invoke the **mmgetacl** or **mmputacl** command.

**User response:** Contact your system administrator.

---

**6027-962** Failed to list fileset *filesetName*.

**Explanation:** Failed to list specific fileset.

**User response:** None.

---

**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 `mmpuACL` 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 [E:nnn]**      **Error 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 [W]**      **Attention: **RULE 'ruleName'** attempts to redefine **EXTERNAL POOL** or **LIST** literal '*poolName*', ignored.**

**Explanation:** Execution continues as if the specified rule was not present.

**User response:** Correct or remove the policy rule.

---

---

**6027-998 [E] Error in FLR/PDR serving for client**  
*clientHostNameAndPortNumber:*  
**FLRs=***numOfFileListRecords*  
**PDRs=***numOfPolicyDecisionResponses*  
**pdrs=***numOfPolicyDecisionResponseRecords*

**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 [E] 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 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. The **mmcrfsc** command should never call the **mmcrfsc** command without a valid file system size argument. Contact the IBM Support Center.

---

**6027-1028** Incorrect value for *-name* flag.

**Explanation:** An incorrect argument was specified with an option that requires one of a limited number of allowable options (for example, *-s* or any of the **yes** | **no** options).

**User response:** Use one of the valid values for the specified option.

---

**6027-1029** Incorrect characters in integer field for *-name* 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 *-optionLetter* option. Valid range is from *value* to *value*

**Explanation:** The value specified with an option was below the minimum.

**User response:** Use an integer in the valid range for the indicated option.

---

**6027-1031** Value above maximum for option *-optionLetter*. Valid range is from *value* to *value*.

**Explanation:** The value specified with an option was above the maximum.

**User response:** Use an integer in the valid 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  $\leq$  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.

---

**6027-1084**    **The specified block size (*valueK*) exceeds the maximum allowed block size currently in effect (*valueK*). Either specify a smaller value for the -B parameter, or increase the maximum block size by issuing: **mmchconfig maxblocksize=*valueK*** and restart 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.

---

**6027-1119**    **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 parameter has a value that is not valid.

**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 subset of the parameters were changed.**

**Explanation:** An interrupt was received in **mmchfs** before all of the requested changes could be completed.

**User response:** Use **mmfsfs** to see what the currently active settings are. Reissue the command if you want to change additional parameters.

---

**6027-1135**    **Restripping may not have finished.**

**Explanation:** An interrupt occurred during restripping.

**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.

---

**6027-1152**    **Block size must be 64K, 128K, 256K, 512K, 1M, 2M, 4M, 8M or 16M.**

**Explanation:** The specified block size value is not valid.

**User response:** Specify a valid block size value.

---

**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. Check 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. Check 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. Check 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.

---

**6027-1168** Inode size must be 512, 1K or 4K.

**Explanation:** The specified inode size is not valid.

**User response:** Specify a valid inode size.

---

**6027-1169** *attribute* must be *value*.

**Explanation:** The specified value of the given attribute is not valid.

**User response:** Specify a valid value.

---

**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:** The internal **tscrfs** command failed.

**User response:** Check the error message from the command that failed.

---

**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:** Check 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.

---

**6027-1218** Node *nodeName* does not have a GPFS server license designation.

**Explanation:** The function that you are assigning to the node requires the node to have a GPFS server license.

**User response:** Use the `mmchlicense` command to assign a valid GPFS license to the node or specify a different node.

---

**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 76.

---

**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 **mmremoteccluster** 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 **mmremoteccluster** 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.

---

**6027-1254** Warning: Not all nodes have proper GPFS license designations. Use the `mmchlicense` command to designate licenses as needed.

**Explanation:** Not all nodes in the cluster have valid license designations.

**User response:** Use `mmlslicense` to see the current license designations. Use `mmchlicense` to assign valid GPFS licenses to all nodes as needed.

---

**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** File *fileName* does not contain any *typeOfStanza* stanzas.

**Explanation:** The input file should contain at least one specified stanza.

**User response:** Correct the input file and reissue the command.

---

**6027-1261** *descriptorField* must be specified in *descriptorType* descriptor.

**Explanation:** A required field of the descriptor was empty. The incorrect descriptor is displayed following this message.

**User response:** Correct the input and reissue the command.

---

**6027-1262** Unable to obtain the GPFS configuration file lock. Retrying ...

**Explanation:** A command requires the lock for the GPFS system data but was not able to obtain it.

**User response:** None. Informational message only.

---

**6027-1263** Unable to obtain the GPFS configuration file lock.

**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 "GPFS cluster configuration data files are locked" on page 76, and then reissue the command.

---

**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 89.

---

**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**    **Unexpected error obtaining the *lockName* lock.**

**Explanation:** GPFS cannot obtain the specified lock.

**User response:** Examine any previous error messages. Correct any problems and reissue the command. If the problem persists, perform problem determination and contact the IBM Support Center.

---

**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 `mmremotecoluster` 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 `mmremotecoluster` 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 `mmremotecoluster` 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** The **-N** option 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 value *value* specified for failure group.

**Explanation:** The specified failure group is not valid.

**User response:** Correct the problem and reissue the command.

**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** This function is not available in the GPFS Express Edition.

**Explanation:** The requested function is not part of the GPFS Express Edition.

**User response:** Install the GPFS Standard Edition on all nodes in the cluster, and then reissue the command.

---

**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** Prerequisite libraries not found or correct version not installed. Ensure *productName* is properly installed.

**Explanation:** The specified software product is missing or is not properly installed.

**User response:** Verify that the product is installed properly.

---

**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** *feature* is not available in the GPFS Express Edition.

**Explanation:** The specified function or feature is not part of the GPFS Express Edition.

**User response:** Install the GPFS Standard Edition on all nodes in the cluster, and then reissue the command.

---

**6027-1309** Storage pools are not available in the GPFS Express Edition.

**Explanation:** Support for multiple storage pools is not part of the GPFS Express Edition.

**User response:** Install the GPFS Standard Edition on all nodes in the cluster, and then reissue the command.

---

**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.

---

**6027-1338** 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.

---

**6027-1352** *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.

---

**6027-1357** 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 **mmaddisk** 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.

---

**6027-1370** 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.

---

**6027-1377**     **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 77.

---

**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.

---

**6027-1450**    **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 [E]** **Open *devicetype 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 instance 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 then take needed repair action on the specified disk.

---

**6027-1501 [X]** **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 [X]** **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**  
*fileSystem.*

**Explanation:** The **mmadddisk** command successfully completed.

**User response:** None. Informational message only.

---

**6027-1504**    **File *name* could not be run with err error.**

**Explanation:** A failure occurred while trying to run an external program.

**User response:** Make sure the file exists. If it does, check its 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=** 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=** 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=** 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 [E]** **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 [E]** **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 [E]** **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.

---

**6027-1513**    *DiskName* **is not an sg device, or sg**  
**driver is older than sg3**

**Explanation:** The disk is not a SCSI disk, or supports SCSI standard older than SCSI 3.

**User response:** Correct the command invocation and try again.

---

**6027-1514**    **ioctl failed with rc=returnCode. Related**  
**values are SCSI status=scsiStatusValue,**  
**host\_status=hostStatusValue,**  
**driver\_status=driverStatsValue.**

**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.

---

**6027-1515**    **READ KEY ioctl failed with**  
**rc=returnCode. Related values are SCSI**  
**status=scsiStatusValue,**  
**host\_status=hostStatusValue,**  
**driver\_status=driverStatsValue.**

**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.

---

**6027-1516**    **REGISTER ioctl failed with**  
**rc=returnCode. Related values are SCSI**  
**status=scsiStatusValue,**  
**host\_status=hostStatusValue,**  
**driver\_status=driverStatsValue.**

**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.

---

**6027-1517**    **READ RESERVE ioctl failed with**  
**rc=returnCode. Related values are SCSI**  
**status=scsiStatusValue,**  
**host\_status=hostStatusValue,**  
**driver\_status=driverStatsValue.**

**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.

---

**6027-1518**    **RESERVE ioctl failed with rc=returnCode.**  
**Related values are SCSI**  
**status=scsiStatusValue,**  
**host\_status=hostStatusValue,**  
**driver\_status=driverStatsValue.**

**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.

---

**6027-1519**    **INQUIRY ioctl failed with rc=returnCode.**  
**Related values are SCSI**  
**status=scsiStatusValue,**  
**host\_status=hostStatusValue,**  
**driver\_status=driverStatsValue.**

**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.

---

**6027-1520**    **PREEMPT ABORT ioctl failed with rc=returnCode. Related values are SCSI status=scsiStatusValue, host\_status=hostStatusValue, driver\_status=driverStatsValue.**

**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.

---

**6027-1521**    **Can not find register key registerKeyValue at device diskName.**

**Explanation:** Unable to find given register key at the disk.

**User response:** Correct the problem and reissue the command.

---

**6027-1522**    **CLEAR ioctl failed with rc=returnCode. Related values are SCSI status=scsiStatusValue, host\_status=hostStatusValue, driver\_status=driverStatsValue.**

**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.

---

**6027-1523**    **Disk name longer than value is not allowed.**

**Explanation:** The specified disk name is too long.

**User response:** Reissue the command with a valid disk name.

---

**6027-1524**    **The READ\_KEYS ioctl data does not contain the key that was passed as input.**

**Explanation:** A `REGISTER ioctl` call apparently succeeded, but when the device was queried for the key, the key was not found.

**User response:** Check the device subsystem and try to correct the problem. If the problem persists, contact the IBM Support Center.

---

**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 [E] 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 [I] Connect in progress to *ipAddress***

**Explanation:** Connecting sockets between nodes.

**User response:** None. Information message only.

---

**6027-1539 [E] 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 [A] Try and buy license has expired!**

**Explanation:** Self explanatory.

**User response:** Purchase a GPFS license to continue using GPFS.

---

**6027-1541 [N] 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 [A] 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 [W] Sum of **prefetchthreads**(*value*), **worker1threads**(*value*) and **nsdMaxWorkerThreads** (*value*) exceeds *value*. Reducing them to *value*, *value* and *value*.**

**Explanation:** The sum of **prefetchthreads**, **worker1threads**, and **nsdMaxWorkerThreads** exceeds the permitted value.

**User response:** Accept the calculated values or reduce the individual settings using **mmchconfig prefetchthreads=newvalue** or **mmchconfig worker1threads=newvalue**. or **mmchconfig nsdMaxWorkerThreads=newvalue**. After using **mmchconfig**, the new settings will not take affect until the GPFS daemon is restarted.

---

**6027-1545 [A] 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 [W] 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 [A] 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 [A] 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 Linux” on page 79 for details.

---

**6027-1549 [A] 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 [W] 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.

---



---

**6027-1559 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 `mmaddisk` or `mmddisk` 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.

---

**6027-1570** 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 `mmcluster` 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.

---

**6027-1587** 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 **mmcluster** 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 **mmclscluster** 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 87.

---

**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 could not 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 found while checking node descriptor *descriptor***

**Explanation:** A node 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.

---



---

**6027-1630**    **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. If no other messages are shown, examine the GPFS log files in the **/var/adm/ras** directory on each node.

---

**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 device *deviceName* refers to an existing NSD *name***

**Explanation:** The specified disk device 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 [I]** 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 *IBM Spectrum Scale: 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.

---

**6027-1699**    **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 [I]**    **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 [N]**    **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 [I]**    **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.

## 6027-1715 • 6027-1731 [E]

**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 [N]** **Close connection to *ipAddress***

**Explanation:** Connection socket closed.

**User response:** None. Informational message only.

---

**6027-1717 [E]** **Error initializing the configuration server, err *value***

**Explanation:** The configuration server module could not be initialized due to lack of system resources.

**User response:** Check system memory.

---

**6027-1718 [E]** **Could not run command *name*, err *value***

**Explanation:** The GPFS daemon failed to run the specified command.

**User response:** Verify correct installation.

---

**6027-1724 [E]** **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 [E]** **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 [E]** **The administrator of the cluster named *clusterName* requires authentication. Contact the administrator to obtain the clusters 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 [E]** **The administrator of the cluster named *clusterName* does not require authentication. Unregister the clusters 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 [E]** **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 [E]** **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 [E]** **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 [E]** **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 [X] Remote mounts are not enabled within this cluster.**

**Explanation:** Remote mounts cannot be performed in this cluster.

**User response:** See the *IBM Spectrum Scale: 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 [E] 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 [E] 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 [N] 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 [N] Close connection to *ipAddress* (*errorString*).**

**Explanation:** Connection socket closed.

**User response:** None. Informational message only.

---

**6027-1738 [E] Close connection to *ipAddress* (*errorString*). Attempting reconnect.**

**Explanation:** Connection socket closed.

**User response:** None. Informational message only.

---

**6027-1739 [X] Accept socket connection failed: *err value*.**

**Explanation:** The Accept socket connection received an unexpected error.

**User response:** None. Informational message only.

---

**6027-1740 [E] 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 [E] 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 [E] 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.

---

**6027-1743 [W] Failed to load GSKit library *path*: (*dlerror*) *errorMessage***

**Explanation:** The GPFS daemon could not load the library required to secure the node-to-node communications.

**User response:** Verify that the **gpfs.gskit** package was properly installed.

---

**6027-1744 [I] GSKit library loaded and initialized.**

**Explanation:** The GPFS daemon successfully loaded the library required to secure the node-to-node communications.

**User response:** None. Informational message only.

---

---

**6027-1745 [E] Unable to resolve symbol for routine:**  
*functionName (dlerror) errorMessage*

**Explanation:** An error occurred while resolving a symbol required for transport-level security.

**User response:** Verify that the gpfs.gskit package was properly installed.

---

**6027-1746 [E] Failed to load or initialize GSKit library: error value**

**Explanation:** An error occurred during the initialization of the transport-security code.

**User response:** Verify that the gpfs.gskit package was properly installed.

---

**6027-1747 [W] The TLS handshake with node *ipAddress* failed with error *value* (*handshakeType*).**

**Explanation:** An error occurred while trying to establish a secure connection with another GPFS node.

**User response:** Examine the error messages to obtain information about the error. Under normal circumstances, the retry logic will ensure that the connection is re-established. If this error persists, record the error code and contact the IBM Support Center.

---

**6027-1748 [W] A secure receive from node *ipAddress* failed with error *value*.**

**Explanation:** An error occurred while receiving an encrypted message from another GPFS node.

**User response:** Examine the error messages to obtain information about the error. Under normal circumstances, the retry logic will ensure that the connection is re-established and the message is received. If this error persists, record the error code and contact the IBM Support Center.

---

**6027-1749 [W] A secure send to node *ipAddress* failed with error *value*.**

**Explanation:** An error occurred while sending an encrypted message to another GPFS node.

**User response:** Examine the error messages to obtain information about the error. Under normal circumstances, the retry logic will ensure that the connection is re-established and the message is sent. If this error persists, record the error code and contact the IBM Support Center.

---

**6027-1750 [N] The *handshakeType* TLS handshake with node *ipAddress* was cancelled: connection reset by peer (return code *value*).**

**Explanation:** A secure connection could not be

established because the remote GPFS node closed the connection.

**User response:** None. Informational message only.

---

**6027-1751 [N] A secure send to node *ipAddress* was cancelled: connection reset by peer (return code *value*).**

**Explanation:** Securely sending a message failed because the remote GPFS node closed the connection.

**User response:** None. Informational message only.

---

**6027-1752 [N] A secure receive to node *ipAddress* was cancelled: connection reset by peer (return code *value*).**

**Explanation:** Securely receiving a message failed because the remote GPFS node closed the connection.

**User response:** None. Informational message only.

---

**6027-1753 [E] The crypto library with FIPS support is not available for this architecture. Disable FIPS mode and reattempt the operation.**

**Explanation:** GPFS is operating in FIPS mode, but the initialization of the cryptographic library failed because FIPS mode is not yet supported on this architecture.

**User response:** Disable FIPS mode and attempt the operation again.

---

**6027-1754 [E] Failed to initialize the crypto library in FIPS mode. Ensure that the crypto library package was correctly installed.**

**Explanation:** GPFS is operating in FIPS mode, but the initialization of the cryptographic library failed.

**User response:** Ensure that the packages required for encryption are properly installed on each node in the cluster.

---

**6027-1803 [E] 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 [E] 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 [N]** Rediscovered nsd server access to *name*.

**Explanation:** A server rediscovered access to the specified disk.

**User response:** None.

---

**6027-1806 [X]** A Persistent Reserve could not be established on device name (*deviceName*): *errorLine*.

**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 [E]** 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.

---

**6027-1808 [A]** Unable to reserve space for NSD buffers. Increase pagepool size to at least *requiredPagePoolSize* MB. Refer to the GPFS: Administration and Programming Reference for more information on selecting an appropriate pagepool size.

**Explanation:** The `pagepool` usage for an NSD buffer ( $4 * \text{maxblocksize}$ ) is limited by factor `nsdBufSpace`. The value of `nsdBufSpace` can be in the range of 10-70. The default value is 30.

**User response:** Use the `mmchconfig` command to decrease the value of `maxblocksize` or to increase the value of `pagepool` or `nsdBufSpace`.

---

**6027-1809 [E]** The defined server *serverName* for NSD *NsdName* couldn't be resolved.

**Explanation:** The host name of the NSD server could not be resolved by `gethostbyname()`.

**User response:** Fix the host name resolution.

---

**6027-1810 [I]** Vdisk server recovery: delay *number sec.* for safe recovery.

**Explanation:** Wait for the existing disk lease to expire before performing vdisk server recovery.

**User response:** None.

---

**6027-1811 [I]** Vdisk server recovery: delay complete.

**Explanation:** Done waiting for existing disk lease to expire before performing vdisk server recovery.

**User response:** None.

---

**6027-1812 [E]** Rediscovery failed for *name*.

**Explanation:** A server failed to rediscover access to the specified disk.

**User response:** Check the disk access issues and run the command again.

---

**6027-1813 [A]** Error reading volume identifier (for *objectName name*) from configuration file.

**Explanation:** The volume identifier for the named recovery group or vdisk could not be read from the `mmsdrfs` file. This should never occur.

**User response:** Check for damage to the `mmsdrfs` file.

---

**6027-1814 [E]** Vdisk *vdiskName* cannot be associated with its recovery group *recoveryGroupName*. This vdisk will be ignored.

**Explanation:** The named vdisk cannot be associated with its recovery group.

**User response:** Check for damage to the `mmsdrfs` file.

---

**6027-1815 [A]** Error reading volume identifier (for NSD *name*) from configuration file.

**Explanation:** The volume identifier for the named NSD could not be read from the `mmsdrfs` file. This should never occur.

**User response:** Check for damage to the `mmsdrfs` file.

---

**6027-1816 [E]** The defined server *serverName* for recovery group *recoveryGroupName* could not be resolved.

**Explanation:** The hostname of the NSD server could not be resolved by `gethostbyname()`.

**User response:** Fix hostname resolution.

---

**6027-1817 [E]** Vdisks are defined, but no recovery groups are defined.

**Explanation:** There are vdisks defined in the `mmsdrfs` file, but no recovery groups are defined. This should never occur.

**User response:** Check for damage to the `mmsdrfs` file.

---

---

**6027-1818 [I] Relinquished recovery group**  
*recoveryGroupName (err errorCode).*

**Explanation:** This node has relinquished serving the named recovery group.

**User response:** None.

---

**6027-1819 Disk descriptor for name refers to an existing pdisk.**

**Explanation:** The **mmcrrecoverygroup** command or **mmaddpdisk** command found an existing pdisk.

**User response:** Correct the input file, or use the **-v** option.

---

**6027-1820 Disk descriptor for name refers to an existing NSD.**

**Explanation:** The **mmcrrecoverygroup** command or **mmaddpdisk** command found an existing NSD.

**User response:** Correct the input file, or use the **-v** option.

---

**6027-1821 Error errno writing disk descriptor on name.**

**Explanation:** The **mmcrrecoverygroup** command or **mmaddpdisk** command got an error writing the disk descriptor.

**User response:** Perform disk diagnostics.

---

**6027-1822 Error errno reading disk descriptor on name.**

**Explanation:** The **tspreparedpdisk** command got an error reading the disk descriptor.

**User response:** Perform disk diagnostics.

---

**6027-1823 Path error, name and name are the same disk.**

**Explanation:** The **tspreparedpdisk** command got an error during path verification. The pdisk descriptor file is miscoded.

**User response:** Correct the pdisk descriptor file and reissue the command.

---

**6027-1824 [X] An unexpected Device Mapper path dmDevice (nsdId) has been detected. The new path does not have a Persistent Reserve set up. Server disk diskName will be put offline**

**Explanation:** A new device mapper path is detected or a previously failed path is activated after the local device discovery has finished. This path lacks a Persistent Reserve, and cannot be used. All device

paths must be active at mount time.

**User response:** Check the paths to all disks making up the file system. Repair any paths to disks which have failed. Rediscover the paths for the NSD.

---

**6027-1825 [A] Unrecoverable NSD checksum error on I/O to NSD disk nsdName, using server serverName. Exceeds retry limit number.**

**Explanation:** The allowed number of retries was exceeded when encountering an NSD checksum error on I/O to the indicated disk, using the indicated server.

**User response:** There may be network issues that require investigation.

---

**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 Device with major/minor numbers 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.

---

**6027-1906** *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.

---

**6027-1908** *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.

---

**6027-1922** *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:** Use `mmauth show ciphers` to display 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.

---

**6027-1945** 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** Use `mmanth genkey` to recover the file `fileName`, or to generate and commit a new key.

**Explanation:** The specified file was not found.

**User response:** Recover the file, or generate a new key by running: `mmanth genkey propagate` or generate a new key by running `mmanth genkey new`, followed by the `mmanth 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** Propagating the cluster configuration data to all affected nodes.

**Explanation:** The cluster configuration data is being sent to the rest of the nodes in the cluster.

**User response:** This is an informational message.

---

**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.

---

**6027-1967** 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.

---

**6027-1974** 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.

---

**6027-1993**    **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 *physicalDiskName*.

**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 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 non-zero return code**

**Explanation:** An ioctl failed with a non-zero return code.

**User response:** Correct the problem, if possible. Otherwise, contact the IBM Support Center.

---

**6027-2049 [X]** **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 [E]** **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**    **Error found while processing stanza**

**Explanation:** A stanza was found to be unsatisfactory in some way.

**User response:** Check the preceding messages, if any, and correct the condition that caused the stanza to be rejected.

---

**6027-2109**    **Failed while processing disk stanza on node** *nodeName.*

**Explanation:** A disk stanza was found to be unsatisfactory in some way.

**User response:** Check the preceding messages, if any, and correct the condition that caused the stanza to be rejected.

---

**6027-2110**    **Missing required parameter** *parameter*

**Explanation:** The specified parameter is required for this command.

**User response:** Specify the missing information and reissue the command.

---

**6027-2111**    **The following disks were not deleted:**  
*diskList*

**Explanation:** The command could not delete the specified disks. Check the preceding messages for error information.

**User response:** Correct the problems and reissue the command.

---

**6027-2112**    **Permission failure. Option** *option*  
**requires root authority to run.**

**Explanation:** The specified command option requires root authority.

**User response:** Log on as root and reissue the 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.

---

**6027-2116 [E]** **The GPFS daemon must be active on the recovery group server nodes.**

**Explanation:** The command requires that the GPFS daemon be active on the recovery group server nodes.

**User response:** Ensure GPFS is running on the recovery group server nodes and reissue the command.

---

**6027-2117 [E]** *object name* **already exists.**

**Explanation:** The user attempted to create an object with a name that already exists.

**User response:** Correct the name and reissue the command.

---

**6027-2118 [E]** **The** *parameter* **is invalid or missing in the pdisk descriptor.**

**Explanation:** The pdisk descriptor is not valid. The bad descriptor is displayed following this message.

**User response:** Correct the input and reissue the command.

---

**6027-2119 [E]** **Recovery group** *name* **not found.**

**Explanation:** The specified recovery group was not found.

**User response:** Correct the input and reissue the command.

---

**6027-2120 [E] Unable to delete recovery group *name* on nodes *nodeNames*.**

**Explanation:** The recovery group could not be deleted on the specified nodes.

**User response:** Perform problem determination.

---

**6027-2121 [I] Recovery group *name* deleted on node *nodeName*.**

**Explanation:** The recovery group has been deleted.

**User response:** This is an informational message.

---

**6027-2122 [E] The number of spares (*numberOfSpares*) must be less than the number of pdisks (*numberOfPdisks*) being created.**

**Explanation:** The number of spares specified must be less than the number of pdisks that are being created.

**User response:** Correct the input and reissue the command.

---

**6027-2123 [E] The GPFS daemon is down on the *vdiskName* servers.**

**Explanation:** The GPFS daemon was down on the vdisk servers when **mmdelvdisk** was issued.

**User response:** Start the GPFS daemon on the specified nodes and issue the specified **mmdelvdisk** command.

---

**6027-2124 [E] Vdisk *vdiskName* is still NSD *nsdName*. Use the **mmdelnsd** command.**

**Explanation:** The specified vdisk is still an NSD.

**User response:** Use the **mmdelnsd** command.

---

**6027-2125 [E] *nsdName* is a vdisk-based NSD and cannot be used as a tiebreaker disk.**

**Explanation:** Vdisk-based NSDs cannot be specified as tiebreaker disks.

**User response:** Correct the input and reissue the command.

---

**6027-2126 [I] No recovery groups were found.**

**Explanation:** A command searched for recovery groups but found none.

**User response:** None. Informational message only.

---

**6027-2127 [E] Disk descriptor *descriptor* refers to an existing pdisk.**

**Explanation:** The specified disk descriptor refers to an existing pdisk.

---

**User response:** Specify another disk that is not an existing pdisk.

---

**6027-2128 [E] The *attribute* attribute must be configured to use *hostname* as a recovery group server.**

**Explanation:** The specified GPFS configuration attributes must be configured to use the node as a recovery group server.

**User response:** Use the **mmchconfig** command to set the attributes, then reissue the command.

---

**6027-2129 [E] Vdisk block size (*blockSize*) must match the file system block size (*blockSize*).**

**Explanation:** The specified NSD is a vdisk with a block size that does not match the block size of the file system.

**User response:** Reissue the command using block sizes that match.

---

**6027-2130 [E] Could not find an active server for recovery group *name*.**

**Explanation:** A command was issued that acts on a recovery group, but no active server was found for the specified recovery group.

**User response:** Perform problem determination.

---

**6027-2131 [E] Cannot create an NSD on a log vdisk.**

**Explanation:** The specified disk is a log vdisk; it cannot be used for an NSD.

**User response:** Specify another disk that is not a log vdisk.

---

**6027-2132 [E] Log vdisk *vdiskName* cannot be deleted while there are other vdisks in recovery group *name*.**

**Explanation:** The specified disk is a log vdisk; it must be the last vdisk deleted from the recovery group.

**User response:** Delete the other vdisks first.

---

**6027-2133 [E] Unable to delete recovery group *name*; vdisks are still defined.**

**Explanation:** Cannot delete a recovery group while there are still vdisks defined.

**User response:** Delete all the vdisks first.

---

---

**6027-2134** Node *nodeName* cannot be used as an NSD server for Persistent Reserve disk *diskName* because it is not a Linux node.

**Explanation:** There was an attempt to enable Persistent Reserve for a disk, but not all of the NSD server nodes are running Linux.

**User response:** Correct the configuration and enter the command again.

---

**6027-2135** All nodes in the cluster must be running AIX to enable Persistent Reserve for SAN attached disk *diskName*.

**Explanation:** There was an attempt to enable Persistent Reserve for a SAN-attached disk, but not all nodes in the cluster are running AIX.

**User response:** Correct the configuration and run the command again.

---

**6027-2136** All NSD server nodes must be running AIX to enable Persistent Reserve for disk *diskName*.

**Explanation:** There was an attempt to enable Persistent Reserve for the specified disk, but not all NSD servers are running AIX.

**User response:** Correct the configuration and enter the command again.

---

**6027-2137** An attempt to clear the Persistent Reserve reservations on disk *diskName* failed.

**Explanation:** You are importing a disk into a cluster in which Persistent Reserve is disabled. An attempt to clear the Persistent Reserve reservations on the disk failed.

**User response:** Correct the configuration and enter the command again.

---

**6027-2138** The cluster must be running either all AIX or all Linux nodes to change Persistent Reserve disk *diskName* to a SAN-attached disk.

**Explanation:** There was an attempt to redefine a Persistent Reserve disk as a SAN attached disk, but not all nodes in the cluster were running either all AIX or all Linux nodes.

**User response:** Correct the configuration and enter the command again.

---

**6027-2139** NSD server nodes must be running either all AIX or all Linux to enable Persistent Reserve for disk *diskName*.

**Explanation:** There was an attempt to enable Persistent Reserve for a disk, but not all NSD server nodes were running all AIX or all Linux nodes.

**User response:** Correct the configuration and enter the command again.

---

**6027-2140** All NSD server nodes must be running AIX or all running Linux to enable Persistent Reserve for disk *diskName*.

**Explanation:** Attempt to enable Persistent Reserve for a disk while not all NSD server nodes are running AIX or all running Linux.

**User response:** Correct the configuration first.

---

**6027-2141** Disk *diskName* is not configured as a regular hdisk.

**Explanation:** In an AIX only cluster, Persistent Reserve is supported for regular hdisks only.

**User response:** Correct the configuration and enter the command again.

---

**6027-2142** Disk *diskName* is not configured as a regular generic disk.

**Explanation:** In a Linux only cluster, Persistent Reserve is supported for regular generic or device mapper virtual disks only.

**User response:** Correct the configuration and enter the command again.

---

**6027-2143** Mount point *mountPoint* can not be part of automount directory *automountDir*.

**Explanation:** The mount point cannot be the parent directory of the automount directory.

**User response:** Specify a mount point that is not the parent of the automount directory.

---

**6027-2144 [E]** The *lockName* lock for file system *fileSystem* is busy.

**Explanation:** More than one process is attempting to obtain the specified lock.

**User response:** Repeat the command. If the problem persists, verify that there are no blocked processes.

---



---

**6027-2145 [E] Internal remote command 'mmremote command' no longer supported.**

**Explanation:** A GPFS administration command invoked an internal remote command which is no longer supported. Backward compatibility for remote commands are only supported for release 3.4 and newer.

**User response:** All nodes within the cluster must be at release 3.4 or newer. If all the cluster nodes meet this requirement, contact the IBM Support Center.

---

**6027-2147 [E] BlockSize must be specified in disk descriptor.**

**Explanation:** The blockSize positional parameter in a vdisk descriptor was empty. The bad disk descriptor is displayed following this message.

**User response:** Correct the input and reissue the command.

---

**6027-2148 [E] nodeName is not a valid recovery group server for recoveryGroupName.**

**Explanation:** The server name specified is not one of the defined recovery group servers.

**User response:** Correct the input and reissue the command.

---

**6027-2149 [E] Could not get recovery group information from an active server.**

**Explanation:** A command that needed recovery group information failed; the GPFS daemons may have become inactive or the recovery group is temporarily unavailable.

**User response:** Reissue the command.

---

**6027-2150 The archive system client backupProgram could not be found or is not executable.**

**Explanation:** TSM dsmc or other specified backup or archive system client could not be found.

**User response:** Verify that TSM is installed, dsmc can be found in the installation location or that the archiver client specified is executable.

---

**6027-2151 The path directoryPath is not contained in the snapshot snapshotName.**

**Explanation:** The directory path supplied is not contained in the snapshot named with the -S parameter.

**User response:** Correct the directory path or snapshot name supplied, or omit -S and the snapshot name in the command.

---

**6027-2152 The path directoryPath containing image archives was not found.**

**Explanation:** The directory path supplied does not contain the expected image files to archive into TSM.

**User response:** Correct the directory path name supplied.

---

**6027-2153 The archiving system backupProgram exited with status return code. Image backup files have been preserved in globalWorkDir**

**Explanation:** Archiving system executed and returned a non-zero exit status due to some error.

**User response:** Examine archiver log files to discern the cause of the archiver's failure. Archive the preserved image files from the indicated path.

---

**6027-2154 Unable to create a policy file for image backup in policyFilePath.**

**Explanation:** A temporary file could not be created in the global shared directory path.

**User response:** Check or correct the directory path name supplied.

---

**6027-2155 File system fileSystem must be mounted read only for restore.**

**Explanation:** The empty file system targeted for restoration must be mounted in read only mode during restoration.

**User response:** Unmount the file system on all nodes and remount it read only, then try the command again.

---

**6027-2156 The image archive index ImagePath could not be found.**

**Explanation:** The archive image index could be found in the specified path

**User response:** Check command arguments for correct specification of image path, then try the command again.

---

**6027-2157 The image archive index ImagePath is corrupt or incomplete.**

**Explanation:** The archive image index specified is damaged.

**User response:** Check the archive image index file for corruption and remedy.

---

---

**6027-2158** Disk usage must be `dataOnly`, `metadataOnly`, `descOnly`, `dataAndMetadata`, `vdiskLog`, `vdiskLogTip`, `vdiskLogTipBackup`, or `vdiskLogReserved`.

**Explanation:** The disk usage positional parameter in a vdisk descriptor has a value that is not valid. The bad disk descriptor is displayed following this message.

**User response:** Correct the input and reissue the command.

---

**6027-2159 [E]** *parameter is not valid or missing in the vdisk descriptor.*

**Explanation:** The vdisk descriptor is not valid. The bad descriptor is displayed following this message.

**User response:** Correct the input and reissue the command.

---

**6027-2160 [E]** Vdisk *vdiskName* is already mapped to NSD *nsdName*.

**Explanation:** The command cannot create the specified NSD because the underlying vdisk is already mapped to a different NSD.

**User response:** Correct the input and reissue the command.

---

**6027-2161 [E]** NAS servers cannot be specified when creating an NSD on a vdisk.

**Explanation:** The command cannot create the specified NSD because servers were specified and the underlying disk is a vdisk.

**User response:** Correct the input and reissue the command.

---

**6027-2162 [E]** Cannot set `nsdRAIDTracks` to zero; *nodeName* is a recovery group server.

**Explanation:** `nsdRAIDTracks` cannot be set to zero while the node is still a recovery group server.

**User response:** Modify or delete the recovery group and reissue the command.

---

**6027-2163 [E]** Vdisk *name* not found in the daemon. Recovery may be occurring. The disk will not be deleted.

**Explanation:** GPFS cannot find the specified vdisk. This can happen if recovery is taking place and the recovery group is temporarily inactive.

**User response:** Reissue the command. If the recovery group is damaged, specify the `-p` option.

---

**6027-2164 [E]** Disk descriptor for *name* refers to an existing pdisk.

**Explanation:** The specified pdisk already exists.

**User response:** Correct the command invocation and try again.

---

**6027-2165 [E]** Node *nodeName* cannot be used as a server of both vdisks and non-vdisk NSDs.

**Explanation:** The command specified an action that would have caused vdisks and non-vdisk NSDs to be defined on the same server. This is not a supported configuration.

**User response:** Correct the command invocation and try again.

---

**6027-2166 [E]** GPFS Native RAID is not configured.

**Explanation:** GPFS Native RAID is not configured on this node.

**User response:** Reissue the command on the appropriate node.

---

**6027-2167 [E]** Device *deviceName* does not exist or is not active on this node.

**Explanation:** The specified device does not exist or is not active on the node.

**User response:** Reissue the command on the appropriate node.

---

**6027-2168 [E]** The GPFS cluster must be shut down before downloading firmware to port cards.

**Explanation:** The GPFS daemon must be down on all nodes in the cluster before attempting to download firmware to a port card.

**User response:** Stop GPFS on all nodes and reissue the command.

---

**6027-2169** Unable to disable Persistent Reserve on the following disks: *diskList*

**Explanation:** The command was unable to disable Persistent Reserve on the specified disks.

**User response:** Examine the disks and additional error information to determine if the disks should support Persistent Reserve. Correct the problem and reissue the command.

---

---

**6027-2170 [E] Recovery group *recoveryGroupName* does not exist or is not active.**

**Explanation:** A command was issued to a recovery group that does not exist or is not in the active state.

**User response:** Reissue the command with a valid recovery group name or wait for the recovery group to become active.

---

**6027-2171 [E] *objectType* *objectName* already exists in the cluster.**

**Explanation:** The file system being imported contains an object with a name that conflicts with the name of an existing object in the cluster.

**User response:** If possible, remove the object with the conflicting name.

---

**6027-2172 [E] Errors encountered while importing GPFS Native RAID objects.**

**Explanation:** Errors were encountered while trying to import a GPFS Native RAID based file system. No file systems will be imported.

**User response:** Check the previous error messages and if possible, correct the problems.

---

**6027-2173 [I] Use `mmchrecoverygroup` to assign and activate servers for the following recovery groups (automatically assigns NSD servers as well): *recoveryGroupList***

**Explanation:** The `mmimportfs` command imported the specified recovery groups. These must have servers assigned and activated.

**User response:** After the `mmimportfs` command finishes, use the `mmchrecoverygroup` command to assign NSD server nodes as needed.

---

**6027-2174 Option *option* can be specified only in conjunction with *option*.**

**Explanation:** The cited option cannot be specified by itself.

**User response:** Correct the input and reissue the command.

---

**6027-2175 [E] Exported path *exportPath* does not exist**

**Explanation:** The directory or one of the components in the directory path to be exported does not exist.

**User response:** Correct the input and reissue the command.

---



---

**6027-2176 [E] `mmchattr` for *fileName* failed.**

**Explanation:** The command to change the attributes of the file failed.

**User response:** Check the previous error messages and correct the problems.

---

**6027-2177 [E] Cannot create file *fileName*.**

**Explanation:** The command to create the specified file failed.

**User response:** Check the previous error messages and correct the problems.

---

**6027-2178 File *fileName* does not contain any NSD descriptors or stanzas.**

**Explanation:** The input file should contain at least one NSD descriptor or stanza.

**User response:** Correct the input file and reissue the command.

---

**6027-2181 [E] Failover is allowed only for single-writer, independent-writer filesets.**

**Explanation:** The fileset AFM mode is not compatible with the requested operation.

**User response:** Check the previous error messages and correct the problems.

---

**6027-2182 [E] Resync is allowed only for single-writer filesets.**

**Explanation:** The fileset AFM mode is not compatible with the requested operation.

**User response:** Check the previous error messages and correct the problems.

---

**6027-2183 [E] Peer snapshots using `mmpsnap` are allowed only for single-writer or primary filesets.**

**Explanation:** The fileset AFM mode is not compatible with the requested operation.

**User response:** Check the previous error messages and correct the problems.

---

**6027-2184 [E] If the recovery group is damaged, issue `mmdelrecoverygroup` *name* `-p`.**

**Explanation:** No active servers were found for the recovery group that is being deleted. If the recovery group is damaged the `-p` option is needed.

**User response:** Perform diagnosis and reissue the command.

---

---

**6027-2185 [E]** There are no `pdisk` stanzas in the input file *fileName*.

**Explanation:** The `mmcrrecoverygroup` input stanza file has no `pdisk` stanzas.

**User response:** Correct the input file and reissue the command.

---

**6027-2186 [E]** There were no valid `vdisk` stanzas in the input file *fileName*.

**Explanation:** The `mmcrvdisk` input stanza file has no valid `vdisk` stanzas.

**User response:** Correct the input file and reissue the command.

---

**6027-2187 [E]** Could not get `pdisk` information for the following recovery groups:  
*recoveryGroupList*

**Explanation:** An `mmlspdisk all` command could not query all of the recovery groups because some nodes could not be reached.

**User response:** None.

---

**6027-2188** Unable to determine the local node identity.

**Explanation:** The command is not able to determine the identity of the local node. This can be the result of a disruption in the network over which the GPFS daemons communicate.

**User response:** Ensure the GPFS daemon network (as identified in the output of the `mmlscluster` command on a good node) is fully operational and reissue the command.

---

**6027-2189 [E]** Action *action* is allowed only for read-only filesets.

**Explanation:** The specified action is only allowed for read-only filesets.

**User response:** None.

---

**6027-2190 [E]** Cannot prefetch file *fileName*. The file does not belong to fileset *fileset*.

**Explanation:** The requested file does not belong to the fileset.

**User response:** None.

---

**6027-2191 [E]** `Vdisk` *vdiskName* not found in recovery group *recoveryGroupName*.

**Explanation:** The `mmdelvdisk` command was invoked with the `--recovery-group` option to delete one or more `vdisks` from a specific recovery group. The specified

`vdisk` does not exist in this recovery group.

**User response:** Correct the input and reissue the command.

---

**6027-2193 [E]** Recovery group *recoveryGroupName* must be active on the primary server *serverName*.

**Explanation:** The recovery group must be active on the specified node.

**User response:** Use the `mmchrecoverygroup` command to activate the group and reissue the command.

---

**6027-2194 [E]** The state of fileset *filesetName* is Expired; prefetch cannot be performed.

**Explanation:** The prefetch operation cannot be performed on filesets that are in the Expired state.

**User response:** None.

---

**6027-2195 [E]** Error getting snapshot ID for *snapshotName*.

**Explanation:** The command was unable to obtain the resync snapshot ID.

**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.

---

**6027-2196 [E]** Resync is allowed only when the fileset queue is in active state.

**Explanation:** This operation is allowed only when the fileset queue is in active state.

**User response:** None.

---

**6027-2197 [E]** Empty file encountered when running the `mmafmctl flushPending` command.

**Explanation:** The `mmafmctl flushPending` command did not find any entries in the file specified with the `--list-file` option.

**User response:** Correct the input file and reissue the command.

---

**6027-2198 [E]** Cannot run the `mmafmctl flushPending` command on directory *dirName*.

**Explanation:** The `mmafmctl flushPending` command cannot be issued on this directory.

**User response:** Correct the input and reissue the command.

---

**6027-2199 [E] No enclosures were found.**

**Explanation:** A command searched for disk enclosures but none were found.

**User response:** None.

---

**6027-2200 [E] Cannot have multiple nodes updating firmware for the same enclosure. Enclosure *serialNumber* is already being updated by node *nodeName*.**

**Explanation:** The `mmchenclosure` command was called with multiple nodes updating the same firmware.

**User response:** Correct the node list and reissue the command.

---

**6027-2201 [E] The `mmafmctl flushPending` command completed with errors.**

**Explanation:** An error occurred while flushing the queue.

**User response:** Examine the GPFS log to identify the cause.

---

**6027-2202 [E] There is a SCSI-3 PR reservation on disk *diskname*. `mmcrnsd` cannot format the disk because the cluster is not configured as PR enabled.**

**Explanation:** The specified disk has a SCSI-3 PR reservation, which prevents the `mmcrnsd` command from formatting it.

**User response:** Clear the PR reservation by following the instructions in "Clearing a leftover Persistent Reserve reservation" on page 139.

---

**6027-2203 Node *nodeName* is not a gateway node.**

**Explanation:** The specified node is not a gateway node.

**User response:** Designate the node as a gateway node or specify a different node on the command line.

---

**6027-2204 AFM target map *mapName* is already defined.**

**Explanation:** A request was made to create an AFM target map with the cited name, but that map name is already defined.

**User response:** Specify a different name for the new AFM target map or first delete the current map definition and then recreate it.

---

**6027-2205 There are no AFM target map definitions.**

**Explanation:** A command searched for AFM target map definitions but found none.

**User response:** None. Informational message only.

---

**6027-2206 AFM target map *mapName* is not defined.**

**Explanation:** The cited AFM target map name is not known to GPFS.

**User response:** Specify an AFM target map known to GPFS.

---

**6027-2207 Node *nodeName* is being used as a gateway node for the AFM cluster *clusterName*.**

**Explanation:** The specified node is defined as a gateway node for the specified AFM cluster.

**User response:** If you are trying to delete the node from the GPFS cluster or delete the gateway node role, you must remove it from the export server map.

---

**6027-2208 [E] *commandName* is already running in the cluster.**

**Explanation:** Only one instance of the specified command is allowed to run.

**User response:** None.

---

**6027-2209 [E] Unable to list *objectName* on node *nodeName*.**

**Explanation:** A command was unable to list the specific object that was requested.

**User response:** None.

---

**6027-2210 [E] Unable to build a storage enclosure inventory file on node *nodeName*.**

**Explanation:** A command was unable to build a storage enclosure inventory file. This is a temporary file that is required to complete the requested command.

**User response:** None.

---

**6027-2211 [E] Error collecting firmware information on node *nodeName*.**

**Explanation:** A command was unable to gather firmware information from the specified node.

**User response:** Ensure the node is active and retry the command.

---

---

**6027-2212 [E] Firmware update file *updateFile* was not found.**

**Explanation:** The `mmchfirmware` command could not find the specified firmware update file to load.

**User response:** Locate the firmware update file and retry the command.

---

**6027-2213 [E] Pdisk path redundancy was lost while updating enclosure firmware.**

**Explanation:** The `mmchfirmware` command lost paths after loading firmware and rebooting the Enclosure Services Module.

**User response:** Wait a few minutes and then retry the command. GPFS might need to be shut down to finish updating the enclosure firmware.

---

**6027-2214 [E] Timeout waiting for firmware to load.**

**Explanation:** A storage enclosure firmware update was in progress, but the update did not complete within the expected time frame.

**User response:** Wait a few minutes, and then use the `mmlsfirmware` command to ensure the operation completed.

---

**6027-2215 [E] Storage enclosure *serialNumber* not found.**

**Explanation:** The specified storage enclosure was not found.

**User response:** None.

---

**6027-2216 Quota management is disabled for file system *fileSystem*.**

**Explanation:** Quota management is disabled for the specified file system.

**User response:** Enable quota management for the file system.

---

**6027-2217 [E] Error *errno* updating firmware for drives *driveList*.**

**Explanation:** The firmware load failed for the specified drives. Some of the drives may have been updated.

**User response:** None.

---

**6027-2218 [E] Storage enclosure *serialNumber* component *componentType* component ID *componentId* not found.**

**Explanation:** The `mmchenclousure` command could not find the component specified for replacement.

**User response:** Use the `mmlsenclousure` command to

determine valid input and then retry the command.

---

**6027-2219 [E] Storage enclosure *serialNumber* component *componentType* component ID *componentId* did not fail. Service is not required.**

**Explanation:** The component specified for the `mmchenclousure` command does not need service.

**User response:** Use the `mmlsenclousure` command to determine valid input and then retry the command.

---

**6027-2220 [E] Recovery group *name* has pdisks with missing paths. Consider using the `-v no` option of the `mmchrecoverygroup` command.**

**Explanation:** The `mmchrecoverygroup` command failed because all the servers could not see all the disks, and the primary server is missing paths to disks.

**User response:** If the disks are cabled correctly, use the `-v no` option of the `mmchrecoverygroup` command.

---

**6027-2221 [E] Error determining redundancy of enclosure *serialNumber* `ESM` *esmName*.**

**Explanation:** The `mmchrecoverygroup` command failed. Check the following error messages.

**User response:** Correct the problem and retry the command.

---

**6027-2222 [E] Storage enclosure *serialNumber* already has a newer firmware version: *firmwareLevel*.**

**Explanation:** The `mmchfirmware` command found a newer level of firmware on the specified storage enclosure.

**User response:** If the intent is to force on the older firmware version, use the `-v no` option.

---

**6027-2223 [E] Storage enclosure *serialNumber* is not redundant. Shutdown GPFS in the cluster and retry the `mmchfirmware` command.**

**Explanation:** The `mmchfirmware` command found a non-redundant storage enclosure. Proceeding could cause loss of data access.

**User response:** Shutdown GPFS in the cluster and retry the `mmchfirmware` command.

---

**6027-2224 [E] Peer snapshot creation failed. Error code *errorCode*.**

**Explanation:** For an active fileset, check the AFM target configuration for peer snapshots. Ensure there is at least one gateway node configured for the cluster.

Examine the preceding messages and the GPFS log for additional details.

**User response:** Correct the problems and reissue the command.

**6027-2225 [E] Peer snapshot successfully deleted at cache. The delete snapshot operation failed at home. Error code *errorCode*.**

**Explanation:** For an active fileset, check the AFM target configuration for peer snapshots. Ensure there is at least one gateway node configured for the cluster. Examine the preceding messages and the GPFS log for additional details.

**User response:** Correct the problems and reissue the command.

**6027-2226 [E] Invalid firmware update file.**

**Explanation:** An invalid firmware update file was specified for the `mmchfirmware` command.

**User response:** Reissue the command with a valid update file.

**6027-2227 [E] Failback is allowed only for independent-writer filesets.**

**Explanation:** Failback operation is allowed only for independent-writer filesets.

**User response:** Check the fileset mode.

**6027-2228 [E] The daemon version (*daemonVersion*) on node *nodeName* is lower than the daemon version (*daemonVersion*) on node *nodeName*.**

**Explanation:** A command was issued that requires nodes to be at specific levels, but the affected GPFS servers are not at compatible levels to support this operation.

**User response:** Update the GPFS code on the specified servers and retry the command.

**6027-2229 [E] Cache Eviction/Prefetch is not allowed for Primary and Secondary mode filesets.**

**Explanation:** Cache eviction/prefetch is not allowed for primary and secondary mode filesets.

**User response:** None.

**6027-2230 [E] `afmTarget=newTargetString` is not allowed. To change the AFM target, use `mmafmctl failover` with the `--target-only` option. For primary filesets, use `mmafmctl changeSecondary`.**

**Explanation:** The `mmchfileset` command cannot be used to change the NFS server or IP address of the home cluster.

**User response:** To change the AFM target, use the `mmafmctl failover` command and specify the `--target-only` option. To change the AFM target for primary filesets, use the `mmafmctl changeSecondary` command.

**6027-2231 [E] The specified block size *blockSize* is smaller than the system page size *pageSize*.**

**Explanation:** The file system block size cannot be smaller than the system memory page size.

**User response:** Specify a block size greater than or equal to the system memory page size.

**6027-2232 [E] Peer snapshots are allowed only for targets using the NFS protocol.**

**Explanation:** The `mmpsnap` command can be used to create snapshots only for filesets that are configured to use the NFS protocol.

**User response:** Specify a valid fileset target.

**6027-2233 [E] Fileset *filesetName* in file system *filesystemName* does not contain peer snapshot *snapshotName*. The delete snapshot operation failed at cache. Error code *errorCode*.**

**Explanation:** The specified snapshot name was not found. The command expects the name of an existing peer snapshot of the active fileset in the specified file system.

**User response:** Reissue the command with a valid peer snapshot name.

**6027-2234 [E] Use the `mmafmctl converttoprimary` command for converting to primary fileset.**

**Explanation:** Converting to a primary fileset is not allowed directly.

**User response:** Check the previous error messages and correct the problems.

**6027-2235 [E] Only independent filesets can be converted to secondary filesets.**

**Explanation:** Converting to secondary filesets is allowed only for independent filesets.

**User response:** None.

---

**6027-2236 [E]** The CPU architecture on this node does not support tracing in *traceMode* mode. Switching to *traceMode* mode.

**Explanation:** The CPU does not have constant time stamp counter capability, which is required for overwrite trace mode. The trace has been enabled in blocking mode.

**User response:** Update the configuration parameters to use the trace facility in blocking mode or replace this node with modern CPU architecture.

---

**6027-2237 [W]** An image backup made from the live file system may not be usable for image restore. Specify a valid global snapshot for image backup.

**Explanation:** The `mmimgbackup` command should always be used with a global snapshot to make a consistent image backup of the file system.

**User response:** Correct the command invocation to include the `-S` option to specify either a global snapshot name or a directory path that includes the snapshot root directory for the file system and a valid global snapshot name.

---

**6027-2238 [E]** Use the `mmafmctl convertToSecondary` command for converting to secondary.

**Explanation:** Converting to secondary is allowed by using the `mmafmctl convertToSecondary` command.

**User response:** None.

---

**6027-2239 [E]** Drive `serialNumber` *serialNumber* is being managed by server *nodeName*. Reissue the `mmchfirmware` command for server *nodeName*.

**Explanation:** The `mmchfirmware` command was issued to update a specific disk drive which is not currently being managed by this node.

**User response:** Reissue the command specifying the active server.

---

**6027-2240 [E]** Option is not supported for a secondary fileset.

**Explanation:** This option cannot be set for a secondary fileset.

**User response:** None.

---

**6027-2241 [E]** Node *nodeName* is not a CES node.

**Explanation:** A Cluster Export Service command specified a node that is not defined as a CES node.

**User response:** Reissue the command specifying a CES node.

---

**6027-2242 [E]** Error in configuration file.

**Explanation:** The `mmnfs export load loadCfgFile` command found an error in the NFS configuration files.

**User response:** Correct the configuration file error.

---

**6027-2245 [E]** To change the AFM target, use `mmafmctl changeSecondary` for the primary.

**Explanation:** Failover with the `targetonly` option can be run on a primary fileset.

**User response:** None.

---

**6027-2246 [E]** Timeout executing function: *functionName* (`return code=returnCode`).

**Explanation:** The `executeCommandWithTimeout` function was called but it timed out.

**User response:** Correct the problem and issue the command again.

---

**6027-2247 [E]** Creation of *exchangeDir* failed.

**Explanation:** A Cluster Export Service command was unable to create the CCR exchange directory.

**User response:** Correct the problem and issue the command again.

---

**6027-2248 [E]** CCR command failed: *command*

**Explanation:** A CCR update command failed.

**User response:** Correct the problem and issue the command again.

---

**6027-2249 [E]** Error getting next *nextName* from CCR.

**Explanation:** An expected value from CCR was not obtained.

**User response:** Issue the command again.

---

**6027-2250 [E]** Error putting next *nextName* to CCR, new ID: *newExpid* version: *version*

**Explanation:** A CCR value update failed.

**User response:** Issue the command again.

---

**6027-2251 [E]** Error retrieving configuration file: *configFile*

**Explanation:** Error retrieving configuration file from CCR.

**User response:** Issue the command again.



---

**6027-2252 [E] Error reading export configuration file (return code: *returnCode*).**

**Explanation:** A CES command was unable to read the export configuration file.

**User response:** Correct the problem and issue the command again.

---

**6027-2253 [E] Error creating the internal export data objects (return code *returnCode*).**

**Explanation:** A CES command was unable to create an export data object.

**User response:** Correct the problem and issue the command again.

---

**6027-2254 [E] Error creating single export output, export *exportPath* not found (return code *returnCode*).**

**Explanation:** A CES command was unable to create a single export print output.

**User response:** Correct the problem and reissue the command.

---

**6027-2255 [E] Error creating export output (return code: *returnCode*).**

**Explanation:** A CES command was unable to create the export print output.

**User response:** Correct the problem and issue the command again.

---

**6027-2256 [E] Error creating the internal export output file string array (return code: *returnCode*).**

**Explanation:** A CES command was unable to create the array for print output.

**User response:** Correct the problem and issue the command again.

---

**6027-2257 [E] Error deleting export, export *exportPath* not found (return code: *returnCode*).**

**Explanation:** A CES command was unable to delete an export. The *exportPath* was not found.

**User response:** Correct the problem and issue the command again.

---

**6027-2258 [E] Error writing export configuration file to CCR (return code: *returnCode*).**

**Explanation:** A CES command was unable to write configuration file to CCR.

**User response:** Correct the problem and issue the command again.

---



---

**6027-2259 [E] The path *exportPath* to create the export does not exist (return code:*returnCode*).**

**Explanation:** A CES command was unable to create an export because the path does not exist.

**User response:** Correct the problem and issue the command again.

---

**6027-2260 [E] The path *exportPath* to create the export is invalid (return code: *returnCode*).**

**Explanation:** A CES command was unable to create an export because the path is invalid.

**User response:** Correct the problem and issue the command again.

---

**6027-2261 [E] Error creating new export object, invalid data entered (return code: *returnCode*).**

**Explanation:** A CES command was unable to add an export because the input data is invalid.

**User response:** Correct the problem and issue the command again.

---

**6027-2262 [E] Error creating new export object; getting new export ID (return code: *returnCode*).**

**Explanation:** A CES command was unable to add an export. A new export ID was not obtained.

**User response:** Correct the problem and issue the command again.

---

**6027-2263 [E] Error adding export; new export path *exportPath* already exists.**

**Explanation:** A CES command was unable to add an export because the path already exists.

**User response:** Correct the problem and issue the command again.

---

**6027-2264 [E] The --servers option is only used to provide names for primary and backup server configurations. Provide a maximum of two server names.**

**Explanation:** An input node list has too many nodes specified.

**User response:** Verify the list of nodes and shorten the list to the supported number.

---

**6027-2265 [E] Cannot convert fileset to secondary fileset.**

**Explanation:** Fileset cannot be converted to a secondary fileset.

**User response:** None.

---

---

**6027-2266 [E]** The snapshot names that start with `psnap-rpo` or `psnap0-rpo` are reserved for RPO.

**Explanation:** The specified snapshot name starts with `psnap-rpo` or `psnap0-rpo`, which are reserved for RPO snapshots.

**User response:** Use a different snapshot name for the `mmcrsnapshot` command.

---

**6027-2267 [I]** Fileset `filesetName` in file system `fileSystem` is either unlinked or being deleted. Home delete-snapshot operation was not queued.

**Explanation:** The command expects that the peer snapshot at home is not deleted because the fileset at cache is either unlinked or being deleted.

**User response:** Delete the snapshot at home manually.

---

**6027-2268 [E]** This is already a secondary fileset.

**Explanation:** The fileset is already a secondary fileset.

**User response:** None.

---

**6027-2269 [E]** Adapter `adapterIdentifier` was not found.

**Explanation:** The specified adapter was not found.

**User response:** Specify an existing adapter and reissue the command.

---

**6027-2270 [E]** Error `errno` updating firmware for adapter `adapterIdentifier`.

**Explanation:** The firmware load failed for the specified adapter.

**User response:** None.

---

**6027-2271 [E]** Error locating the reference client IP `ipAddress`, return code: `returnCode`

**Explanation:** The reference IP address for reordering a client could not be found for the given export path.

**User response:** Correct the problem and try again.

---

**6027-2272 [E]** Error removing the requested IP address `ipAddress` from a client declaration, return code: `returnCode`

**Explanation:** One of the specified IP addresses to remove could not be found in any client declaration for the given export path.

**User response:** Correct the problem and try again.

---

**6027-2273 [E]** Error adding the requested IP address `ipAddress` to a client declaration, return code: `returnCode`

**Explanation:** One of the specified IP addresses to add could not be applied for the given export path.

**User response:** Correct the problem and try again.

---

**6027-2274 [E]** Error changing the requested IP address `ipAddress` of a client declaration, return code: `returnCode`

**Explanation:** The client change could not be applied for the given export path.

**User response:** Correct the problem and try again.

---

**6027-2275 [E]** Unable to determine the status of DASD device `dasdDevice`

**Explanation:** The `dasdview` command failed.

**User response:** Examine the preceding messages, correct the problem, and reissue the command.

---

**6027-2276 [E]** The specified DASD device `dasdDevice` is not properly formatted. It is not an ECKD-type device, or it has a format other than CDL or LDL, or it has a block size other than 4096.

**Explanation:** The specified device is not properly formatted.

**User response:** Correct the problem and reissue the command.

---

**6027-2277 [E]** Unable to determine if DASD device `dasdDevice` is partitioned.

**Explanation:** The `fdasd` command failed.

**User response:** Examine the preceding messages, correct the problem, and reissue the command.

---

**6027-2278 [E]** Cannot partition DASD device `dasdDevice`; it is already partitioned.

**Explanation:** The specified DASD device is already partitioned.

**User response:** Remove the existing partitions, or reissue the command using the desired partition name.

---

**6027-2279 [E]** Unable to partition DASD device `dasdDevice`

**Explanation:** The `fdasd` command failed.

**User response:** Examine the preceding messages, correct the problem, and reissue the command.

---

**6027-2280 [E]** The DASD device with bus ID *busID* cannot be found or it is in use.

**Explanation:** The `chccwdev` command failed.

**User response:** Examine the preceding messages, correct the problem, and reissue the command.

---

**6027-2281 [E]** Error *errno* updating firmware for enclosure *enclosureIdentifier*.

**Explanation:** The firmware load failed for the specified enclosure.

**User response:** None.

---

**6027-2282 [E]** Action *action* is not allowed for secondary filesets.

**Explanation:** The specified action is not allowed for secondary filesets.

**User response:** None.

---

**6027-2283 [E]** Node *nodeName* is already a CES node.

**Explanation:** An `mmchnode` command attempted to enable CES services on a node that is already part of the CES cluster.

**User response:** Reissue the command specifying a node that is not a CES node.

---

**6027-2284 [E]** The fileset `afmshowhomesnapshot` value is 'yes'. The fileset mode cannot be changed.

**Explanation:** The fileset `afmshowhomesnapshot` attribute value is `yes`. The fileset mode change is not allowed.

**User response:** First change the attribute `afmshowhomesnapshot` value to `no`, and then issue the command again to change the mode.

---

**6027-2285 [E]** Deletion of initial snapshot *snapshotName* of fileset *filesetName* in file system *fileSystem* failed. The delete fileset operation failed at cache. Error code *errorCode*.

**Explanation:** The deletion of the initial snapshot `psnap0` of *filesetName* failed. The primary and secondary filesets cannot be deleted without deleting the initial snapshot.

**User response:** None.

---



---

**6027-2286 [E]** RPO peer snapshots using `mmpsnap` are allowed only for primary filesets.

**Explanation:** RPO snapshots can be created only for primary filesets.

**User response:** Reissue the command with a valid primary fileset or without the `--rpo` option.

---

**6027-2287** The fileset needs to be linked to change `afmShowHomeSnapshot` to 'no'.

**Explanation:** The `afmShowHomeSnapshot` value cannot be changed to `no` if the fileset is unlinked.

**User response:** Link the fileset and reissue the command.

---

**6027-2288 [E]** Option *optionName* is not supported for AFM filesets.

**Explanation:** IAM modes are not supported for AFM filesets.

**User response:** None.

---

**6027-2289 [E]** Peer snapshot creation failed while running *subCommand*. Error code *errorCode*

**Explanation:** For an active fileset, check the AFM target configuration for peer snapshots. Ensure there is at least one gateway node configured for the cluster. Examine the preceding messages and the GPFS log for additional details.

**User response:** Correct the problems and reissue the command.

---

**6027-2290 [E]** The comment string should be less than 50 characters long.

**Explanation:** The comment/prefix string of the snapshot is longer than 50 characters.

**User response:** Reduce the comment string size and reissue the command.

---

**6027-2291 [E]** Peer snapshot creation failed while generating snapshot name. Error code *errorCode*

**Explanation:** For an active fileset, check the AFM target configuration for peer snapshots. Ensure there is at least one gateway node configured for the cluster. Examine the preceding messages and the GPFS log for additional details.

**User response:** Correct the problems and reissue the command.

---

---

**6027-2292 [E]** The initial snapshot *psnap0Name* does not exist. The peer snapshot creation failed. Error code *errorCode*

**Explanation:** For an active fileset, check the AFM target configuration for peer snapshots. Ensure the initial peer snapshot exists for the fileset. Examine the preceding messages and the GPFS log for additional details.

**User response:** Verify that the fileset is a primary fileset and that it has *psnap0* created and try again.

---

**6027-2293 [E]** The peer snapshot creation failed because fileset *filesetName* is in *filesetState* state.

**Explanation:** For an active fileset, check the AFM target configuration for peer snapshots. Ensure there is at least one gateway node configured for the cluster. Examine the preceding messages and the GPFS log for additional details.

**User response:** None. The fileset needs to be in active or dirty state.

---

**6027-2294 [E]** Removing older peer snapshots failed while obtaining snap IDs. Error code *errorCode*

**Explanation:** Ensure the fileset exists. Examine the preceding messages and the GPFS log for additional details.

**User response:** Verify that snapshots exist for the given fileset.

---

**6027-2295 [E]** Removing older peer snapshots failed while obtaining old snap IDs. Error code *errorCode*

**Explanation:** Ensure the fileset exists. Examine the preceding messages and the GPFS log for additional details.

**User response:** Verify that snapshots exist for the given fileset.

---

**6027-2296 [E]** Need a target to convert to the primary fileset.

**Explanation:** Need a target to convert to the primary fileset.

**User response:** Specify a target to convert to the primary fileset.

---

**6027-2297 [E]** The check-metadata and nocheck-metadata options are not supported for a non-AFM fileset.

**Explanation:** The check-metadata and

nocheck-metadata options are not supported for a non-AFM fileset.

**User response:** None.

---

**6027-2298 [E]** Only independent filesets can be converted to primary or secondary.

**Explanation:** Only independent filesets can be converted to primary or secondary.

**User response:** Specify an independent fileset.

---

**6027-2299 [E]** Issue the `mmafmctl getstate` command to check fileset state and if required issue `mmafmctl convertToPrimary`.

**Explanation:** Issue the `mmafmctl getstate` command to check fileset state and if required issue `mmafmctl convertToPrimary`.

**User response:** Issue the `mmafmctl getstate` command to check fileset state and if required issue `mmafmctl convertToPrimary`.

---

**6027-2300 [E]** The check-metadata and nocheck-metadata options are not supported for the primary fileset.

**Explanation:** The check-metadata and nocheck-metadata options are not supported for the primary fileset.

**User response:** None.

---

**6027-2301 [E]** The inband option is not supported for the primary fileset.

**Explanation:** The inband option is not supported for the primary fileset.

**User response:** None.

---

**6027-2302 [E]** AFM target cannot be changed for the primary fileset.

**Explanation:** AFM target cannot be changed for the primary fileset.

**User response:** None.

---

**6027-2303 [E]** The inband option is not supported for an AFM fileset.

**Explanation:** The inband option is not supported for an AFM fileset.

**User response:** None.

---

**6027-2304 [E]** Target cannot be changed for an AFM fileset.

**Explanation:** Target cannot be changed for an AFM fileset.

**User response:** None.

---

**6027-2305 [E]** The `mmafmctl convertToPrimary` command is not allowed for this primary fileset.

**Explanation:** The `mmafmctl convertToPrimary` command is not allowed for the primary fileset because it is not in `PrimInitFail` state.

**User response:** None.

---

**6027-2306 [E]** Failed to check for cached files while doing primary conversion from `filesetMode` mode.

**Explanation:** Failed to check for cached files while doing primary conversion.

**User response:** None.

---

**6027-2307 [E]** Uncached files present, run `prefetch` first.

**Explanation:** Uncached files present.

**User response:** Run `prefetch` and then do the conversion.

---

**6027-2308 [E]** Uncached files present, run `prefetch` first using policy output: `nodeDirFileOut`.

**Explanation:** Uncached files present.

**User response:** Run `prefetch` first using policy output.

---

**6027-2309 [E]** Conversion to primary not allowed for `filesetMode` mode.

**Explanation:** Conversion to primary not allowed for this mode.

**User response:** None.

---

**6027-2310 [E]** This option is available only for a primary fileset.

**Explanation:** This option is available only for a primary fileset.

**User response:** None.

---

**6027-2311 [E]** The `target-only` option is not allowed for a promoted primary without a target.

**Explanation:** The `target-only` option is not allowed for a promoted primary without a target.

**User response:** None.

---

**6027-2312 [E]** Need a target to setup the new secondary.

**Explanation:** Target is required to setup the new secondary.

**User response:** None.

---

**6027-2313 [E]** The `target-only` and `inband` options are not allowed together.

**Explanation:** The `target-only` and `inband` options are not allowed together.

**User response:** None.

---

**6027-2314 [E]** Could not run `commandName`. Verify that the Object protocol was installed.

**Explanation:** The `mmcesobjlscfg` command cannot find a prerequisite command on the system.

**User response:** Install the missing command and try again.

---

**6027-2315 [E]** Could not determine CCR file for service `serviceName`

**Explanation:** For the given service name, there is not a corresponding file in the CCR.

**User response:** None.

---

**6027-2316 [E]** Unable to retrieve file `fileName` from CCR using `command` command. Verify that the Object protocol is correctly installed.

**Explanation:** There was an error downloading a file from the CCR repository.

**User response:** Correct the error and try again.

---

**6027-2317 [E]** Unable to parse version number of file `fileName` from `mmccr` output

**Explanation:** The current version should be printed by `mmccr` when a file is extracted. The command could not read the version number from the output and failed.

**User response:** Investigate the failure in the CCR and fix the problem.

---

**6027-2318 [E]** Could not put `localFilePath` into the CCR as `ccrName`

**Explanation:** There was an error when trying to do an `fput` of a file into the CCR.

**User response:** Investigate the error and fix the problem.

---

**6027-2319 [I] Version mismatch during upload of *fileName* (*version*). Retrying.**

**Explanation:** The file could not be uploaded to the CCR because another process updated it in the meantime. The file will be downloaded, modified, and uploaded again.

**User response:** None. The upload will automatically be tried again.

---

**6027-2320 *directoryName* does not resolve to a directory in *deviceName*. The directory must be within the specified file system.**

**Explanation:** The cited directory does not belong to the specified file system.

**User response:** Correct the directory name and reissue the command.

---

**6027-2321 [E] AFM primary or secondary filesets cannot be created for file system *fileSystem* because version is less than *supportedVersion*.**

**Explanation:** The AFM primary or secondary filesets are not supported for a file system version that is less than 14.20.

**User response:** Upgrade the file system and reissue the command.

---

**6027-2322 [E] The OBJ service cannot be enabled because it is not installed. The file *fileName* was not found.**

**Explanation:** The node could not enable the CES OBJ service because of a missing binary or configuration file.

**User response:** Install the required software and retry the command.

---

**6027-2323 [E] The OBJ service cannot be enabled because the number of CES IPs below the minimum of *minValue* expected.**

**Explanation:** The value of CES IPs was below the minimum.

**User response:** Add at least *minValue* CES IPs to the cluster.

---

**6027-2324 [E] The object store for *serviceName* is either not a GPFS type or *mountPoint* does not exist.**

**Explanation:** The object store is not available at this time.

**User response:** Verify that *serviceName* is a GPFS type.

Verify that the *mountPoint* exists, the file system is mounted, or the fileset is linked.

---

**6027-2325 [E] File *fileName* does not exist in CCR. Verify that the Object protocol is correctly installed.**

**Explanation:** There was an error verifying Object config and ring files in the CCR repository.

**User response:** Correct the error and try again.

---

**6027-2326 [E] The OBJ service cannot be enabled because attribute *attributeName* for a CES IP has not been defined. Verify that the Object protocol is correctly installed.**

**Explanation:** There was an error verifying *attributeName* on CES IPs.

**User response:** Correct the error and try again.

---

**6027-2327 The snapshot *snapshotName* is the wrong scope for use in *targetType* backup**

**Explanation:** The snapshot specified is the wrong scope.

**User response:** Please provide a valid snapshot name for this backup type.

---

**6027-2329 [E] The fileset attributes cannot be set for the primary fileset with caching disabled.**

**Explanation:** The fileset attributes cannot be set for the primary fileset with caching disabled.

**User response:** None.

---

**6027-2330 [E] The outband option is not supported for AFM filesets.**

**Explanation:** The outband option is not supported for AFM filesets.

**User response:** None.

---

**6027-2331 [E] CCR value *ccrValue* not defined. The OBJ service cannot be enabled if identity authentication is not configured.**

**Explanation:** Object authentication type was not found.

**User response:** Configure identity authentication and try again.

---

**6027-2332 [E] Only regular independent filesets are converted to secondary filesets.**

**Explanation:** Only regular independent filesets can be converted to secondary filesets.

**User response:** Specify a regular independent fileset and run the command again.

---

**6027-2333 [E] Failed to disable *serviceName* service. Ensure *authType* authentication is removed.**

**Explanation:** Disable CES service failed because authentication was not removed.

**User response:** Remove authentication and retry.

---

**6027-2334 [E] Fileset *indFileset* cannot be changed because it has a dependent fileset *depFileset***

**Explanation:** Filesets with dependent filesets cannot be converted to primary or secondary.

**User response:** This operation cannot proceed until all the dependent filesets are unlinked.

---

**6027-2335 [E] Failed to convert fileset, because the policy to detect special files is failing.**

**Explanation:** The policy to detect special files is failing.

**User response:** Retry the command later.

---

**6027-2336 [E] Immutable/append-only files or clones copied from a snapshot are present, hence conversion is disallowed**

**Explanation:** Conversion is disallowed if immutable/append-only files or clones copied from a snapshot are present.

**User response:** Files should not be immutable/append-only.

---

**6027-2337 [E] Conversion to primary is not allowed at this time. Retry the command later.**

**Explanation:** Conversion to primary is not allowed at this time.

**User response:** Retry the command later.

---

**6027-2338 [E] Conversion to primary is not allowed because the state of the fileset is *filesetState*.**

**Explanation:** Conversion to primary is not allowed with the current state of the fileset.

**User response:** Retry the command later.

---



---

**6027-2339 [E] Orphans are present, run prefetch first.**

**Explanation:** Orphans are present.

**User response:** Run prefetch on the fileset and then do the conversion.

---

**6027-2340 [E] Fileset was left in PrimInitFail state. Take the necessary actions.**

**Explanation:** The fileset was left in PrimInitFail state.

**User response:** Take the necessary actions.

---

**6027-2341 [E] This operation can be done only on a primary fileset**

**Explanation:** This is not a primary fileset.

**User response:** None.

---

**6027-2342 [E] Failover/resync is currently running so conversion is not allowed**

**Explanation:** Failover/resync is currently running so conversion is not allowed.

**User response:** Retry the command later after failover/resync completes.

---

**6027-2343 [E] DR Setup cannot be done on a fileset with mode *filesetMode*.**

**Explanation:** Setup cannot be done on a fileset with this mode.

**User response:** None.

---

**6027-2344 [E] The GPFS daemon must be active on the node from which the *mmcmd* is executed with option *--inode-criteria* or *-o*.**

**Explanation:** The GPFS daemon needs to be active on the node where the command is issued with *--inode-criteria* or *-o* options.

**User response:** Run the command where the daemon is active.

---

**6027-2345 [E] The provided snapshot name must be unique to list filesets in a specific snapshot**

**Explanation:** The *mmlsfileset* command received a snapshot name that is not unique.

**User response:** Correct the command invocation or remove the duplicate named snapshots and try again.

---

---

**6027-2346 [E] The local node is not a CES node.**

**Explanation:** A local Cluster Export Service command was invoked on a node that is not defined as a Cluster Export Service node.

**User response:** Reissue the command on a CES node.

---

**6027-2347 [E] Error changing export, export *exportPath* not found.**

**Explanation:** A CES command was unable to change an export. The *exportPath* was not found.

**User response:** Correct problem and issue the command again.

---

**6027-2348 [E] A device for *directoryName* does not exist or is not active on this node.**

**Explanation:** The device containing the specified directory does not exist or is not active on the node.

**User response:** Reissue the command with a correct directory or on an appropriate node.

---

**6027-2349 [E] The fileset for *junctionName* does not exist in the *targetType* specified.**

**Explanation:** The fileset to back up cannot be found in the file system or snapshot specified.

**User response:** Reissue the command with a correct name for the fileset, snapshot, or file system.

---

**6027-2350 [E] The fileset for *junctionName* is not linked in the *targetType* specified.**

**Explanation:** The fileset to back up is not linked in the file system or snapshot specified.

**User response:** Relink the fileset in the file system. Optionally create a snapshot and reissue the command with a correct name for the fileset, snapshot, and file system.

---

**6027-2351 [E] One or more unlinked filesets (*filesetName*s) exist in the *targetType* specified. Check your filesets and try again.**

**Explanation:** The file system to back up contains one or more filesets that are unlinked in the file system or snapshot specified.

**User response:** Relink the fileset in the file system. Optionally create a snapshot and reissue the command with a correct name for the fileset, snapshot, and file system.

---



---

**6027-2352 The snapshot *snapshotName* could not be found for use by *commandName***

**Explanation:** The snapshot specified could not be located.

**User response:** Please provide a valid snapshot name.

---

**6027-2353 [E] The snapshot name cannot be generated.**

**Explanation:** The snapshot name cannot be generated.

**User response:** None.

---

**6027-2354 Node *nodeName* must be disabled as a CES node before trying to remove it from the GPFS cluster.**

**Explanation:** The specified node is defined as a CES node.

**User response:** Disable the CES node and try again.

---

**6027-2355 [E] Unable to reload *moduleName*. Node *hostname* should be rebooted.**

**Explanation:** Host adapter firmware was updated so the the specified module needs to be unloaded and reloaded. Linux does not display the new firmware level until the module is reloaded.

**User response:** Reboot the node.

---

**6027-2356 [E] Node *nodeName* is being used as a recovery group 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.

---

**6027-2357 [E] Root fileset cannot be converted to primary fileset.**

**Explanation:** Root fileset cannot be converted to the primary fileset.

**User response:** None.

---

**6027-2358 [E] Root fileset cannot be converted to secondary fileset.**

**Explanation:** Root fileset cannot be converted to the secondary fileset.

**User response:** None.

---



---

**6027-2359 [I] Attention: *command* is now enabled. This attribute can no longer be modified.**

**Explanation:** Indefinite retention protection is enabled. This value can not be changed in the future.

**User response:** None.

---

**6027-2360 [E] The current value of *command* is *attrName*. This value cannot be changed.**

**Explanation:** Indefinite retention protection is enabled for this cluster and this attribute cannot be changed.

**User response:** None.

---

**6027-2361 [E] *command* is enabled. File systems cannot be deleted.**

**Explanation:** When indefinite retention protection is enabled the file systems cannot be deleted.

**User response:** None.

---

**6027-2362 [E] The current value of *command* is *attrName*. No changes made.**

**Explanation:** The current value and the request value are the same. No changes made.

**User response:** None.

---

**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.

---

**6027-2576 [E] 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.

---

**6027-2600 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.

---

**6027-2601 Snapshot name *dirName* already exists.**

**Explanation:** by the **tscrsnapshot** command.

**User response:** Delete existing file/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:** Delete an existing snapshot, then issue the create snapshot command again.

---

**6027-2608** Cannot create new snapshot: state of fileset *filesetName* is inconsistent (*badState*).

**Explanation:** An operation on the cited fileset is incomplete.

**User response:** Complete pending fileset actions, then issue the create snapshot command again.

---

**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 in state *snapshotState*.

**Explanation:** The snapshot cannot be deleted while it is in the cited transition state because of an in-progress snapshot operation.

**User response:** Wait for the in-progress operation to complete and then reissue the 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 specified snapshot 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 [N]** 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 [N]** 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.

---

**6027-2623 [E]** Error deleting snapshot *snapshotName* 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.

---

**6027-2633** 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.

---

**6027-2646**    **Unable to sync all nodes while *quiesced*, *rc=returnCode***

**Explanation:** This message is issued by the **tscrsnapshot** command.

**User response:** Restart failing nodes or switches 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.

**User response:** Change the file system format version by issuing `mmchfs -V`.

---

**6027-2649** Fileset *filesetName* contains user files and cannot be deleted unless the `-f` option is specified.

**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 `mmdelfileset` 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. Specify `-f` to force unlink.

**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 *diskAddresses* disk addresses.

**Explanation:** Too many disk addresses specified.

**User response:** Provide less than 256 disk addresses to the command.

---

**6027-2667 [I]** Allowing block allocation for file system *fileSystem* that makes a file ill-replicated due to insufficient *resource* 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 [X]** 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 [X]** 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).

---

**6027-2677** Fileset *filesetName* has pending changes that need to be synced.

**Explanation:** A user is trying to change a caching option for a fileset while it has local changes that are not yet synced with the home server.

**User response:** Perform AFM recovery before reissuing the command.

---

**6027-2678** File system *fileSystem* is mounted on nodes *nodes* or fileset *filesetName* is not unlinked.

**Explanation:** A user is trying to change a caching feature for a fileset while the file system is still mounted or the fileset is still linked.

**User response:** Unmount the file system from all nodes or unlink the fileset before reissuing the command.

---

**6027-2679** Mount of *fileSystem* failed because mount event not handled by any data management application.

**Explanation:** The mount failed because the file system is enabled for DMAPI events (`-z yes`), but there was no data management application running to handle the event.

**User response:** Make sure the DM application (for example HSM or HPSS) is running before the file system is mounted.

---

**6027-2680** AFM filesets cannot be created for file system *fileSystem*.

**Explanation:** The current file system format version does not support AFM-enabled filesets; the `-p` option cannot be used.

**User response:** Change the file system format version by issuing `mmchfs -V`.

---

**6027-2681** Snapshot *snapshotName* has linked independent filesets

**Explanation:** The specified snapshot is not in a valid state.

**User response:** Correct the problem and reissue the command.

---

**6027-2682 [E]** Set quota file attribute error (*reasonCode*)*explanation*

**Explanation:** While mounting a file system a new quota file failed to be created due to inconsistency with the current degree of replication or the number of failure groups.

**User response:** Disable quotas. Check and correct the degree of replication and the number of failure groups. Re-enable quotas.

---

**6027-2683** Fileset *filesetName* in 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.

---

**6027-2684** File system *fileSystem* does not contain global snapshot *snapshotName*, *err* = *number*

**Explanation:** An incorrect snapshot name was specified.

**User response:** Select a valid snapshot and issue the command again.

---

**6027-2685** Total file system capacity allows *minMaxInodes* inodes in *fileSystem*. Currently the total inode limits used by all the inode spaces in *inodeSpace* is *inodeSpaceLimit*. There must be at least *number* inodes available to create a new inode space. Use the `mmlsfileset -L` command to show the maximum inode limits of each fileset. Try reducing the maximum inode limits for some of the inode spaces in *fileSystem*.

**Explanation:** The number of inodes available is too small to create a new inode space.

**User response:** Reduce the maximum inode limits and issue the command again.

---

**6027-2688** Only independent filesets can be configured as AFM filesets. The `--inode-space=new` option is required.

**Explanation:** Only independent filesets can be configured for caching.

**User response:** Specify the `--inode-space=new` option.

---

**6027-2689** The value for `--block-size` must be the keyword `auto` or the value must be of the form `[n]K`, `[n]M`, `[n]G` or `[n]T`, where `n` is an optional integer in the range 1 to 1023.

**Explanation:** An invalid value was specified with the `--block-size` option.

**User response:** Reissue the command with a valid option.

---

**6027-2690** Fileset *filesetName* can only be linked within its own inode space.

**Explanation:** A dependent fileset can only be linked within its own inode space.

**User response:** Correct the junction path and reissue the command.

---

**6027-2691** The `fastea` feature needs to be enabled for file system *fileSystem* before creating AFM filesets.

**Explanation:** The current file system on-disk format does not support storing of extended attributes in the file's inode. This is required for AFM-enabled filesets.

**User response:** Use the `mmmigratefs` command to enable the fast extended-attributes feature.

---

**6027-2692** Error encountered while processing the input file.

**Explanation:** The `tscrsnapshot` command encountered an error while processing the input file.

**User response:** Check and validate the fileset names listed in the input file.

---

**6027-2693** Fileset junction name *junctionName* conflicts with the current setting of `mmsnapdir`.

**Explanation:** The fileset junction name conflicts with the current setting of `mmsnapdir`.

**User response:** Select a new junction name or a new directory for the link and reissue the `mmlinkfileset` command.

---

**6027-2694 [I]** The requested maximum number of inodes is already at *number*.

**Explanation:** The specified number of nodes is already in effect.

**User response:** This is an informational message.

---



---

**6027-2695 [E]** The number of inodes to preallocate cannot be higher than the maximum number of inodes.

**Explanation:** The specified number of nodes to preallocate is not valid.

**User response:** Correct the `--inode-limit` argument then retry the command.

---

**6027-2696 [E]** The number of inodes to preallocate cannot be lower than the *number* inodes already allocated.

**Explanation:** The specified number of nodes to preallocate is not valid.

**User response:** Correct the `--inode-limit` argument then retry the command.

---

**6027-2697** Fileset at *junctionPath* has pending changes that need to be synced.

**Explanation:** A user is trying to change a caching option for a fileset while it has local changes that are not yet synced with the home server.

**User response:** Perform AFM recovery before reissuing the command.

---

**6027-2698** File system *fileSystem* is mounted on *nodes* `nodes` or fileset at *junctionPath* is not unlinked.

**Explanation:** A user is trying to change a caching feature for a fileset while the file system is still mounted or the fileset is still linked.

**User response:** Unmount the file system from all nodes or unlink the fileset before reissuing the command.

---

**6027-2699** Cannot create a new independent fileset until an existing one is deleted. File system *fileSystem* has a limit of *maxNumber* independent filesets.

**Explanation:** An attempt to create an independent fileset for the cited file system failed because it would exceed the cited limit.

**User response:** Remove unneeded independent filesets and reissue the command.

---

**6027-2700 [E]** 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 [I] Recovered *number* nodes.**

**Explanation:** The asynchronous part (phase 2) of node failure recovery has completed.

**User response:** None. Informational message only.

---

**6027-2707 [I] 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 [E] 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 [E] 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 [E] 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. *errorString*

**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. *errorString*

**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 [I]** 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: `GPFSMDPORTRANGE='range'`. Using default.

**Explanation:** The GPFS command port range format is `llll[-hhhhh]`, where `llll` 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 [E]** 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 [N]** 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 [I]** *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 [N]** 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 [I]** 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.

---

**6027-2727** *fileSystem*: quota manager is not available.

**Explanation:** An attempt was made to perform a quota command without a quota manager running. This could be caused by a conflicting offline `mmfsck` command.

**User response:** Reissue the command once the conflicting program has ended.

---

**6027-2728 [N]** Connection from *node* rejected because it does not support IPv6

**Explanation:** A connection request was received from a node that does not support Internet Protocol Version 6 (IPv6), and at least one node in the cluster is configured with an IPv6 address (not an IPv4-mapped one) as its primary address. Since the connecting node

will not be able to communicate with the IPv6 node, it is not permitted to join the cluster.

**User response:** Upgrade the connecting node to a version of GPFS that supports IPv6, or delete all nodes with IPv6-only addresses from the cluster.

---

**6027-2729** Value *value* for option *optionName* is out of range. Valid values are *value* through *value*.

**Explanation:** An out of range value was specified for the specified option.

**User response:** Correct the command line.

---

**6027-2730 [E]** Node *nodeName* failed to take over as cluster manager.

**Explanation:** An attempt to takeover as cluster manager failed.

**User response:** Make sure that GPFS is running on a sufficient number of quorum nodes.

---

**6027-2731** Failed to locate a working cluster manager.

**Explanation:** The cluster manager has failed or changed. The new cluster manager has not been appointed.

**User response:** Check the internode communication configuration and ensure enough GPFS nodes are up to make a quorum.

---

**6027-2732** Attention: No data disks remain in the system pool. Use `mmapplypolicy` to migrate all data left in the system pool to other storage pool.

**Explanation:** The `mmchdisk` command has been issued but no data disks remain in the system pool. Warn user to use `mmapplypolicy` to move data to other storage pool.

**User response:** None. Informational message only.

---

**6027-2733** The file system name (*fsname*) is longer than the maximum allowable length (*maxLength*).

**Explanation:** The file system name is invalid because it is longer than the maximum allowed length of 255 characters.

**User response:** Specify a file system name whose length is 255 characters or less and reissue the command.

---

**6027-2734 [E]** Disk failure from node *nodeName* Volume *name*. Physical volume *name*.

**Explanation:** An I/O request to a disk or a request to fence a disk has failed in such a manner that GPFS can no longer use the disk.

**User response:** Check the disk hardware and the software subsystems in the path to the disk.

---

**6027-2735 [E]** Not a manager

**Explanation:** This node is not a manager or no longer a manager of the type required to proceed with the operation. This could be caused by the change of manager in the middle of the operation.

**User response:** Retry the operation.

---

**6027-2736** The value for `--block-size` must be the keyword `auto` or the value must be of the form `nK`, `nM`, `nG` or `nT`, where `n` is an optional integer in the range 1 to 1023.

**Explanation:** An invalid value was specified with the `--block-size` option.

**User response:** Reissue the command with a valid option.

---

**6027-2738** Editing quota limits for the root user is not permitted

**Explanation:** The `root` user was specified for quota limits 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.

---

**6027-2739** Editing quota limits for *groupName* group not permitted.

**Explanation:** The `system` group was specified for quota limits 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.

---

**6027-2740 [I]** Starting new election as previous `clmgr` is expelled

**Explanation:** This node is taking over as `clmgr` without challenge as the old `clmgr` is being expelled.

**User response:** None. Informational message only.

---

**6027-2741 [W] This node can not continue to be cluster manager**

**Explanation:** This node invoked the user-specified callback handler for event **tiebreakerCheck** and it returned a non-zero value. This node cannot continue to be the cluster manager.

**User response:** None. Informational message only.

---

**6027-2742 [I] CallExitScript: exit script *exitScript* on event *eventName* returned code *returnCode*, quorumloss.**

**Explanation:** This node invoked the user-specified callback handler for the **tiebreakerCheck** event and it returned a non-zero value. The user-specified action with the error is **quorumloss**.

**User response:** None. Informational message only.

---

**6027-2743 Permission denied.**

**Explanation:** The command is invoked by an unauthorized user.

**User response:** Retry the command with an authorized user.

---

**6027-2744 [D] Invoking tiebreaker callback script**

**Explanation:** The node is invoking the callback script due to change in quorum membership.

**User response:** None. Informational message only.

---

**6027-2745 [E] File system is not mounted.**

**Explanation:** A command was issued, which requires that the file system be mounted.

**User response:** Mount the file system and reissue the command.

---

**6027-2746 [E] Too many disks unavailable for this server to continue serving a RecoveryGroup.**

**Explanation:** RecoveryGroup panic: Too many disks unavailable to continue serving this RecoveryGroup. This server will resign, and failover to an alternate server will be attempted.

**User response:** Ensure the alternate server took over. Determine what caused this event and address the situation. Prior messages may help determine the cause of the event.

---

**6027-2747 [E] Inconsistency detected between the local node number retrieved from 'mmsdrfs' (*nodeNumber*) and the node number retrieved from 'mmfs.cfg' (*nodeNumber*).**

**Explanation:** The node number retrieved by obtaining the list of nodes in the **mmsdrfs** file did not match the node number contained in **mmfs.cfg**. There may have been a recent change in the IP addresses being used by network interfaces configured at the node.

**User response:** Stop and restart GPFS daemon.

---

**6027-2748 Terminating because a conflicting program on the same inode space *inodeSpace* is running.**

**Explanation:** A program detected that it must terminate because a conflicting program is running.

**User response:** Reissue the command after the conflicting program ends.

---

**6027-2749 Specified locality group '*number*' does not match disk '*name*' locality group '*number*'. To change locality groups in an SNC environment, please use the **mmeldisk** and **mmaddisk** commands.**

**Explanation:** The locality group specified on the **mmchdisk** command does not match the current locality group of the disk.

**User response:** To change locality groups in an SNC environment, use the **mmeldisk** and **mmaddisk** commands.

---

**6027-2750 [I] Node *nodeName* is now the Group Leader.**

**Explanation:** A new cluster Group Leader has been assigned.

**User response:** None. Informational message only.

---

**6027-2751 [I] Starting new election: Last elected: *NodeNumber* Sequence: *SequenceNumber***

**Explanation:** A new disk election will be started. The disk challenge will be skipped since the last elected node was either none or the local node.

**User response:** None. Informational message only.

---

**6027-2752 [I] This node got elected. Sequence: *SequenceNumber***

**Explanation:** Local node got elected in the disk election. This node will become the cluster manager.

**User response:** None. Informational message only.

---

---

**6027-2753 [N] Responding to disk challenge:**  
**response:** *ResponseValue*. **Error code:**  
*ErrorCode*.

**Explanation:** A disk challenge has been received, indicating that another node is attempting to become a Cluster Manager. Issuing a challenge response, to confirm the local node is still alive and will remain the Cluster Manager.

**User response:** None. Informational message only.

---

**6027-2754 [X] Challenge thread did not respond to challenge in time: took *TimeIntervalSecs* seconds.**

**Explanation:** Challenge thread took too long to respond to a disk challenge. Challenge thread will exit, which will result in the local node losing quorum.

**User response:** None. Informational message only.

---

**6027-2755 [N] Another node committed disk election with sequence *CommittedSequenceNumber* (our sequence was *OurSequenceNumber*).**

**Explanation:** Another node committed a disk election with a sequence number higher than the one used when this node used to commit an election in the past. This means that the other node has become, or is becoming, a Cluster Manager. To avoid having two Cluster Managers, this node will lose quorum.

**User response:** None. Informational message only.

---

**6027-2756 Attention: In file system *FileSystemName*, *FileSetName* (Default) *QuotaLimitType*(*QuotaLimit*) for *QuotaTypeUserName/GroupName/FileSetName* is too small. Suggest setting it higher than *minQuotaLimit*.**

**Explanation:** Users set too low quota limits. It will cause unexpected quota behavior. *MinQuotaLimit* is computed through:

1. for block:  $QUOTA\_THRESHOLD * MIN\_SHARE\_BLOCKS * subblocksize$
2. for inode:  $QUOTA\_THRESHOLD * MIN\_SHARE\_INODES$

**User response:** Users should reset quota limits so that they are more than *MinQuotaLimit*. It is just a warning. Quota limits will be set anyway.

---

**6027-2757 [E] The peer snapshot is in progress. Queue cannot be flushed now.**

**Explanation:** The Peer Snapshot is in progress. Queue cannot be flushed now.

**User response:** Reissue the command once the peer snapshot has ended.

---

**6027-2758 [E] The AFM target does not support this operation. Run *mmafmconfig* on the AFM target cluster.**

**Explanation:** The *.afmctl* file is probably not present on the AFM target cluster.

**User response:** Run *mmafmconfig* on the AFM target cluster to configure the AFM target cluster.

---

**6027-2759 [N] Disk lease period expired in cluster *ClusterName*. Attempting to reacquire lease.**

**Explanation:** The disk lease period expired, which will prevent the local node from being able to perform disk I/O. This can be caused by a temporary communication outage.

**User response:** If message is repeated then the communication outage should be investigated.

---

**6027-2760 [N] Disk lease reacquired in cluster *ClusterName*.**

**Explanation:** The disk lease has been reacquired, and disk I/O will be resumed.

**User response:** None. Informational message only.

---

**6027-2761 Unable to run *command* on '*fileSystem*' while the file system is mounted in restricted mode.**

**Explanation:** A command that can alter 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.

---

**6027-2762 Unable to run *command* on '*fileSystem*' 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.

---

**6027-2763 Unable to start *command* on '*fileSystem*' 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.

---

---

**6027-2764** Terminating *command* on *fileSystem* because a conflicting program *name* is running.

**Explanation:** A program detected that it must terminate because a conflicting program is running.

**User response:** Reissue the command after the conflicting program ends.

---

**6027-2765** *command* on '*fileSystem*' is finished waiting. Processing continues ... *name*

**Explanation:** A program detected that it can now continue the processing since a conflicting program has ended.

**User response:** None. Informational message only.

---

**6027-2766 [I]** User script has chosen to expel node *nodeName* instead of node *nodeName*.

**Explanation:** User has specified a callback script that is invoked whenever a decision is about to be taken on what node should be expelled from the active cluster. As a result of the execution of the script, GPFS will reverse its decision on what node to expel.

**User response:** None.

---

**6027-2767 [E]** Error *errorNumber* while accessing tiebreaker devices.

**Explanation:** An error was encountered while reading from or writing to the tiebreaker devices. When such error happens while the cluster manager is checking for challenges, it will cause the cluster manager to lose cluster membership.

**User response:** Verify the health of the tiebreaker devices.

---

**6027-2770** Disk *diskName* belongs to a write-affinity enabled storage pool. Its failure group cannot be changed.

**Explanation:** The failure group specified on the **mmchdisk** command does not match the current failure group of the disk.

**User response:** Use the **mmdeldisk** and **mmaddisk** commands to change failure groups in a write-affinity enabled storage pool.

---

**6027-2771** *fileSystem*: Default per-fileset quotas are disabled for *quotaType*.

**Explanation:** A command was issued to modify default fileset-level quota, but default quotas are not enabled.

**User response:** Ensure the **--perfileset-quota** option is in effect for the file system, then use the

**mmdefquotaon** command to enable default fileset-level quotas. After default quotas are enabled, issue the failed command again.

---

**6027-2772** Cannot close disk *name*.

**Explanation:** Could not access the specified disk.

**User response:** Check the disk hardware and the path to the disk. Refer to "Unable to access disks" on page 131.

---

**6027-2773** *fileSystem*:filesetName: default quota for *quotaType* is disabled.

**Explanation:** A command was issued to modify default quota, but default quota is not enabled.

**User response:** Ensure the **-Q yes** option is in effect for the file system, then enable default quota with the **mmdefquotaon** command.

---

**6027-2774** *fileSystem*: Per-fileset quotas are not enabled.

**Explanation:** A command was issued to modify fileset-level quota, but per-fileset quota management is not enabled.

**User response:** Ensure that the **--perfileset-quota** option is in effect for the file system and reissue the command.

---

**6027-2775** Storage pool named *poolName* does not exist.

**Explanation:** The **mmlspool** command was issued, but the specified storage pool does not exist.

**User response:** Correct the input and reissue the command.

---

**6027-2776** Attention: A disk being stopped reduces the degree of system metadata replication (*value*) or data replication (*value*) to lower than tolerable.

**Explanation:** The **mmchdisk stop** command was issued, but the disk cannot be stopped because of the current file system metadata and data replication factors.

**User response:** Make more disks available, delete unavailable disks, or change the file system metadata replication factor. Also check the current value of the **unmountOnDiskFail** configuration parameter.

---

**6027-2777 [E]** Node *nodeName* is being expelled because of an expired lease. Pings sent: *pingsSent*. Replies received: *pingRepliesReceived*.

**Explanation:** The node listed did not renew its lease

in a timely fashion and is being expelled from the cluster.

**User response:** Check the network connection between this node and the node listed in the message.

**6027-2778 [I] Node *nodeName*: ping timed out. Pings sent: *pingsSent*. Replies received: *pingRepliesReceived*.**

**Explanation:** Ping timed out for the node listed, which should be the cluster manager. A new cluster manager will be chosen while the current cluster manager is expelled from the cluster.

**User response:** Check the network connection between this node and the node listed in the message.

**6027-2779 [E] Challenge thread stopped.**

**Explanation:** A tiebreaker challenge thread stopped because of an error. Cluster membership will be lost.

**User response:** Check for additional error messages. File systems will be unmounted, then the node will rejoin the cluster.

**6027-2780 [E] Not enough quorum nodes reachable: *reachableNodes*.**

**Explanation:** The cluster manager cannot reach a sufficient number of quorum nodes, and therefore must resign to prevent cluster partitioning.

**User response:** Determine if there is a network outage or if too many nodes have failed.

**6027-2781 [E] Lease expired for *numSecs* seconds (shutdownOnLeaseExpiry).**

**Explanation:** Disk lease expired for too long, which results in the node losing cluster membership.

**User response:** None. The node will attempt to rejoin the cluster.

**6027-2782 [E] This node is being expelled from the cluster.**

**Explanation:** This node received a message instructing it to leave the cluster, which might indicate communication problems between this node and some other node in the cluster.

**User response:** None. The node will attempt to rejoin the cluster.

**6027-2783 [E] New leader elected with a higher ballot number.**

**Explanation:** A new group leader was elected with a higher ballot number, and this node is no longer the

leader. Therefore, this node must leave the cluster and rejoin.

**User response:** None. The node will attempt to rejoin the cluster.

**6027-2784 [E] No longer a cluster manager or lost quorum while running a group protocol.**

**Explanation:** Cluster manager no longer maintains quorum after attempting to run a group protocol, which might indicate a network outage or node failures.

**User response:** None. The node will attempt to rejoin the cluster.

**6027-2785 [X] A severe error was encountered during cluster probe.**

**Explanation:** A severe error was encountered while running the cluster probe to determine the state of the nodes in the cluster.

**User response:** Examine additional error messages. The node will attempt to rejoin the cluster.

**6027-2786 [E] Unable to contact any quorum nodes during cluster probe.**

**Explanation:** This node has been unable to contact any quorum nodes during cluster probe, which might indicate a network outage or too many quorum node failures.

**User response:** Determine whether there was a network outage or whether quorum nodes failed.

**6027-2787 [E] Unable to contact enough other quorum nodes during cluster probe.**

**Explanation:** This node, a quorum node, was unable to contact a sufficient number of quorum nodes during cluster probe, which might indicate a network outage or too many quorum node failures.

**User response:** Determine whether there was a network outage or whether quorum nodes failed.

**6027-2788 [E] Attempt to run leader election failed with error *errorNumber*.**

**Explanation:** This node attempted to run a group leader election but failed to get elected. This failure might indicate that two or more quorum nodes attempted to run the election at the same time. As a result, this node will lose cluster membership and then attempt to rejoin the cluster.

**User response:** None. The node will attempt to rejoin the cluster.

---

**6027-2789 [E] Tiebreaker script returned a non-zero value.**

**Explanation:** The tiebreaker script, invoked during group leader election, returned a non-zero value, which results in the node losing cluster membership and then attempting to rejoin the cluster.

**User response:** None. The node will attempt to rejoin the cluster.

---

**6027-2790 Attention: Disk parameters were changed. Use the mmrestripefs command with the -r option to relocate data and metadata.**

**Explanation:** The mmchdisk command with the change option was issued.

**User response:** Issue the mmrestripefs -r command to relocate data and metadata.

---

**6027-2791 Disk *diskName* does not belong to file system *deviceName*.**

**Explanation:** The input disk name does not belong to the specified file system.

**User response:** Correct the command line.

---

**6027-2792 The current file system version does not support default per-fileset quotas.**

**Explanation:** The current version of the file system does not support default fileset-level quotas.

**User response:** Use the mmchfs -V command to activate the new function.

---

**6027-2793 [E] Contents of local *fileName* file are invalid. Node may be unable to be elected group leader.**

**Explanation:** In an environment where tie-breaker disks are used, the contents of the ballot file have become invalid, possibly because the file has been overwritten by another application. This node will be unable to be elected group leader.

**User response:** Run mmcommon resetTiebreaker, which will ensure the GPFS daemon is down on all quorum nodes and then remove the given file on this node. After that, restart the cluster on this and on the other nodes.

---

**6027-2794 [E] Invalid content of disk paxos sector for disk *diskName*.**

**Explanation:** In an environment where tie-breaker disks are used, the contents of either one of the tie-breaker disks or the ballot files became invalid, possibly because the file has been overwritten by another application.

---

**User response:** Examine mmfs.log file on all quorum nodes for indication of a corrupted ballot file. If 6027-2793 is found then follow instructions for that message. If problem cannot be resolved, shut down GPFS across the cluster, undefine, and then redefine the tiebreakerdisks configuration variable, and finally restart the cluster.

---

**6027-2795 An error occurred while executing 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.

---

**6027-2796 [W] Callback event *eventName* is not supported on this node; processing continues ...**

**Explanation:** informational

**User response:**

---

**6027-2797 [I] Node *nodeName*: lease request received late. Pings sent: *pingsSent*. Maximum pings missed: *maxPingsMissed*.**

**Explanation:** The cluster manager reports that the lease request from the given node was received late, possibly indicating a network outage.

**User response:** Check the network connection between this node and the node listed in the message.

---

**6027-2798 [E] The node *nodeName* does not have a valid Extended License to run the requested command.**

**Explanation:** The file system manager node does not have a valid extended license to run ILM, AFM, or CNFS commands.

**User response:** Make sure gpfs.ext package is installed correctly on file system manager node and try again.

---

**6027-2799 Option '*option*' is incompatible with option '*option*'.**

**Explanation:** The options specified on the command are incompatible.

**User response:** Do not specify these two options together.

---

**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 [I] Loaded policy '*policyFileName* or *filesystemName*': *summaryOfPolicyRules***

**Explanation:** The specified loaded policy has the specified policy rules.

**User response:** None. Informational message only.

---

**6027-2806 [E] Error while validating policy '*policyFileName* or *filesystemName*': *rc=errorCode: errorDetailsString***

**Explanation:** An error occurred while validating the specified 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 [W] 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 *filesystemName*': *summaryOfPolicyRules***

**Explanation:** The specified validated policy has the specified policy rules.

**User response:** None. Informational message only.

---

**6027-2810 [W] 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.

---

**6027-2821** Rule '*ruleName*' (*ruleNumber*) specifies a **THRESHOLD** for **EXTERNAL POOL** '*externalPoolName*'. This is not supported.

**Explanation:** GPFS does not support the **THRESHOLD** clause within a migrate rule that names an external pool in the **FROM POOL** clause.

**User response:** Correct or remove the rule.

---

---

**6027-2822** This file system does not support fast extended attributes, which are needed for encryption.

**Explanation:** Fast extended attributes need to be supported by the file system for encryption to be activated.

**User response:** Enable the fast extended attributes feature in this file system.

---

**6027-2823 [E]** Encryption activated in the file system, but node not enabled for encryption.

**Explanation:** The file system is enabled for encryption, but this node is not.

**User response:** Ensure the GPFS encryption packages are installed. Verify if encryption is supported on this node architecture.

---

**6027-2824** This file system version does not support encryption rules.

**Explanation:** This file system version does not support encryption.

**User response:** Update the file system to a version which supports encryption.

---

**6027-2825** Duplicate encryption set name '*setName*'.

**Explanation:** The given set name is duplicated in the policy file.

**User response:** Ensure each set name appears only once in the policy file.

---

**6027-2826** The encryption set '*setName*' requested by rule '*rule*' could not be found.

**Explanation:** The given set name used in the rule cannot be found.

**User response:** Verify if the set name is correct. Add the given set if it is missing from the policy.

---

**6027-2827 [E]** Error in evaluation of encryption policy for file *fileName: %s*

**Explanation:** An error occurred while evaluating the encryption rules in the given policy file.

**User response:** Examine the other error messages produced while evaluating the policy file.

---

**6027-2828 [E]** Encryption not supported on Windows. Encrypted file systems are not allowed when Windows nodes are present in the cluster.

**Explanation:** Self-explanatory.

**User response:** To activate encryption, ensure there are

no Windows nodes in the cluster.

---

**6027-2950 [E]** Trace value '*value*' after class '*class*' must be from 0 to 14.

**Explanation:** The specified trace value is not recognized.

**User response:** Specify a valid trace integer value.

---

**6027-2951** Value *value* for `worker1Threads` must be  $\leq$  than the original setting *value*

**Explanation:** An attempt to dynamically set `worker1Threads` found the value out of range. The dynamic value must be  $2 \leq \textit{value} \leq$  the original setting when the GPFS daemon was started.

---

**6027-2952 [E]** Unknown assert class '*assertClass*'.

**Explanation:** The assert class is not recognized.

**User response:** Specify a valid assert class.

---

**6027-2953 [E]** Non-numeric assert value '*value*' after class '*class*'.

**Explanation:** The specified assert value is not recognized.

**User response:** Specify a valid assert integer value.

---

**6027-2954 [E]** Assert value '*value*' after class '*class*' must be from 0 to 127.

**Explanation:** The specified assert value is not recognized.

**User response:** Specify a valid assert integer value.

---

**6027-2955 [W]** Time-of-day may have jumped back. Late by *delaySeconds* seconds to wake certain threads.

**Explanation:** Time-of-day may have jumped back, which has resulted in some threads being awakened later than expected. It is also possible that some other factor has caused a delay in waking up the threads.

**User response:** Verify if there is any problem with network time synchronization, or if time-of-day is being incorrectly set.

---

**6027-2956 [E]** Invalid crypto engine type (`encryptionCryptoEngineType`): *cryptoEngineType*.

**Explanation:** The specified value for `encryptionCryptoEngineType` is incorrect.

**User response:** Specify a valid value for `encryptionCryptoEngineType`.

---

**6027-2957 [E] Invalid cluster manager selection choice (clusterManagerSelection):** *clusterManagerSelection*.

**Explanation:** The specified value for **clusterManagerSelection** is incorrect.

**User response:** Specify a valid value for **clusterManagerSelection**.

---

**6027-2958 [E] Invalid NIST compliance type (nistCompliance):** *nistComplianceValue*.

**Explanation:** The specified value for **nistCompliance** is incorrect.

**User response:** Specify a valid value for **nistCompliance**.

---

**6027-2959 [E] The CPU architecture on this node does not support tracing in traceMode mode. Switching to traceMode mode.**

**Explanation:** The CPU does not have constant time stamp counter capability required for overwrite trace mode. The trace has been enabled in blocking mode.

**User response:** Update configuration parameters to use trace facility in blocking mode or replace this node with modern CPU architecture.

---

**6027-3101 Pdisk rotation rate invalid in option 'option'.**

**Explanation:** When parsing disk lists, the pdisk rotation rate is not valid.

**User response:** Specify a valid rotation rate (SSD, NVRAM, or 1025 through 65535).

---

**6027-3102 Pdisk FRU number too long in option 'option', maximum length length.**

**Explanation:** When parsing disk lists, the pdisk FRU number is too long.

**User response:** Specify a valid FRU number that is shorter than or equal to the maximum length.

---

**6027-3103 Pdisk location too long in option 'option', maximum length length.**

**Explanation:** When parsing disk lists, the pdisk location is too long.

**User response:** Specify a valid location that is shorter than or equal to the maximum length.

---

**6027-3105 Pdisk nPathActive invalid in option 'option'.**

**Explanation:** When parsing disk lists, the **nPathActive** value is not valid.

**User response:** Specify a valid **nPathActive** value (0 to 255).

---

**6027-3106 Pdisk nPathTotal invalid in option 'option'.**

**Explanation:** When parsing disk lists, the **nPathTotal** value is not valid.

**User response:** Specify a valid *nPathTotal* value (0 to 255).

---

**6027-3107 Pdisk nsdFormatVersion invalid in option 'name1name2'.**

**Explanation:** The **nsdFormatVersion** that is entered while parsing the disk is invalid.

**User response:** Specify valid **nsdFormatVersion**, 1 or 2.

---

**6027-3200 AFM ERROR: command pCacheCmd fileset filesetName fileids [parentId.childId.tParentId.targetId,ReqCmd] original error oerr application error aerr remote error remoteError**

**Explanation:** AFM operations on a particular file failed.

**User response:** For asynchronous operations that are queued, run the **mmafmctl** command with the **resumeQueued** option after fixing the problem at the home cluster.

---

**6027-3201 AFM ERROR DETAILS: type: remoteCmdType snapshot name snapshotName snapshot ID snapshotId**

**Explanation:** Peer snapshot creation or deletion failed.

**User response:** Fix snapshot creation or deletion error.

---

**6027-3204 AFM: Failed to set xattr on inode inodeNum error err, ignoring.**

**Explanation:** Setting extended attributes on an inode failed.

**User response:** None.

---

**6027-3205 AFM: Failed to get xattrs for inode inodeNum, ignoring.**

**Explanation:** Getting extended attributes on an inode failed.

**User response:** None.

---

**6027-3209** Home NFS mount of *host:path* failed with error *err*

**Explanation:** NFS mounting of path from the home cluster failed.

**User response:** Make sure the exported path can be mounted over NFSv3.

---

**6027-3210** Cannot find AFM control file for fileset *filesetName* in the exported file system at home. ACLs and extended attributes will not be synchronized. Sparse files will have zeros written for holes.

**Explanation:** Either home path does not belong to GPFS, or the AFM control file is not present in the exported path.

**User response:** If the exported path belongs to a GPFS file system, run the `mmafmconfig` command with the `enable` option on the export path at home.

---

**6027-3211** Change in home export detected. Caching will be disabled.

**Explanation:** A change in home export was detected or the home path is stale.

**User response:** Ensure the exported path is accessible.

---

**6027-3212** AFM ERROR: Cannot enable AFM for fileset *filesetName* (error *err*)

**Explanation:** AFM was not enabled for the fileset because the root file handle was modified, or the remote path is stale.

**User response:** Ensure the remote export path is accessible for NFS mount.

---

**6027-3213** Cannot find snapshot link directory name for exported file system at home for fileset *filesetName*. Snapshot directory at home will be cached.

**Explanation:** Unable to determine the snapshot directory at the home cluster.

**User response:** None.

---

**6027-3214 [E]** AFM: Unexpiration of fileset *filesetName* failed with error *err*. Use `mmafmctl` to manually unexpire the fileset.

**Explanation:** Unexpiration of fileset failed after a home reconnect.

**User response:** Run the `mmafmctl` command with the `unexpire` option on the fileset.

---

**6027-3215 [W]** AFM: Peer snapshot delayed due to long running execution of operation to remote cluster for fileset *filesetName*. Peer snapshot continuing to wait.

**Explanation:** Peer snapshot command timed out waiting to flush messages.

**User response:** None.

---

**6027-3216** Fileset *filesetName* encountered an error synchronizing with the remote cluster. Cannot synchronize with the remote cluster until AFM recovery is executed.

**Explanation:** Cache failed to synchronize with home because of an out of memory or conflict error. Recovery, resynchronization, or both will be performed by GPFS to synchronize cache with the home.

**User response:** None.

---

**6027-3217** AFM ERROR Unable to unmount NFS export for fileset *filesetName*

**Explanation:** NFS unmount of the path failed.

**User response:** None.

---

**6027-3220** AFM: Home NFS mount of *host:path* failed with error *err* for file system *fileSystem* fileset id *filesetName*. Caching will be disabled and the mount will be tried again after *mountRetryTime* seconds, on next request to gateway

**Explanation:** NFS mount of the home cluster failed. The mount will be tried again after *mountRetryTime* seconds.

**User response:** Make sure the exported path can be mounted over NFSv3.

---

**6027-3221** AFM: Home NFS mount of *host:path* succeeded for file system *fileSystem* fileset *filesetName*. Caching is enabled.

**Explanation:** NFS mount of the path from the home cluster succeeded. Caching is enabled.

**User response:** None.

---

**6027-3224 [I]** AFM: Failed to set extended attributes on file system *fileSystem* inode *inodeNum* error *err*, ignoring.

**Explanation:** Setting extended attributes on an inode failed.

**User response:** None.

---

---

**6027-3225 [I] AFM: Failed to get extended attributes for file system *fileSystem* inode *inodeNum*, ignoring.**

**Explanation:** Getting extended attributes on an inode failed.

**User response:** None.

---

**6027-3226 [I] AFM: Cannot find control file for file system *fileSystem* fileset *filesetName* in the exported file system at home. ACLs and extended attributes will not be synchronized. Sparse files will have zeros written for holes.**

**Explanation:** Either the home path does not belong to GPFS, or the AFM control file is not present in the exported path.

**User response:** If the exported path belongs to a GPFS file system, run the **mmafmconfig** command with the **enable** option on the export path at home.

---

**6027-3227 [E] AFM: Cannot enable AFM for file system *fileSystem* fileset *filesetName* (error *err*)**

**Explanation:** AFM was not enabled for the fileset because the root file handle was modified, or the remote path is stale.

**User response:** Ensure the remote export path is accessible for NFS mount.

---

**6027-3228 [E] AFM: Unable to unmount NFS export for file system *fileSystem* fileset *filesetName***

**Explanation:** NFS unmount of the path failed.

**User response:** None.

---

**6027-3229 [E] AFM: File system *fileSystem* fileset *filesetName* encountered an error synchronizing with the remote cluster. Cannot synchronize with the remote cluster until AFM recovery is executed.**

**Explanation:** The cache failed to synchronize with home because of an out of memory or conflict error. Recovery, resynchronization, or both will be performed by GPFS to synchronize the cache with the home.

**User response:** None.

---

**6027-3230 [I] AFM: Cannot find snapshot link directory name for exported file system at home for file system *fileSystem* fileset *filesetName*. Snapshot directory at home will be cached.**

**Explanation:** Unable to determine the snapshot

directory at the home cluster.

**User response:** None.

---

**6027-3232 type AFM: *pCacheCmd* file system *fileSystem* fileset *filesetName* file IDs [*parentId.childId.tParentId.targetId,flag*] name *sourceName* origin error *err***

**Explanation:** AFM operations on a particular file failed.

**User response:** For asynchronous operations that are requested, run the **mmafmctl** command with the **resumeRequeued** option after fixing the problem at the home cluster.

---

**6027-3233 [I] AFM: Previous error repeated *repeatNum* times.**

**Explanation:** Multiple AFM operations have failed.

**User response:** None.

---

**6027-3234 [E] AFM: Unable to start thread to unexpire filesets.**

**Explanation:** Failed to start thread for unexpiration of fileset.

**User response:** None.

---

**6027-3235 [I] AFM: Stopping recovery for the file system *fileSystem* fileset *filesetName***

**Explanation:** AFM recovery terminated because the current node is no longer MDS for the fileset.

**User response:** None.

---

**6027-3236 [E] AFM: Recovery on file system *fileSystem* fileset *filesetName* failed with error *err*. Recovery will be retried on next access after recovery retry interval (*timeout* seconds) or manually resolve known problems and recover the fileset.**

**Explanation:** Recovery failed to complete on the fileset. The next access will restart recovery.

**Explanation:** AFM recovery failed. Fileset will be temporarily put into dropped state and will be recovered on accessing fileset after timeout mentioned in the error message. User can recover the fileset manually by running **mmafmctl** command with **recover** option after rectifying any known errors leading to failure.

**User response:** None.

---

---

**6027-3239 [E] AFM: Remote command *remoteCmdType* on file system *fileSystem* snapshot *snapshotName* snapshot ID *snapshotId* failed.**

**Explanation:** A failure occurred when creating or deleting a peer snapshot.

**User response:** Examine the error details and retry the operation.

---

**6027-3240 [E] AFM: *pCacheCmd* file system *fileSystem* fileset *filesetName* file IDs [*parentId.childId.tParentId.targetId.flag*] error *err***

**Explanation:** Operation failed to execute on home in independent-writer mode.

**User response:** None.

---

**6027-3241 [I] AFM: GW queue transfer started for file system *fileSystem* fileset *filesetName*. Transferring to *nodeAddress*.**

**Explanation:** An old GW initiated the queue transfer because a new GW node joined the cluster, and the fileset now belongs to the new GW node.

**User response:** None.

---

**6027-3242 [I] AFM: GW queue transfer started for file system *fileSystem* fileset *filesetName*. Receiving from *nodeAddress*.**

**Explanation:** An old MDS initiated the queue transfer because this node joined the cluster as GW and the fileset now belongs to this node.

**User response:** None.

---

**6027-3243 [I] AFM: GW queue transfer completed for file system *fileSystem* fileset *filesetName*. error *error***

**Explanation:** A GW queue transfer completed.

**User response:** None.

---

**6027-3244 [I] AFM: Home mount of *afmTarget* succeeded for file system *fileSystem* fileset *filesetName*. Caching is enabled.**

**Explanation:** A mount of the path from the home cluster succeeded. Caching is enabled.

**User response:** None.

---

**6027-3245 [E] AFM: Home mount of *afmTarget* failed with error *error* for file system *fileSystem* fileset ID *filesetName*. Caching will be disabled and the mount will be tried again after *mountRetryTime* seconds, on the next request to the gateway.**

**Explanation:** A mount of the home cluster failed. The mount will be tried again after *mountRetryTime* seconds.

**User response:** Verify that the *afmTarget* can be mounted using the specified protocol.

---

**6027-3246 [I] AFM: Prefetch recovery started for the file system *fileSystem* fileset *filesetName*.**

**Explanation:** Prefetch recovery started.

**User response:** None.

---

**6027-3247 [I] AFM: Prefetch recovery completed for the file system *fileSystem* fileset *filesetName*. error *error***

**Explanation:** Prefetch recovery completed.

**User response:** None.

---

**6027-3248 [E] AFM: Cannot find the control file for fileset *filesetName* in the exported file system at home. This file is required to operate in primary mode. The fileset will be disabled.**

**Explanation:** Either the home path does not belong to GPFS, or the AFM control file is not present in the exported path.

**User response:** If the exported path belongs to a GPFS file system, run the `mmafmconfig` command with the `enable` option on the export path at home.

---

**6027-3249 [E] AFM: Target for fileset *filesetName* is not a secondary-mode fileset or file system. This is required to operate in primary mode. The fileset will be disabled.**

**Explanation:** The AFM target is not a secondary fileset or file system.

**User response:** The AFM target fileset or file system should be converted to secondary mode.

---

**6027-3250 [E] AFM: Refresh intervals cannot be set for fileset.**

**Explanation:** Refresh intervals are not supported on primary and secondary-mode filesets.

**User response:** None.

---

---

**6027-3252 [I] AFM: Home has been restored for cache *filesetName*. Synchronization with home will be resumed.**

**Explanation:** A change in home export was detected that caused the home to be restored. Synchronization with home will be resumed.

**User response:** None.

---

**6027-3253 [E] AFM: Change in home is detected for cache *filesetName*. Synchronization with home is suspended until the problem is resolved.**

**Explanation:** A change in home export was detected or the home path is stale.

**User response:** Ensure the exported path is accessible.

---

**6027-3254 [W] AFM: Home is taking longer than expected to respond for cache *filesetName*. Synchronization with home is temporarily suspended.**

**Explanation:** A pending message from gateway node to home is taking longer than expected to respond. This could be the result of a network issue or a problem at the home site.

**User response:** Ensure the exported path is accessible.

---

**6027-3255 [E] AFM: Target for fileset *filesetName* is a secondary-mode fileset or file system. Only a primary-mode, read-only or local-update mode fileset can operate on a secondary-mode fileset. The fileset will be disabled.**

**Explanation:** The AFM target is a secondary fileset or file system. Only a primary-mode, read-only, or local-update fileset can operate on a secondary-mode fileset.

**User response:** Use a secondary-mode fileset as the target for the primary-mode, read-only or local-update mode fileset.

---

**6027-3256 [I] AFM: The RPO peer snapshot was missed for file system *fileSystem* fileset *filesetName*.**

**Explanation:** The periodic RPO peer snapshot was not taken in time for the primary fileset.

**User response:** None.

---

**6027-3257 [E] AFM: Unable to start thread to verify primary filesets for RPO.**

**Explanation:** Failed to start thread for verification of primary filesets for RPO.

**User response:** None.

---

**6027-3300 Attribute *afmShowHomeSnapshot* cannot be changed for a single-writer fileset.**

**Explanation:** Changing *afmShowHomeSnapshot* is not supported for single-writer filesets.

**User response:** None.

---

**6027-3301 Unable to quiesce all nodes; some processes are busy or holding required resources.**

**Explanation:** A timeout occurred on one or more nodes while trying to quiesce the file system during a snapshot command.

**User response:** Check the GPFS log on the file system manager node.

---

**6027-3302 Attribute *afmShowHomeSnapshot* cannot be changed for a *afmMode* fileset.**

**Explanation:** Changing *afmShowHomeSnapshot* is not supported for single-writer or independent-writer filesets.

**User response:** None.

---

**6027-3303 Cannot restore snapshot; quota management is active for *fileSystem*.**

**Explanation:** File system quota management is still active. The file system must be unmounted when restoring global snapshots.

**User response:** Unmount the file system and reissue the restore command.

---

**6027-3304 Attention: Disk space reclaim on *number* of *number* regions in *fileSystem* returned errors.**

**Explanation:** Free disk space reclaims on some regions failed during *tsreclaim* run. Typically this is due to the lack of space reclaim support by the disk controller or operating system. It may also be due to utilities such as *mmdefragfs* or *mmfsck* running concurrently.

**User response:** Verify that the disk controllers and the operating systems in the cluster support thin-provisioning space reclaim. Or, rerun the *mmfsctl reclaimSpace* command after *mmdefragfs* or *mmfsck* completes.

---

**6027-3305 AFM Fileset *filesetName* cannot be changed as it is in *beingDeleted* state**

**Explanation:** The user specified a fileset to *tchfileset* that cannot be changed.



**User response:** None. You cannot change the attributes of the root fileset.

---

**6027-3306 Fileset cannot be changed because it is unlinked.**

**Explanation:** The fileset cannot be changed when it is unlinked.

**User response:** Link the fileset and then try the operation again.

---

**6027-3307 Fileset cannot be changed.**

**Explanation:** Fileset cannot be changed.

**User response:** None.

---

**6027-3308 This AFM option cannot be set for a secondary fileset.**

**Explanation:** This AFM option cannot be set for a secondary fileset. The fileset cannot be changed.

**User response:** None.

---

**6027-3309 The AFM attribute specified cannot be set for a primary fileset.**

**Explanation:** This AFM option cannot be set for a primary fileset. The fileset cannot be changed.

**User response:** None.

---

**6027-3310 A secondary fileset cannot be changed.**

**Explanation:** A secondary fileset cannot be changed.

**User response:** None.

---

**6027-3311 A primary fileset cannot be changed.**

**Explanation:** A primary fileset cannot be changed.

**User response:** None.

---

**6027-3312 No inode was found matching the criteria.**

**Explanation:** No inode was found matching the criteria.

**User response:** None.

---

**6027-3313 File system scan RESTARTED due to resume of all disks being emptied.**

**Explanation:** The parallel inode traversal (PIT) phase is restarted with a file system restripe.

**User response:** None.

---

**6027-3314 File system scan RESTARTED due to new disks to be emptied.**

**Explanation:** The file system restripe was restarted after a new disk was suspended.

**User response:** None.

---

**6027-3315 File system scan CANCELLED due to new disks to be emptied or resume of all disks being emptied.**

**Explanation:** The parallel inode traversal (PIT) phase is cancelled during the file system restripe.

**User response:** None.

---

**6027-3316 Unable to create file system because there is not enough space for the log files. Number of log files: *numberOfLogFiles*. Log file size: *logFileSize*. Change one or more of the following as suggested and try again:**

**Explanation:** There is not enough space available to create all the required log files. This can happen when the storage pool is not large enough.

**User response:** Refer to the details given and correct the file system parameters.

---

**6027-3317 Warning: file system is not 4K aligned due to small *reasonString*. Native 4K sector disks cannot be added to this file system unless the disk that is used is *dataOnly* and the data block size is at least 128K.**

**Explanation:** The file system is created with a small inode or block size. Native 4K sector disk cannot be added to the file system, unless the disk that is used is *dataOnly* and the data block size is at least 128K.

**User response:** None.

---

**6027-3318 Fileset *filesetName* cannot be deleted as it is in compliant mode and it contains user files.**

**Explanation:** An attempt was made to delete a non-empty fileset that is in compliant mode.

**User response:** None.

---

**6027-3319 The AFM attribute *optionName* cannot be set for a primary fileset.**

**Explanation:** This AFM option cannot be set for a primary fileset. Hence, the fileset cannot be changed.

**User response:** None.

---

**6027-3320** *commandName:*  
**indefiniteRetentionProtection is enabled. File system cannot be deleted.**

**Explanation:** Indefinite retention is enabled for the file system so it cannot be deleted.

**User response:** None.

---

**6027-3400** **Attention: The file system is at risk. The specified replication factor does not tolerate unavailable metadata disks.**

**Explanation:** The default metadata replication was reduced to one while there were unavailable, or stopped, metadata disks. This condition prevents future file system manager takeover.

**User response:** Change the default metadata replication, or delete unavailable disks if possible.

---

**6027-3401** **Failure group value for disk *diskName* is not valid.**

**Explanation:** An explicit failure group must be specified for each disk that belongs to a write affinity enabled storage pool.

**User response:** Specify a valid failure group.

---

**6027-3402 [X]** **An unexpected device mapper path *dmDevice (nsdId)* was detected. The new path does not have Persistent Reserve enabled. The local access to disk *diskName* will be marked as down.**

**Explanation:** A new device mapper path was detected, or a previously failed path was activated after the local device discovery was finished. This path lacks a Persistent Reserve and cannot be used. All device paths must be active at mount time.

**User response:** Check the paths to all disks in the file system. Repair any failed paths to disks then rediscover the local disk access.

---

**6027-3404 [E]** **The current file system version does not support write caching.**

**Explanation:** The current file system version does not allow the write caching option.

**User response:** Use **mmchfs -V** to convert the file system to version 14.04 (4.1.0.0) or higher and reissue the command.

---

**6027-3405 [E]** **Cannot change the rapid repair, *"fileSystemName"* is mounted on number node(s).**

**Explanation:** Rapid repair can only be changed on unmounted file systems.

**User response:** Unmount the file system before running this command.

---

**6027-3406** **Error: Cannot add 4K native dataOnly disk *diskName* to non-4K aligned file system unless the file system version is at least 4.1.1.4.**

**Explanation:** An attempt was made through the **mmadddisk** command to add a 4K native disk to a non-4K aligned file system while the file system version is not at 4.1.1.4 or later.

**User response:** Upgrade the file system to 4.1.1.4 or later, and then retry the command.

---

**6027-3450** **Error *errorNumber* when purging key (file system *fileSystem*). Key name format possibly incorrect.**

**Explanation:** An error was encountered when purging a key from the key cache. The specified key name might have been incorrect, or an internal error was encountered.

**User response:** Ensure that the key name specified in the command is correct.

---

**6027-3451** **Error *errorNumber* when emptying cache (file system *fileSystem*).**

**Explanation:** An error was encountered when purging all the keys from the key cache.

**User response:** Contact the IBM Support Center.

---

**6027-3452 [E]** **Unable to create encrypted file *fileName* (inode *inodeNumber*, fileset *filesetNumber*, file system *fileSystem*).**

**Explanation:** Unable to create a new encrypted file. The key required to encrypt the file might not be available.

**User response:** Examine the error message following this message for information on the specific failure.

---

**6027-3453 [E]** **Unable to open encrypted file: inode *inodeNumber*, fileset *filesetNumber*, file system *fileSystem*.**

**Explanation:** Unable to open an existing encrypted file. The key used to encrypt the file might not be available.

**User response:** Examine the error message following this message for information on the specific failure.

---

**6027-3457 [E] Unable to rewrap key with name**  
*Keyname (inode inodeNumber, fileset  
 filesetNumber, file system fileSystem).*

**Explanation:** Unable to rewrap the key for a specified file because of an error with the key name.

**User response:** Examine the error message following this message for information on the specific failure.

---

**6027-3458 [E] Invalid length for the Keyname string.**

**Explanation:** The *Keyname* string has an incorrect length. The length of the specified string was either zero or it was larger than the maximum allowed length.

**User response:** Verify the *Keyname* string.

---

**6027-3459 [E] Not enough memory.**

**Explanation:** Unable to allocate memory for the *Keyname* string.

**User response:** Restart GPFS. Contact the IBM Support Center.

---

**6027-3460 [E] Incorrect format for the Keyname string.**

**Explanation:** An incorrect format was used when specifying the *Keyname* string.

**User response:** Verify the format of the *Keyname* string.

---

**6027-3461 [E] Error code: *errorNumber*.**

**Explanation:** An error occurred when processing a key ID.

**User response:** Contact the IBM Support Center.

---

**6027-3462 [E] Unable to rewrap key: original key name: *originalKeyname*, new key name: *newKeyname* (inode *inodeNumber*, fileset *filesetNumber*, file system *fileSystem*).**

**Explanation:** Unable to rewrap the key for a specified file, possibly because the existing key or the new key cannot be retrieved from the key server.

**User response:** Examine the error message following this message for information on the specific failure.

---

**6027-3463 [E] Rewrap error.**

**Explanation:** An internal error occurred during key rewrap.

**User response:** Examine the error messages surrounding this message. Contact the IBM Support Center.

---

**6027-3464 [E] New key is already in use.**

**Explanation:** The new key specified in a key rewrap is already being used.

**User response:** Ensure that the new key specified in the key rewrap is not being used by the file.

---

**6027-3465 [E] Cannot retrieve original key.**

**Explanation:** Original key being used by the file cannot be retrieved from the key server.

**User response:** Verify that the key server is available, the credentials to access the key server are correct, and that the key is defined on the key server.

---

**6027-3466 [E] Cannot retrieve new key.**

**Explanation:** Unable to retrieve the new key specified in the rewrap from the key server.

**User response:** Verify that the key server is available, the credentials to access the key server are correct, and that the key is defined on the key server.

---

**6027-3468 [E] Rewrap error code *errorNumber*.**

**Explanation:** Key rewrap failed.

**User response:** Record the error code and contact the IBM Support Center.

---

**6027-3469 [E] Encryption is enabled but the crypto module could not be initialized. Error code: *number*. Ensure that the GPFS crypto package was installed.**

**Explanation:** Encryption is enabled, but the cryptographic module required for encryption could not be loaded.

**User response:** Ensure that the packages required for encryption are installed on each node in the cluster.

---

**6027-3470 [E] Cannot create file *fileName*: extended attribute is too large: *numBytesRequired* bytes (*numBytesAvailable* available) (fileset *filesetNumber*, file system *fileSystem*).**

**Explanation:** Unable to create an encryption file because the extended attribute required for encryption is too large.

**User response:** Change the encryption policy so that the file key is wrapped fewer times, reduce the number of keys used to wrap a file key, or create a file system with a larger inode size.

---

---

6027-3471 [E] **At least one key must be specified.**

**Explanation:** No key name was specified.

**User response:** Specify at least one key name.

---

6027-3472 [E] **Could not combine the keys.**

**Explanation:** Unable to combine the keys used to wrap a file key.

**User response:** Examine the keys being used. Contact the IBM Support Center.

---

6027-3473 [E] **Could not locate the RKM.conf file.**

**Explanation:** Unable to locate the RKM.conf configuration file.

**User response:** Contact the IBM Support Center.

---

6027-3474 [E] **Could not open *fileType* file ('*fileName*' was specified).**

**Explanation:** Unable to open the specified configuration file. Encryption files will not be accessible.

**User response:** Ensure that the specified configuration file is present on all nodes.

---

6027-3475 [E] **Could not read file '*fileName*'.**

**Explanation:** Unable to read the specified file.

**User response:** Ensure that the specified file is accessible from the node.

---

6027-3476 [E] **Could not seek through file '*fileName*'.**

**Explanation:** Unable to seek through the specified file. Possible inconsistency in the local file system where the file is stored.

**User response:** Ensure that the specified file can be read from the local node.

---

6027-3477 [E] **Could not wrap the FEK.**

**Explanation:** Unable to wrap the file encryption key.

**User response:** Examine other error messages. Verify that the encryption policies being used are correct.

---

6027-3478 [E] **Insufficient memory.**

**Explanation:** Internal error: unable to allocate memory.

**User response:** Restart GPFS. Contact the IBM Support Center.

---

---

6027-3479 [E] **Missing combine parameter string.**

**Explanation:** The combine parameter string was not specified in the encryption policy.

**User response:** Verify the syntax of the encryption policy.

---

6027-3480 [E] **Missing encryption parameter string.**

**Explanation:** The encryption parameter string was not specified in the encryption policy.

**User response:** Verify the syntax of the encryption policy.

---

6027-3481 [E] **Missing wrapping parameter string.**

**Explanation:** The wrapping parameter string was not specified in the encryption policy.

**User response:** Verify the syntax of the encryption policy.

---

6027-3482 [E] **'*combineParameter*' could not be parsed as a valid combine parameter string.**

**Explanation:** Unable to parse the combine parameter string.

**User response:** Verify the syntax of the encryption policy.

---

6027-3483 [E] **'*encryptionParameter*' could not be parsed as a valid encryption parameter string.**

**Explanation:** Unable to parse the encryption parameter string.

**User response:** Verify the syntax of the encryption policy.

---

6027-3484 [E] **'*wrappingParameter*' could not be parsed as a valid wrapping parameter string.**

**Explanation:** Unable to parse the wrapping parameter string.

**User response:** Verify the syntax of the encryption policy.

---

6027-3485 [E] **The Keyname string cannot be longer than *number* characters.**

**Explanation:** The specified *Keyname* string has too many characters.

**User response:** Verify that the specified *Keyname* string is correct.

---

---

**6027-3486 [E] The KMIP library could not be initialized.**

**Explanation:** The KMIP library used to communicate with the key server could not be initialized.

**User response:** Restart GPFS. Contact the IBM Support Center.

---

**6027-3487 [E] The RKM ID cannot be longer than *number* characters.**

**Explanation:** The remote key manager ID cannot be longer than the specified length.

**User response:** Use a shorter remote key manager ID.

---

**6027-3488 [E] The length of the key ID cannot be zero.**

**Explanation:** The length of the specified key ID string cannot be zero.

**User response:** Specify a key ID string with a valid length.

---

**6027-3489 [E] The length of the RKM ID cannot be zero.**

**Explanation:** The length of the specified RKM ID string cannot be zero.

**User response:** Specify an RKM ID string with a valid length.

---

**6027-3490 [E] The maximum size of the RKM.conf file currently supported is *number* bytes.**

**Explanation:** The RKM.conf file is larger than the size that is currently supported.

**User response:** User a smaller RKM.conf configuration file.

---

**6027-3491 [E] The string '*Keyname*' could not be parsed as a valid key name.**

**Explanation:** The specified string could not be parsed as a valid key name.

**User response:** Specify a valid *Keyname* string.

---

**6027-3493 [E] *numKeys* keys were specified but a maximum of *numKeysMax* is supported.**

**Explanation:** The maximum number of specified key IDs was exceeded.

**User response:** Change the encryption policy to use fewer keys.

---



---

**6027-3494 [E] Unrecognized cipher mode.**

**Explanation:** Unable to recognize the specified cipher mode.

**User response:** Specify one of the valid cipher modes.

---

**6027-3495 [E] Unrecognized cipher.**

**Explanation:** Unable to recognize the specified cipher.

**User response:** Specify one of the valid ciphers.

---

**6027-3496 [E] Unrecognized combine mode.**

**Explanation:** Unable to recognize the specified combine mode.

**User response:** Specify one of the valid combine modes.

---

**6027-3497 [E] Unrecognized encryption mode.**

**Explanation:** Unable to recognize the specified encryption mode.

**User response:** Specify one of the valid encryption modes.

---

**6027-3498 [E] Invalid key length.**

**Explanation:** An invalid key length was specified.

**User response:** Specify a valid key length for the chosen cipher mode.

---

**6027-3499 [E] Unrecognized wrapping mode.**

**Explanation:** Unable to recognize the specified wrapping mode.

**User response:** Specify one of the valid wrapping modes.

---

**6027-3500 [E] Duplicate Keyname string '*keyIdentifier*'.**

**Explanation:** A given *Keyname* string has been specified twice.

**User response:** Change the encryption policy to eliminate the duplicate.

---

**6027-3501 [E] Unrecognized combine mode (*combineMode*).**

**Explanation:** The specified combine mode was not recognized.

**User response:** Specify a valid combine mode.

---

---

**6027-3502 [E] Unrecognized cipher mode ('cipherMode').**

**Explanation:** The specified cipher mode was not recognized.

**User response:** Specify a valid cipher mode.

---

**6027-3503 [E] Unrecognized cipher ('cipher').**

**Explanation:** The specified cipher was not recognized.

**User response:** Specify a valid cipher.

---

**6027-3504 [E] Unrecognized encryption mode ('mode').**

**Explanation:** The specified encryption mode was not recognized.

**User response:** Specify a valid encryption mode.

---

**6027-3505 [E] Invalid key length ('keyLength').**

**Explanation:** The specified key length was incorrect.

**User response:** Specify a valid key length.

---

**6027-3506 [E] Mode 'mode1' is not compatible with mode 'mode2', aborting.**

**Explanation:** The two specified encryption parameters are not compatible.

**User response:** Change the encryption policy and specify compatible encryption parameters.

---

**6027-3509 [E] Key 'keyID:RKMID' could not be fetched (RKM reported error errorNumber).**

**Explanation:** The key with the specified name cannot be fetched from the key server.

**User response:** Examine the error messages to obtain information about the failure. Verify connectivity to the key server and that the specified key is present at the server.

---

**6027-3510 [E] Could not bind symbol symbolName (errorDescription).**

**Explanation:** Unable to find the location of a symbol in the library.

**User response:** Contact the IBM Support Center.

---

**6027-3512 [E] The specified type 'type' for backend 'backend' is invalid.**

**Explanation:** An incorrect type was specified for a key server backend.

**User response:** Specify a correct backend type in RKM.conf.

---

---

**6027-3513 [E] Duplicate backend 'backend'.**

**Explanation:** A duplicate backend name was specified in RKM.conf.

**User response:** Specify unique RKM backends in RKM.conf.

---

**6027-3517 [E] Could not open library (libName).**

**Explanation:** Unable to open the specified library.

**User response:** Verify that all required packages are installed for encryption. Contact the IBM Support Center.

---

**6027-3518 [E] The length of the RKM ID string is invalid (must be between 0 and length characters).**

**Explanation:** The length of the RKM backend ID is invalid.

**User response:** Specify an RKM backend ID with a valid length.

---

**6027-3519 [E] 'numAttempts' is not a valid number of connection attempts.**

**Explanation:** The value specified for the number of connection attempts is incorrect.

**User response:** Specify a valid number of connection attempts.

---

**6027-3520 [E] 'sleepInterval' is not a valid sleep interval.**

**Explanation:** The value specified for the sleep interval is incorrect.

**User response:** Specify a valid sleep interval value (in microseconds).

---

**6027-3521 [E] 'timeout' is not a valid connection timeout.**

**Explanation:** The value specified for the connection timeout is incorrect.

**User response:** Specify a valid connection timeout (in seconds).

---

**6027-3522 [E] 'url' is not a valid URL.**

**Explanation:** The specified string is not a valid URL for the key server.

**User response:** Specify a valid URL for the key server.

---

---

**6027-3524 [E] *'tenantName'* is not a valid tenantName.**

**Explanation:** An incorrect value was specified for the tenant name.

**User response:** Specify a valid tenant name.

---

**6027-3527 [E] Backend *'backend'* could not be initialized (error *errorNumber*).**

**Explanation:** Key server backend could not be initialized.

**User response:** Examine the error messages. Verify connectivity to the server. Contact the IBM Support Center.

---

**6027-3528 [E] Unrecognized wrapping mode (*'wrapMode'*).**

**Explanation:** The specified key wrapping mode was not recognized.

**User response:** Specify a valid key wrapping mode.

---

**6027-3529 [E] An error was encountered while processing file *'fileName'*:**

**Explanation:** An error was encountered while processing the specified configuration file.

**User response:** Examine the error messages that follow and correct the corresponding conditions.

---

**6027-3530 [E] Unable to open encrypted file: key retrieval not initialized (inode *inodeNumber*, fileset *filesetNumber*, file system *fileSystem*).**

**Explanation:** File is encrypted but the infrastructure required to retrieve encryption keys was not initialized, likely because processing of RKM.conf failed.

**User response:** Examine error messages at the time the file system was mounted.

---

**6027-3533 [E] Invalid encryption key derivation function.**

**Explanation:** An incorrect key derivation function was specified.

**User response:** Specify a valid key derivation function.

---

**6027-3534 [E] Unrecognized encryption key derivation function (*'keyDerivation'*).**

**Explanation:** The specified key derivation function was not recognized.

**User response:** Specify a valid key derivation function.

---



---

**6027-3535 [E] Incorrect client certificate label *'clientCertLabel'* for backend *'backend'*.**

**Explanation:** The specified client keypair certificate label is incorrect for the backend.

**User response:** Ensure that the correct client certificate label is used in RKM.conf.

---

**6027-3537 [E] Setting default encryption parameters requires empty combine and wrapping parameter strings.**

**Explanation:** A non-empty combine or wrapping parameter string was used in an encryption policy rule that also uses the default parameter string.

**User response:** Ensure that neither the combine nor the wrapping parameter is set when the default parameter string is used in the encryption rule.

---

**6027-3540 [E] The specified RKM backend type (*rkmType*) is invalid.**

**Explanation:** The specified RKM type in RKM.conf is incorrect.

**User response:** Ensure that only supported RKM types are specified in RKM.conf.

---

**6027-3541 [E] Encryption is not supported on Windows.**

**Explanation:** Encryption cannot be activated if there are Windows nodes in the cluster.

**User response:** Ensure that encryption is not activated if there are Windows nodes in the cluster.

---

**6027-3543 [E] The integrity of the file encrypting key could not be verified after unwrapping; the operation was cancelled.**

**Explanation:** When opening an existing encrypted file, the integrity of the file encrypting key could not be verified. Either the cryptographic extended attributes were damaged, or the master key(s) used to unwrap the FEK have changed.

**User response:** Check for other symptoms of data corruption, and verify that the configuration of the key server has not changed.

---

**6027-3545 [E] Encryption is enabled but there is no valid license. Ensure that the GPFS crypto package was installed properly.**

**Explanation:** The required license is missing for the GPFS encryption package.

**User response:** Ensure that the GPFS encryption package was installed properly.

---

---

**6027-3546 [E] Key 'keyID:rkmID' could not be fetched. The specified RKM ID does not exist; check the RKM.conf settings.**

**Explanation:** The specified RKM ID part of the key name does not exist, and therefore the key cannot be retrieved. The corresponding RKM might have been removed from RKM.conf.

**User response:** Check the set of RKMs specified in RKM.conf.

---

**6027-3547 [E] Key 'keyID:rkmID' could not be fetched. The connection was reset by the peer while performing the TLS handshake.**

**Explanation:** The specified key could not be retrieved from the server, because the connection with the server was reset while performing the TLS handshake.

**User response:** Check connectivity to the server. Check credentials to access the server. Contact the IBM Support Center.

---

**6027-3548 [E] Key 'keyID:rkmID' could not be fetched. The IP address of the RKM could not be resolved.**

**Explanation:** The specified key could not be retrieved from the server because the IP address of the server could not be resolved.

**User response:** Ensure that the hostname of the key server is correct. Verify whether there are problems with name resolutions.

---

**6027-3549 [E] Key 'keyID:rkmID' could not be fetched. The TCP connection with the RKM could not be established.**

**Explanation:** Unable to establish a TCP connection with the key server.

**User response:** Check the connectivity to the key server.

---

**6027-3550 Error when retrieving encryption attribute: errorDescription.**

**Explanation:** Unable to retrieve or decode the encryption attribute for a given file.

**User response:** File could be damaged and may need to be removed if it cannot be read.

---

**6027-3551 Error flushing work file fileName: errorString**

**Explanation:** An error occurred while attempting to flush the named work file or socket.

**User response:** None.

---



---

**6027-3552 Failed to fork a new process to operationString file system.**

**Explanation:** Failed to fork a new process to suspend/resume file system.

**User response:** None.

---

**6027-3553 Failed to sync fileset filesetName.**

**Explanation:** Failed to sync fileset.

**User response:** None.

---

**6027-3554 The restore command encountered an out-of-memory error.**

**Explanation:** The fileset snapshot restore command encountered an out-of-memory error.

**User response:** None.

---

**6027-3555 name must be combined with FileInherit, DirInherit or both.**

**Explanation:** NoPropagateInherit must be accompanied by other inherit flags. Valid values are FileInherit and DirInherit.

**User response:** Specify a valid NFSv4 option and reissue the command.

---

**6027-3556 cmdName error: insufficient memory.**

**Explanation:** The command exhausted virtual memory.

**User response:** Consider some of the command parameters that might affect memory usage. Contact the IBM Support Center.

---

**6027-3557 cmdName error: could not create a temporary file.**

**Explanation:** A temporary file could not be created in the current directory.

**User response:** Ensure that the file system is not full and that files can be created. Contact the IBM Support Center.

---

**6027-3558 cmdName error: could not initialize the key management subsystem (error returnCode).**

**Explanation:** An internal component of the cryptographic library could not be properly initialized.

**User response:** Ensure that the gpfs.gskit package was installed properly. Contact the IBM Support Center.

---



---

**6027-3559** *cmdName* **error: could not create the key database (error *returnCode*).**

**Explanation:** The key database file could not be created.

**User response:** Ensure that the file system is not full and that files can be created. Contact the IBM Support Center.

---

**6027-3560** *cmdName* **error: could not create the new self-signed certificate (error *returnCode*).**

**Explanation:** A new certificate could not be successfully created.

**User response:** Ensure that the supplied canonical name is valid. Contact the IBM Support Center.

---

**6027-3561** *cmdName* **error: could not extract the key item (error *returnCode*).**

**Explanation:** The public key item could not be extracted successfully.

**User response:** Contact the IBM Support Center.

---

**6027-3562** *cmdName* **error: base64 conversion failed (error *returnCode*).**

**Explanation:** The conversion from or to the **BASE64** encoding could not be performed successfully.

**User response:** Contact the IBM Support Center.

---

**6027-3563** *cmdName* **error: could not extract the private key (error *returnCode*).**

**Explanation:** The private key could not be extracted successfully.

**User response:** Contact the IBM Support Center.

---

**6027-3564** *cmdName* **error: could not initialize the ICC subsystem (error *returnCode* *returnCode*).**

**Explanation:** An internal component of the cryptographic library could not be properly initialized.

**User response:** Ensure that the `gpfs.gskit` package was installed properly. Contact the IBM Support Center.

---

**6027-3565** *cmdName* **error: I/O error.**

**Explanation:** A terminal failure occurred while performing I/O.

**User response:** Contact the IBM Support Center.

---



---

**6027-3566** *cmdName* **error: could not open file '*fileName*'.**

**Explanation:** The specified file could not be opened.

**User response:** Ensure that the specified path and file name are correct and that you have sufficient permissions to access the file.

---

**6027-3567** *cmdName* **error: could not convert the private key.**

**Explanation:** The private key material could not be converted successfully.

**User response:** Contact the IBM Support Center.

---

**6027-3568** *cmdName* **error: could not extract the private key information structure.**

**Explanation:** The private key could not be extracted successfully.

**User response:** Contact the IBM Support Center.

---

**6027-3569** *cmdName* **error: could not convert the private key information to DER format.**

**Explanation:** The private key material could not be converted successfully.

**User response:** Contact the IBM Support Center.

---

**6027-3570** *cmdName* **error: could not encrypt the private key information structure (error *returnCode*).**

**Explanation:** The private key material could not be encrypted successfully.

**User response:** Contact the IBM Support Center.

---

**6027-3571** *cmdName* **error: could not insert the key in the keystore, check your system's clock (error *returnCode*).**

**Explanation:** Insertion of the new keypair into the keystore failed because the local date and time are not properly set on your system.

**User response:** Synchronize the local date and time on your system and try this command again.

---

**6027-3572** *cmdName* **error: could not insert the key in the keystore (error *returnCode*).**

**Explanation:** Insertion of the new keypair into the keystore failed.

**User response:** Contact the IBM Support Center.

---

---

**6027-3573** *cmdName* **error: could not insert the certificate in the keystore (error *returnCode*).**

**Explanation:** Insertion of the new certificate into the keystore failed.

**User response:** Contact the IBM Support Center.

---

**6027-3574** *cmdName* **error: could not initialize the digest algorithm.**

**Explanation:** Initialization of a cryptographic algorithm failed.

**User response:** Contact the IBM Support Center.

---

**6027-3575** *cmdName* **error: error while computing the digest.**

**Explanation:** Computation of the certificate digest failed.

**User response:** Contact the IBM Support Center.

---

**6027-3576** *cmdName* **error: could not initialize the SSL environment (error *returnCode*).**

**Explanation:** An internal component of the cryptographic library could not be properly initialized.

**User response:** Ensure that the `gpfs.gskit` package was installed properly. Contact the IBM Support Center.

---

**6027-3577** **Failed to sync fileset *filesetName*.**  
*errString*.

**Explanation:** Failed to sync fileset.

**User response:** Check the error message and try again. If the problem persists, contact the IBM Support Center.

---

**6027-3578 [E]** *pathName* **is not a valid argument for this command. You must specify a path name within a single GPFS snapshot.**

**Explanation:** This message is similar to message number 6027-872, but the *pathName* does not specify a path that can be scanned. The value specified for *pathName* might be a `.snapdir` or similar object.

**User response:** Correct the command invocation and reissue the command.

---

**6027-3579** *cmdName* **error: the cryptographic library could not be initialized in FIPS mode.**

**Explanation:** The cluster is configured to operate in FIPS mode but the cryptographic library could not be initialized in that mode.

**User response:** Verify that the `gpfs.gskit` package has been installed properly and that GPFS supports FIPS

mode on your platform. Contact the IBM Support Center.

---

**6027-3580** **Failed to sync file system: *fileSystem***  
**Error: *errString*.**

**Explanation:** Failed to sync file system.

**User response:** Check the error message and try again. If the problem persists, contact the IBM Support Center.

---

**6027-3581** **Failed to create the operation list file.**

**Explanation:** Failed to create the operation list file.

**User response:** Verify that the file path is correct and check the additional error messages.

---

**6027-3582 [E]** **Compression is not supported for clone or clone-parent files.**

**Explanation:** File compression is not supported as the file being compressed is a clone or a clone parent file.

**User response:** None.

---

**6027-3583 [E]** **Compression is not supported for snapshot files.**

**Explanation:** The file being compressed is within a snapshot and snapshot file compression is not supported.

**User response:** None.

---

**6027-3584 [E]** **Current file system version does not support compression.**

**Explanation:** The current file system version is not recent enough for file compression support.

**User response:** Upgrade the file system to the latest version and retry the command.

---

**6027-3585 [E]** **Compression is not supported for AFM cached files.**

**Explanation:** The file being compressed is cached in an AFM cache fileset and compression is not supported for such files.

**User response:** None.

---

**6027-3586 [E]** **Compression/uncompression failed.**

**Explanation:** Compression or uncompression failed.

**User response:** Refer to the error message below this line for the cause of the compression failure.

---

---

**6027-3587 [E] Aborting compression as the file is opened in hyper allocation mode.**

**Explanation:** Compression operation is not performed because the file is opened in hyper allocation mode.

**User response:** Compress this file after the file is closed.

---

**6027-3588 [E] Aborting compression as the file is currently memory mapped, opened in direct I/O mode, or stored in a horizontal storage pool.**

**Explanation:** Compression operation is not performed because it is inefficient or unsafe to compress the file at this time.

**User response:** Compress this file after the file is no longer memory mapped, opened in direct I/O mode, or stored in a horizontal storage pool.

---

**6027-3589** *cmdName* **error: Cannot set the password twice.**

**Explanation:** An attempt was made to set the password by using different available options.

**User response:** Set the password either through the CLI or by specifying a file that contains it.

---

**6027-3590** *cmdName* **error: Could not access file *fileName* (error *errorCode*).**

**Explanation:** The specified file could not be accessed.

**User response:** Check whether the file name is correct and verify whether you have required access privileges to access the file.

---

**6027-3591** *cmdName* **error: The password specified in file *fileName* exceeds the maximum length of *length* characters.**

**Explanation:** The password stored in the specified file is too long.

**User response:** Pick a shorter password and retry the operation.

---

**6027-3592** *cmdName* **error: Could not read the password from file *fileName*.**

**Explanation:** The password could not be read from the specified file.

**User response:** Ensure that the file can be read.

---



---

**6027-3593 [E] Compression is supported only for regular files.**

**Explanation:** The file is not compressed because compression is supported only for regular files.

**User response:** None.

---

**6027-3700 [E] Key '*keyID*' was not found on RKM ID '*rkmID*'.**

**Explanation:** The specified key could not be retrieved from the key server.

**User response:** Verify that the key is present at the server. Verify that the name of the keys used in the encryption policy is correct.

---

**6027-3701 [E] Key '*keyID:rkmID*' could not be fetched. The authentication with the RKM was not successful.**

**Explanation:** Unable to authenticate with the key server.

**User response:** Verify that the credentials used to authenticate with the key server are correct.

---

**6027-3702 [E] Key '*keyID:rkmID*' could not be fetched. Permission denied.**

**Explanation:** Unable to authenticate with the key server.

**User response:** Verify that the credentials used to authenticate with the key server are correct.

---

**6027-3703 [E] I/O error while accessing the keystore file '*keystoreFileName*'.**

**Explanation:** An error occurred while accessing the keystore file.

**User response:** Verify that the name of the keystore file in RKM.conf is correct. Verify that the keystore file can be read on each node.

---

**6027-3704 [E] The keystore file '*keystoreFileName*' has an invalid format.**

**Explanation:** The specified keystore file has an invalid format.

**User response:** Verify that the format of the keystore file is correct.

---

**6027-3705 [E] Incorrect FEK length after unwrapping; the operation was cancelled.**

**Explanation:** When opening an existing encrypted file, the size of the FEK that was unwrapped did not correspond to the one recorded in the file's extended attributes. Either the cryptographic extended attributes

were damaged, or the master key(s) used to unwrap the FEK have changed.

**User response:** Check for other symptoms of data corruption, and verify that the configuration of the key server has not changed.

---

**6027-3706 [E] The crypto library with FIPS support is not available for this architecture. Disable FIPS mode and reattempt the operation.**

**Explanation:** GPFS is operating in FIPS mode, but the initialization of the cryptographic library failed because FIPS mode is not yet supported on this architecture.

**User response:** Disable FIPS mode and attempt the operation again.

---

**6027-3707 [E] The crypto library could not be initialized in FIPS mode. Ensure that the crypto library package was correctly installed.**

**Explanation:** GPFS is operating in FIPS mode, but the initialization of the cryptographic library failed.

**User response:** Ensure that the packages required for encryption are properly installed on each node in the cluster.

---

**6027-3708 [E] Incorrect passphrase for backend 'backend'.**

**Explanation:** The specified passphrase is incorrect for the backend.

**User response:** Ensure that the correct passphrase is used for the backend in RKM.conf.

---

**6027-3709 [E] Error encountered when parsing line *lineNumber*: expected a new RKM backend stanza.**

**Explanation:** An error was encountered when parsing a line in RKM.conf. Parsing of the previous backend is complete, and the stanza for the next backend is expected.

**User response:** Correct the syntax in RKM.conf.

---

**6027-3710 [E] Error encountered when parsing line *lineNumber*: invalid key 'keyIdentifier'.**

**Explanation:** An error was encountered when parsing a line in RKM.conf.

**User response:** Specify a well-formed stanza in RKM.conf.

---

**6027-3711 [E] Error encountered when parsing line *lineNumber*: invalid key-value pair.**

**Explanation:** An error was encountered when parsing a line in RKM.conf: an invalid key-value pair was found.

**User response:** Correct the specification of the RKM backend in RKM.conf.

---

**6027-3712 [E] Error encountered when parsing line *lineNumber*: incomplete RKM backend stanza 'backend'.**

**Explanation:** An error was encountered when parsing a line in RKM.conf. The specification of the backend stanza was incomplete.

**User response:** Correct the specification of the RKM backend in RKM.conf.

---

**6027-3713 [E] An error was encountered when parsing line *lineNumber*: duplicate key 'key'.**

**Explanation:** A duplicate keyword was found in RKM.conf.

**User response:** Eliminate duplicate entries in the backend specification.

---

**6027-3714 [E] Incorrect permissions for the */var/mmfs/etc/RKM.conf* configuration file on node *nodeName*: the file must be owned by the root user and be in the root group, must be a regular file and be readable and writable by the owner only.**

**Explanation:** The permissions for the */var/mmfs/etc/RKM.conf* configuration file are incorrect. The file must be owned by the root user, must be in the root group, must be a regular file, and must be readable and writeable by the owner only.

**User response:** Fix the permissions on the file and retry the operation.

---

**6027-3715 [E] Error encountered when parsing line *lineNumber*: RKM ID 'RKMID' is too long, it cannot exceed *length* characters.**

**Explanation:** The RKMID chosen at the specified line of */var/mmfs/etc/RKM.conf* contains too many characters.

**User response:** Choose a shorter string for the RKMID.

---

**6027-3716 [E] Key 'keyID:rkmID' could not be fetched. The TLS handshake could not be completed successfully.**

**Explanation:** The specified key could not be retrieved

from the server because the TLS handshake did not complete successfully.

**User response:** Ensure that the configurations of GPFS and the remote key management (RKM) server are compatible when it comes to the version of the TLS protocol used upon key retrieval (GPFS uses the **nistCompliance** configuration variable to control that). In particular, if **nistCompliance=SP800-131A** is set in GPFS, ensure that the TLS v1.2 protocol is enabled in the RKM server. If this does not resolve the issue, contact the IBM Support Center.

---

**6027-3717 [E] Key 'keyID:rkmID' could not be fetched. The RKM is in quarantine after experiencing a fatal error.**

**Explanation:** GPFS has quarantined the remote key management (RKM) server and will refrain from initiating further connections to it for a limited amount of time.

**User response:** Examine the error messages that precede this message to determine the cause of the quarantine.

---

**6027-3718 [E] Key 'keyID:rkmID' could not be fetched. Invalid request.**

**Explanation:** The key could not be fetched because the remote key management (RKM) server reported that the request was invalid.

**User response:** Ensure that the RKM server trusts the client certificate that was used for this request. If this does not resolve the issue, contact the IBM Support Center.

---

**6027-3719 [W] Wrapping parameter string 'oldWrappingParameter' is not safe and will be replaced with 'newWrappingParameter'.**

**Explanation:** The wrapping parameter specified by the policy should no longer be used since it may cause data corruption or weaken the security of the system. For this reason, the wrapping parameter specified in the message will be used instead.

**User response:** Change the policy file and replace the specified wrapping parameter with a more secure one. Consult the *IBM Spectrum Scale: Advanced Administration Guide* for a list of supported wrapping parameters.

---

**6027-3900 Invalid flag 'flagName' in the criteria file.**

**Explanation:** An invalid flag was found in the criteria file.

**User response:** None.

---

**6027-3901 Failed to receive inode list: listName.**

**Explanation:** A failure occurred while receiving an inode list.

**User response:** None.

---

**6027-3902 Check file 'fileName' on fileSystem for inodes that were found matching the criteria.**

**Explanation:** The named file contains the inodes generated by parallel inode traversal (PIT) with interesting flags; for example, **dataUpdateMiss** or **BROKEN**.

**User response:** None.

---

**6027-3903 [W] quotaType quota is disabled or quota file is invalid.**

**Explanation:** The corresponding quota type is disabled or invalid, and cannot be copied.

**User response:** Verify that the corresponding quota type is enabled.

---

**6027-3904 [W] quotaType quota file is not a metadata file. File was not copied.**

**Explanation:** The quota file is not a metadata file, and it cannot be copied in this way.

**User response:** Copy quota files directly.

---

**6027-3905 [E] Specified directory does not exist or is invalid.**

**Explanation:** The specified directory does not exist or is invalid.

**User response:** Check the spelling or validity of the directory.

---

**6027-3906 [W] backupQuotaFile already exists.**

**Explanation:** The destination file for a metadata quota file backup already exists.

**User response:** Move or delete the specified file and retry.

---

**6027-3907 [E] No other quorum node found during cluster probe.**

**Explanation:** The node could not renew its disk lease and there was no other quorum node available to contact.

**User response:** Determine whether there was a network outage, and also ensure the cluster is configured with enough quorum nodes. The node will attempt to rejoin the cluster.

---

**6027-3908** Check file '*fileName*' on *fileSystem* for inodes with broken disk addresses or failures.

**Explanation:** The named file contains the inodes generated by parallel inode traversal (PIT) with interesting flags; for example, **dataUpdateMiss** or **BROKEN**.

**User response:** None.

---

**6027-3909** The file (*backupQuotaFile*) is a quota file in *fileSystem* already.

**Explanation:** The file is a quota file already. An incorrect file name might have been specified.

**User response:** None.

---

**6027-3910 [I]** Delay *number* seconds for safe recovery.

**Explanation:** When disk lease is in use, wait for the existing lease to expire before performing log and token manager recovery.

**User response:** None.

---

**6027-3911** Error reading message from the file system daemon: *errorString* : The system ran out of memory buffers or memory to expand the memory buffer pool.

**Explanation:** The system ran out of memory buffers or memory to expand the memory buffer pool. This prevented the client from receiving a message from the file system daemon.

**User response:** Try again later.

---

**6027-3912 [E]** File *fileName* cannot run with error *errorCode: errorString*.

**Explanation:** The named shell script cannot run.

**User response:** Verify that the file exists and that the access permissions are correct.

---

**6027-3913** Attention: disk *diskName* is a 4K native dataOnly disk and it is used in a non-4K aligned file system. Its usage is not allowed to change from dataOnly.

**Explanation:** An attempt was made through the `mmchdisk` command to change the usage of a 4K native disk in a non-4K aligned file system from `dataOnly` to something else.

**User response:** None.

---



---

**6027-3914 [E]** Current file system version does not support compression.

**Explanation:** File system version is not recent enough for file compression support.

**User response:** Upgrade the file system to the latest version, then retry the command.

---

**6027-4000 [I]** *descriptorType* descriptor on this NSD can be updated by running the following command from the node physically connected to NSD *nsdName*:

**Explanation:** This message is displayed when a descriptor validation thread finds a valid NSD, or disk, or stripe group descriptor but with a different ID. This can happen if a device is reused for another NSD.

**User response:** None. After this message, another message is displayed with a command to fix the problem.

---

**6027-4001 [I]** '`mmfsadm writeDesc <device> descriptorType descriptorId:descriptorId nsdFormatVersion pdiskStatus`', where device is the device name of that NSD.

**Explanation:** This message displays the command that must run to fix the NSD or disk descriptor on that device. The *deviceName* must be supplied by system administrator or obtained from `mmnsd -m` command. The *descriptorId* is a hexadecimal value.

**User response:** Run the command that is displayed on that NSD server node and replace *deviceName* with the device name of that NSD.

---

**6027-4002 [I]** Before running this command, check both NSDs. You might have to delete one of the NSDs.

**Explanation:** Informational message.

**User response:** The system administrator should decide which NSD to keep before running the command to fix it. If you want to keep the NSD found on disk, then you do not run the command. Instead, delete the other NSD found in cache (the NSD ID shown in the command).

---

**6027-4003 [E]** The on-disk *descriptorType* descriptor of *nsdName* *descriptorIdName* *descriptorId:descriptorId* is not valid because of bad *corruptionType*:

**Explanation:** The descriptor validation thread found an on-disk descriptor that is corrupted. GPFS will automatically fix it.

**User response:** None.

---

---

**6027-4004 [D] On-disk NSD descriptor:** *nsdId nsdId nsdMagic nsdMagic nsdFormatVersion nsdFormatVersion on disk nsdChecksum nsdChecksum calculated checksum calculatedChecksum nsdDescSize nsdDescSize firstPaxosSector firstPaxosSector nPaxosSectors nPaxosSectors nsdIsPdisk nsdIsPdisk*

**Explanation:** Description of an on-disk NSD descriptor.

**User response:** None.

---

**6027-4005 [D] Local copy of NSD descriptor:** *nsdId nsdId nsdMagic nsdMagic formatVersion formatVersion nsdDescSize nsdDescSize firstPaxosSector firstPaxosSector nPaxosSectors nPaxosSectors*

**Explanation:** Description of the cached NSD descriptor.

**User response:** None.

---

**6027-4006 [I] Writing NSD descriptor of *nsdName* with local copy:** *nsdId nsdId nsdFormatVersion formatVersion firstPaxosSector firstPaxosSector nPaxosSectors nPaxosSectors nsdDescSize nsdDescSize nsdIsPdisk nsdIsPdisk nsdChecksum nsdChecksum*

**Explanation:** Description of the NSD descriptor that was written.

**User response:** None.

---

**6027-4007** *errorType descriptor on descriptorType nsdId nsdId:nsdId error error*

**Explanation:** This message is displayed after reading and writing NSD, disk and stripe group descriptors.

**User response:** None.

---

**6027-4008 [E] On-disk *descriptorType* descriptor of *nsdName* is valid but has a different UID:** *uid descriptorId:descriptorId on-disk uid descriptorId:descriptorId nsdId nsdId:nsdId*

**Explanation:** While verifying an on-disk descriptor, a valid descriptor was found but with a different ID. This can happen if a device is reused for another NSD with the **mmcrnsd -v no** command.

**User response:** After this message there are more messages displayed that describe the actions to follow.

---

**6027-4009 [E] On-disk NSD descriptor of *nsdName* is valid but has a different ID. ID in cache is *cachedId* and ID on-disk is *ondiskId***

**Explanation:** While verifying an on-disk NSD descriptor, a valid descriptor was found but with a different ID. This can happen if a device is reused for another NSD with the **mmcrnsd -v no** command.

**User response:** After this message, there are more messages displayed that describe the actions to follow.

---

**6027-4010 [I] This corruption can happen if the device is reused by another NSD with the -v option and a file system is created with that reused NSD.**

**Explanation:** Description of a corruption that can happen when an NSD is reused.

**User response:** Verify that the NSD was not reused to create another NSD with the **-v** option and that the NSD was not used for another file system.

---

**6027-4011 [D] On-disk disk descriptor: uid** *descriptorID:descriptorID magic descMagic formatVersion formatVersion descSize descSize checksum on disk diskChecksum calculated checksum calculatedChecksum firstSGDescSector firstSGDescSector nSGDescSectors nSGDescSectors lastUpdateTime lastUpdateTime*

**Explanation:** Description of the on-disk disk descriptor.

**User response:** None.

---

**6027-4012 [D] Local copy of disk descriptor: uid** *descriptorID:descriptorID firstSGDescSector firstSGDescSector nSGDescSectors nSGDescSectors*

**Explanation:** Description of the cached disk descriptor.

**User response:** None.

---

**6027-4013 [I] Writing disk descriptor of *nsdName* with local copy:** *uid descriptorID:descriptorID, magic magic, formatVersion formatVersion firstSGDescSector firstSGDescSector nSGDescSectors nSGDescSectors descSize descSize*

**Explanation:** Writing disk descriptor to disk with local information.

**User response:** None.

---

6027-4014 [D] **Local copy of StripeGroup descriptor:**

**uid** *descriptorID:descriptorID*  
**curFmtVersion** *curFmtVersion*  
**configVersion** *configVersion*

**Explanation:** Description of the cached stripe group descriptor.

**User response:** None.

---

6027-4015 [D] **On-disk StripeGroup descriptor: uid**  
*sgUid:sgUid* **magic** *magic* **curFmtVersion**  
*curFmtVersion* **descSize** *descSize* **on-disk**  
**checksum** *diskChecksum* **calculated**  
**checksum** *calculatedChecksum*  
**configVersion** *configVersion*

**lastUpdateTime** *lastUpdateTime*

**Explanation:** Description of the on-disk stripe group descriptor.

**User response:** None.

---

6027-4016 [E] **Data buffer checksum mismatch during write. File system** *fileSystem* **tag** *tag1 tag2*  
**nBytes** *nBytes* *diskAddresses*

**Explanation:** GPFS detected a mismatch in the checksum of the data buffer content which means content of data buffer was changing while a direct I/O write operation was in progress.

**User response:** None.



---

## Accessibility features for IBM Spectrum Scale

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 IBM Spectrum Scale:

- 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

IBM Knowledge Center, and its related publications, are accessibility-enabled. The accessibility features are described in IBM Knowledge Center ([www.ibm.com/support/knowledgecenter](http://www.ibm.com/support/knowledgecenter)).

---

### Keyboard navigation

This product uses standard Microsoft Windows navigation keys.

---

### IBM and accessibility

See the IBM Human Ability and Accessibility Center ([www.ibm.com/able](http://www.ibm.com/able)) for more information about the commitment that IBM has to accessibility.



---

## Notices

This information was developed for products and services that are offered in the USA.

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 that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right 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 grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing  
IBM Corporation  
North Castle Drive, MD-NC119  
Armonk, NY 10504-1785  
United States of America*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing  
Legal and Intellectual Property Law  
IBM Japan Ltd.  
19-21, Nihonbashi-Hakozakicho, Chuo-ku  
Tokyo 103-8510, 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 websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites 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  
Dept. H6MA/Building 707  
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 of 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.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

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 illustrate 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. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

Portions of this code are derived from IBM Corp. Sample Programs.

© Copyright IBM Corp. \_enter the year or years\_. All rights reserved.

---

## **Trademarks**

IBM, the IBM logo, and [ibm.com](http://ibm.com) are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at Copyright and trademark information at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml).

Intel is a trademark of Intel Corporation or its subsidiaries in the United States and other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of the Open Group in the United States and other countries.

---

## **Terms and conditions for product documentation**

Permissions for the use of these publications are granted subject to the following terms and conditions.

### **Applicability**

These terms and conditions are in addition to any terms of use for the IBM website.

### **Personal use**

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

### **Commercial use**

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

### **Rights**

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

---

## **IBM Online Privacy Statement**

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, See IBM's Privacy Policy at <http://www.ibm.com/privacy> and IBM's Online Privacy Statement at <http://www.ibm.com/privacy/details> the section entitled "Cookies, Web Beacons and Other Technologies" and the "IBM Software Products and Software-as-a-Service Privacy Statement" at <http://www.ibm.com/software/info/product-privacy>.

---

## Glossary

This glossary provides terms and definitions for IBM Spectrum Scale.

The following cross-references are used in this glossary:

- *See* refers you from a nonpreferred term to the preferred term or from an abbreviation to the spelled-out form.
- *See also* refers you to a related or contrasting term.

For other terms and definitions, see the IBM Terminology website ([www.ibm.com/software/globalization/terminology](http://www.ibm.com/software/globalization/terminology)) (opens in new window).

### 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 must be a quorum node. The selection of the cluster manager node favors the quorum-manager node with the lowest node number among the nodes that are operating at that particular time.

**Note:** The cluster manager role is not moved to another node when a node with a lower node number becomes active.

#### 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.

#### dependent fileset

A fileset that shares the inode space of an existing independent fileset.

#### 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*.

#### 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.

#### 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.

**domain**

A logical grouping of resources in a network for the purpose of common management and administration.

**E****ECKD™**

See *extended count key data (ECKD)*.

**ECKD device**

See *extended count key data device (ECKD device)*.

**encryption key**

A mathematical value that allows components to verify that they are in communication with the expected server. Encryption keys are based on a public or private key pair that is created during the installation process. See also *file encryption key*, *master encryption key*.

**extended count key data (ECKD)**

An extension of the count-key-data (CKD) architecture. It includes additional commands that can be used to improve performance.

**extended count key data device (ECKD device)**

A disk storage device that has a data transfer rate faster than some processors can utilize and that is connected to the processor through use of a speed matching buffer. A specialized channel program is needed to communicate with such a device. See also *fixed-block architecture disk device*.

**F****failback**

Cluster recovery from failover following repair. See also *failover*.

**failover**

(1) The assumption of file system duties by another node when a node fails. (2) The process of transferring all control of the ESS to a single cluster in the ESS when the other clusters in the ESS fails. See also *cluster*. (3) The routing of all transactions to a second controller when the first controller fails. See also *cluster*.

**failure group**

A collection of disks that share common access paths or adapter connection, and could all become unavailable through a single hardware failure.

**FEK** See *file encryption key*.

**fileset** A hierarchical grouping of files managed as a unit for balancing workload across a cluster. See also *dependent fileset*, *independent fileset*.

**fileset snapshot**

A snapshot of an independent fileset plus all dependent filesets.

**file clone**

A writable snapshot of an individual file.

**file encryption key (FEK)**

A key used to encrypt sectors of an individual file. See also *encryption key*.

**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.

**fixed-block architecture disk device (FBA disk device)**

A disk device that stores data in blocks of fixed size. These blocks are addressed by block number relative to the beginning of the file. See also *extended count key data device*.

**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

### **global snapshot**

A snapshot of an entire GPFS file system.

### **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.

### **independent fileset**

A fileset that has its own inode space.

### **indirect block**

A block containing pointers to other blocks.

**inode** The internal structure that describes the individual files in the file system. There is one inode for each file.

### **inode space**

A collection of inode number ranges reserved for an independent fileset, which enables more efficient per-fileset functions.

## ISKLM

IBM Security Key Lifecycle Manager. For GPFS encryption, the ISKLM is used as an RKM server to store MEKs.

## 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.

## M

### **master encryption key (MEK)**

A key used to encrypt other keys. See also *encryption key*.

**MEK** See *master encryption key*.

### **metadata**

Data structures that contain information that is needed to access file data. Metadata includes inodes, indirect blocks, and directories. Metadata is 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, service-class, and encryption 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 performed.

**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.

**remote key management server (RKM server)**

A server that is used to store master encryption keys.

**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.

**RKM server**

See *remote key management server*.

**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.

**Scale Out Backup and Restore (SOBAR)**

A specialized mechanism for data protection against disaster only for GPFS file systems that are managed by Tivoli Storage Manager (TSM) Hierarchical Storage Management (HSM).

**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.

**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**

An exact copy of changed data in the active files and directories of a file system or fileset at a single point in time. See also *fileset snapshot*, *global snapshot*.

**source node**

The node on which a data management event is generated.

**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**

**VFS** See *virtual file system*.

**virtual file system (VFS)**

A remote file system that has been mounted so that it is accessible to the local user.

**virtual node (vnode)**

The structure that contains information about a file system object in a virtual file system (VFS).

---

# Index

## Special characters

- /etc/filesystems 96
- /etc/fstab 96
- /etc/hosts 74
- /etc/resolv.conf 93
- /tmp/mmfs 146, 167
- /usr/lpp/mmfs/bin 79
- /usr/lpp/mmfs/bin/runmmfs 34
- /usr/lpp/mmfs/samples/gatherlogs.samples.sh file 3
- /var/adm/ras/mmfs.log.previous 89
- /var/mmfs/etc/mmlock 77
- /var/mmfs/gen/mmsdrfs 77
- .ptrash directory 148
- .rhosts 76
- .snapshots 116, 118, 119

## A

- access
  - to disk 131
- ACCESS\_TIME attribute 55, 56
- accessibility features for IBM Spectrum Scale 317
- active file management, questions related to 148
- administration commands
  - failure 77
- AFM 148
- AFM fileset, changing mode of 148
- AFM, extended attribute size supported by 148
- AFM, messages requeuing 124
- AIX
  - kernel debugger 71
- AIX error logs
  - MMFS\_DISKFAIL 131
  - MMFS\_QUOTA 106
  - unavailable disks 106
- AIX logical volume
  - down 136
- AIX platform
  - gpfs.snap command 24
- application program errors 92
- application programs
  - errors 20, 22, 83, 91
- authentication 26
  - problem determination 75
- authorization error 76
- autofs 99
- autofs mount 98
- autoload option
  - on mmchconfig command 80
  - on mmcrcluster command 80
- automount 98, 103
- automount daemon 98
- automount failure 98, 99, 100
- Availability 151

## B

- back up
  - cluster data 78

## C

- candidate file 51, 54
  - attributes 55
- CES
  - monitoring 11
  - troubleshooting 11
- CES administration 11
- CES collection 13
- CES monitoring 11
- CES service logs 3
- CES tracing 13
- changing mode of AFM fileset 148
- checking, Persistent Reserve 139
- chosen file 51, 53
- CIFS serving, Windows SMB2 protocol 93
- cipherList 103
- Clearing a leftover Persistent Reserve reservation 139
- client node 103
- clock synchronization 2, 112
- cluster
  - deleting a node 90
- cluster configuration information
  - displaying 44
- cluster data
  - backup 78
- Cluster Export Services
  - administration 11
  - issue collection 13
  - monitoring 11
  - tracing 13
- cluster file systems
  - displaying 45
- cluster overload detection 68
- cluster security configuration 101
- cluster state information 43
- commands
  - cluster state information 43
  - conflicting invocation 95
  - errprt 167
  - file system and disk information 49
  - gpfs.snap 23, 24, 25, 167
  - grep 19
  - lslpp 167
  - lslv 145
  - lsdf 50, 105
  - lspv 137
  - lsvg 136
  - lxtrace 33, 34
  - mmadddisk 109, 114, 133, 136, 138
  - mmaddnode 87, 88, 146
  - mmafmctl 124
  - mmafmctl Device getstate 43
  - mmapplypolicy 51, 111, 112, 115, 144
  - mmauth 61, 101
  - mmbackup 116
  - mmchcluster 75
  - mmchconfig 45, 80, 88, 103
  - mmchdisk 96, 106, 109, 114, 115, 127, 130, 131, 133, 135
  - mmcheckquota 21, 57, 92, 106
  - mmchfs 22, 78, 86, 90, 96, 98, 106, 123
  - mmchnode 146

- commands (*continued*)
  - mmchnsd 127
  - mmcommon recoverfs 109
  - mmcommon showLocks 77
  - mmcluster 45, 75, 80, 87, 146
  - mmcrfs 90, 123, 127, 138
  - mmcrnsd 127, 130
  - mmcrsnapshot 117, 118
  - mmdeldisk 109, 114, 133, 136
  - mmdelfileset 113
  - mmdelfs 134, 135
  - mmdelnode 87, 90
  - mmdelnsd 130, 134
  - mmdelsnapshot 117
  - mmdf 86, 110, 136
  - mmdiag 43
  - mmdsh 76
  - mmdumpperfdata 31
  - mmexpelnode 46
  - mmfileid 59, 124, 133
  - mmfsadm 33, 37, 81, 87, 124, 133
  - mmfsc 49, 95, 96, 114, 124, 134, 136, 147
  - mmgetstate 43, 81, 89
  - mmlsattr 113
  - mmlscluster 44, 87, 101, 145
  - mmlsconfig 34, 45, 98
  - mmlsdisk 90, 95, 96, 106, 109, 127, 130, 132, 135, 168
  - mmlsfileset 113
  - mmlsfs 97, 133, 134, 167
  - mmlsmgr 33, 96
  - mmlsmount 50, 80, 91, 95, 105, 106, 127
  - mmlsnsd 57, 128, 136
  - mmlspolicy 112
  - mmlsquota 91, 92
  - mmlssnapshot 116, 117, 118
  - mmmound 49, 95, 106, 138
  - mmpmon 71, 119, 120
  - mmquotaoff 92
  - mmquotaon 92
  - mmrefresh 45, 96, 98
  - mmremotecluster 61, 101, 102
  - mmremotefs 98, 101
  - mmrepquota 92
  - mmrestorefs 117, 118, 119
  - mmrestripefile 112, 115
  - mmrestripefs 115, 133, 136
  - mmrpldisk 109, 114, 138
  - mmsdrrestore 46
  - mmshutdown 44, 46, 80, 81, 83, 98, 99, 108
  - mmsnapdir 116, 118, 119
  - mmstartup 80, 98, 99
  - mmumount 105, 106, 136
  - mmunlinkfileset 113
  - mmwindisk 58
  - mount 95, 96, 98, 134, 138
  - ping 76
  - rcp 75
  - rpm 167
  - rsh 75, 89
  - scp 76
  - ssh 76
  - umount 105, 106
  - varyonvg 138
- commands, administration
  - failure 77
- communication paths
  - unavailable 96

- compiling mmfslinux module 79
- configuration
  - hard loop ID 75
  - performance tuning 75
- configuration data 109
- configuration parameters
  - kernel 79
- configuration problems 73
- configuration variable settings
  - displaying 45
- connectivity problems 76
- contact node address 101
- contact node failure 102
- creating a file, failure 143
- creating a master GPFS log file 2
- cron 146

## D

- data
  - replicated 133
- data always gathered by gpfs.snap 23
  - for a master snapshot 25
  - on AIX 24
  - on all platforms 23
  - on Linux 25
  - on Windows 25
- Data always gathered for an Object on Linux 27
- Data always gathered for authentication on Linux 28
- Data always gathered for CES on Linux 28
- Data always gathered for NFS on Linux 27
- Data always gathered for performance on Linux 29
- Data always gathered for SMB on Linux 26
- data collection 26
- data file issues
  - cluster configuration 77
- data gathered by
  - gpfs.snap on Linux 26
- data integrity 22, 124
- Data Management API (DMAPI)
  - file system will not mount 97
- data replication 132
- data structure 20
- dataOnly attribute 114
- dataStructureDump 34
- dead man switch timer 85
- deadlock
  - automated breakup 66
  - breakup on demand 67
  - cluster overload detection 68
- deadlocks 86, 87
  - automated data collection 65
  - automated detection 63
  - information about 63
- debug data collection
  - CES tracing 13
- delays 86, 87
- DELETE rule 51, 54
- deleting a node
  - from a cluster 90
- descOnly 107
- diagnostic data
  - deadlock diagnostics 41
  - standard diagnostics 41
- directories
  - /tmp/mmfs 146, 167
  - .snapshots 116, 118, 119

- directory that has not been cached, traversing 148
- disabling IPv6
  - for SSH connection delays 93
- disabling Persistent Reserve manually 140
- disaster recovery
  - other problems 89
  - setup problems 88
- disk access 131
- disk commands
  - hang 138
- disk connectivity failure 135
- disk descriptor replica 106
- disk failover 135
- disk failure 135
- disk leasing 85
- disk recovery 135
- disk subsystem
  - failure 127
- disks
  - damaged files 59
  - declared down 130
  - define for GPFS use 136
  - displaying information of 57
  - failure 20, 22, 127
  - media failure 132
  - partial failure 136
  - replacing 109
  - usage 107
- disks down 136
- disks, viewing 58
- displaying disk information 57
- displaying NSD information 128
- DNS server failure 101

## E

- enabling Persistent Reserve manually 140
- encryption issues 143
  - issues with adding encryption policy 143
  - permission denied message 143
- ERRNO I/O error code 89
- error codes
  - EINVAL 112
  - EIO 20, 127, 134
  - ENODEV 83
  - ENOENT 105
  - ENOSPC 110, 134
  - ERRNO I/O 89
  - ESTALE 22, 83, 105
  - NO SUCH DIRECTORY 83
  - NO SUCH FILE 83
- error log
  - MMFS\_LONGDISKIO 21
  - MMFS\_QUOTA 21
- error logs 1
  - example 22
  - MMFS\_ABNORMAL\_SHUTDOWN 20
  - MMFS\_DISKFAIL 20
  - MMFS\_ENVIRON 20
  - MMFS\_FSSTRUCT 20
  - MMFS\_GENERIC 20
  - MMFS\_LONGDISKIO 21
  - MMFS\_QUOTA 21, 57
  - MMFS\_SYSTEM\_UNMOUNT 22
  - MMFS\_SYSTEM\_WARNING 22
  - operating system 19

- error messages
  - 0516-1339 130
  - 0516-1397 130
  - 0516-862 130
  - 6027-1209 83
  - 6027-1242 77
  - 6027-1290 109
  - 6027-1598 87
  - 6027-1615 76
  - 6027-1617 76
  - 6027-1627 91
  - 6027-1628 77
  - 6027-1630 77
  - 6027-1631 77
  - 6027-1632 77
  - 6027-1633 77
  - 6027-1636 128
  - 6027-1661 128
  - 6027-1662 130
  - 6027-1995 116
  - 6027-1996 108
  - 6027-2108 128
  - 6027-2109 128
  - 6027-300 80
  - 6027-306 82
  - 6027-319 81, 82
  - 6027-320 82
  - 6027-321 82
  - 6027-322 82
  - 6027-341 79, 82
  - 6027-342 79, 82
  - 6027-343 79, 82
  - 6027-344 79, 83
  - 6027-361 135
  - 6027-418 107, 135
  - 6027-419 97, 107
  - 6027-435 88
  - 6027-473 107
  - 6027-474 107
  - 6027-482 97, 135
  - 6027-485 135
  - 6027-490 88
  - 6027-506 92
  - 6027-533 86
  - 6027-538 90
  - 6027-549 97
  - 6027-580 97
  - 6027-631 109
  - 6027-632 108, 109
  - 6027-635 108
  - 6027-636 108, 135
  - 6027-638 109
  - 6027-645 97
  - 6027-650 83
  - 6027-663 91
  - 6027-665 80, 91
  - 6027-695 92
  - 6027-953 118
  - ANS1312E 116
  - cluster configuration data file issues 77
  - descriptor replica 88
  - disk media failures 135
  - failed to connect 80, 135
  - file system forced unmount problems 107
  - file system mount problems 97
  - GPFS cluster data recovery 77
  - incompatible version number 81

- error messages (*continued*)
  - mmbackup 116
  - mmfsd ready 80
  - multiple file system manager failures 108
  - network problems 82
  - quorum 88
  - rsh problems 76
  - shared segment problems 81, 82
  - snapshot 116, 117, 118
  - TSM 116
- error numbers
  - application calls 98
  - configuration problems 78
  - data corruption 124
  - EALL\_UNAVAIL = 218 108
  - ECONFIG = 208 78
  - ECONFIG = 215 78, 82
  - ECONFIG = 218 79
  - ECONFIG = 237 78
  - ENO\_MGR = 212 109, 135
  - ENO\_QUOTA\_INST = 237 98
  - EOFFLINE = 208 135
  - EPANIC = 666 107
  - EVALIDATE = 214 124
  - file system forced unmount 107
  - GPFS application calls 135
  - GPFS daemon will not come up 82
  - installation problems 78
  - multiple file system manager failures 109
- errors, application program 92
- errors, Persistent Reserve 138
- errpt command 167
- events
  - Availability 151
  - Reliability 151
  - Serviceability 151
- example
  - error logs 22
- EXCLUDE rule 55
- excluded file 55
  - attributes 55
- extended attribute size supported by AFM 148

## F

- facility
  - Linux kernel crash dump (LKCD) 71
- failure
  - disk 130
    - mmccr command 148
    - mmfsck command 147
    - of disk media 132
    - snapshot 116
  - failure creating a file 143
  - failure group 106
  - failure groups
    - loss of 107
    - use of 106
  - failure, key rewrap 144
  - failure, mount 143
  - failures
    - mmbackup 116
- File Authentication
  - setup problems 93
- file creation failure 143
- file migration
  - problems 113

- File Placement Optimizer (FPO), questions related to 148
- file placement policy 112
- file system 95
  - mount status 108
  - space 110
- file system descriptor 106, 107
  - failure groups 106
  - inaccessible 107
- file system manager
  - cannot appoint 105
  - contact problems
    - communication paths unavailable 96
    - multiple failures 108
- file system mount failure 143
- file system or fileset getting full 148
- file systems
  - cannot be unmounted 50
  - creation failure 90
  - determining if mounted 108
  - discrepancy between configuration data and on-disk data 109
  - do not mount 95
  - does not mount 95
  - does not unmount 104
  - forced unmount 22, 105, 108
  - free space shortage 118
  - listing mounted 50
  - loss of access 91
  - remote 100
  - state after restore 118
  - unable to determine if mounted 108
  - will not mount 49
- FILE\_SIZE attribute 55, 56
- files
  - /etc/filesystems 96
  - /etc/fstab 96
  - /etc/group 21
  - /etc/hosts 74
  - /etc/passwd 21
  - /etc/resolv.conf 93
  - /usr/lpp/mmfs/bin/runmmfs 34
  - /usr/lpp/mmfs/samples/gatherlogs.samples.sh 3
  - /var/adm/ras/mmfs.log.previous 89
  - /var/mmfs/etc/mmlock 77
  - /var/mmfs/gen/mmsdrfs 77
  - .rhosts 76
  - detecting damage 59
  - mmfs.log 2, 80, 81, 83, 95, 99, 100, 101, 102, 103, 104, 105, 167
  - mmsdrbackup 78
  - mmsdrfs 78
    - protocol authentication log 9
- FILESET\_NAME attribute 55, 56
- filesets
  - child 113
  - deleting 113
  - emptying 113
  - errors 114
  - lost+found 114
  - moving contents 113
  - performance 113
  - problems 109
  - snapshots 113
  - unlinking 113
  - usage errors 113
- FPO 148
- FSDesc structure 106



full file system or fileset 148

## G

generate

trace reports 34

generating GPFS trace reports

mmtracectl command 34

GPFS

/tmp/mmfs directory 146

abnormal termination in mmpmon 120

active file management 148

AFM 124

AIX 99

application program errors 92

authentication issues 93

automount 98

automount failure 99

automount failure in Linux 98

Availability 151

checking Persistent Reserve 139

cipherList option has not been set properly 103

clearing a leftover Persistent Reserve reservation 139

client nodes 103

cluster configuration

issues 77, 78

cluster name 101

cluster security configurations 101

cluster state information commands 43, 44, 45, 46

command 23, 24, 25, 43

configuration data 109

contact node address 101

contact nodes down 102

core dumps 38

corrupted data integrity 124

data gathered for protocol on Linux 26, 27, 28, 29, 30

data integrity 124

data integrity may be corrupted 124

deadlocks 63, 65, 66, 67, 68

delays and deadlocks 86

determine if a file system is mounted 108

determining the health of integrated SMB server 122

disaster recovery issues 88

discrepancy between GPFS configuration data and the on-disk data for a file system 109

disk accessing command failure 138

disk connectivity failure 135

disk failure 132, 135

disk information commands 49, 50, 51, 58, 59

disk issues 85, 127

disk media failure 132, 135

disk recovery 135

disk subsystem failures 127

displaying NSD information 128

encryption rules 143

error creating internal storage 147

error encountered while creating NSD disks 127

error encountered while using NSD disks 127

error messages for file system 97, 98

error message 108, 148

error messages 116, 117, 118, 135

error messages for file system forced unmount problems 107

error messages for file system mount status 108

error messages for indirect snapshot errors 116

error messages not directly related to snapshots 116

error messages related to snapshots 117

GPFS (continued)

error message "Function not implemented" 100

error numbers 107, 109, 135

error numbers specific to GPFS application calls 124

errors 112, 113, 138

errors associated with filesets 109

errors associated with policies 109

errors associated with storage pools, 109

errors encountered 115

errors encountered with filesets 114

failure group considerations 106

failures using the mmbackup command 116

file placement optimizer 148

file system 104, 105, 143

file system commands 49, 50, 51, 58, 59

file system failure 95

file system has adequate free space 110

file system is forced to unmount 107

file system is mounted 108

file system issues 95

file system manager appointment fails 109

file system manager failures 109

file system mount problems 97, 98

file system mount status 108

file system mounting 147

file systems manager failure 108

filesets usage 113

forced unmount 105

gpfs.snap 23, 24, 25

guarding against disk failures 132

GUI logs 41

hang in mmpmon 120

health of integrated SMB server 122

ill-placed files 111

incorrect output from mmpmon 120

indirect issues with snapshot 116

installation and configuration issues 73, 74, 77, 79, 80, 81, 82, 83, 85, 87, 89, 92

integrated SMB server 122

issues while working with Samba 123

issues with snapshot 116, 117

key rewrap 144

local node failure 102

locating snapshot 116

log 1, 2

logical volumes are properly defined 136

manually disabling Persistent Reserve 140

manually enabling Persistent Reserve 140

mapping 100

master log file 2

message 6027-648 147

message referring to an existing NSD 130

message requeuing 124

message requeuing in AFM 124

message severity tags 171

messages 173

mmafmctl Device getstate 43

mmapplypolicy -L command 52, 53, 54, 55, 56

mmbackup command 116

mmbackup errors 116

mmdumpperfdata command 31

mmexpelnode command 46

mmfsadm command 33

mmpmon 120

mmpmon command 120

mmpmon output 120

mmremotecluster command 101

- GPFS (*continued*)
  - mount 98, 100, 147
  - mount failure 103, 143
  - mounting cluster 102
  - mounting cluster does not have direct access to the disks 102
  - multipath device 141
  - multiple file system manager failures 108
  - negative values in the 'predicted pool utilizations', 111
  - NFS client 121
  - NFS problems 121
  - NFS V4 121
  - NFS V4 issues 121
  - NFS V4 problem 121
  - no replication 134
  - NO\_SPACE error 110
  - nodes will not start 81
  - NSD creation failure 130
  - NSD disk does not have an NSD server specified 102
  - NSD information 128
  - NSD is down 130
  - NSD server 103
  - NSD subsystem failures 127
  - NSDs built on top of AIX logical volume is down 136
  - offline mmfsck command failure 147
  - old inode data 121
  - on-disk data 109
  - Operating system error logs 19
  - partial disk failure 136
  - permission denied error message 103
  - permission denied failure 144
  - Persistent Reserve errors 138
  - physical disk association 145
  - physical disk association with logical volume 145
  - policies 111, 112
  - predicted pool utilizations 111
  - problem determination hints 145
  - problem determination tips 145
  - problems not directly related to snapshots 116
  - problems while working with Samba in 123
  - problems with locating a snapshot 116
  - problems with non-IBM disks 138
  - protocol service logs 3, 6, 8, 11, 13
  - quorum nodes in cluster 145
  - RAS events 151
  - Reliability 151
  - remote cluster name 101
  - remote command issues 75, 76
  - remote file system 100, 101
  - remote file system does not mount 100, 101
  - remote file system I/O failure 100
  - remote mount failure 103
  - replicated data 133
  - replicated metadata 133, 134
  - replication 132, 134
  - Requeuing message 124
  - requeuing of messages in AFM 124
  - restoring a snapshot 118
  - Samba 123
  - security issues 75
  - Serviceability 151
  - set up 38
  - setup issues 119
  - SMB server health 122
  - snapshot directory name conflict 118
  - snapshot problems 116
  - snapshot status errors 117
- GPFS (*continued*)
  - snapshot usage errors 117
  - some files are 'ill-placed' 111
  - stale inode data 121
  - storage pools 114, 115
  - strict replication 134
  - system load increase in night 146
  - timeout executing function error message 148
  - trace facility 34
  - tracing the mmpmon command 120
  - TSM error messages 116
  - UID mapping 100
  - unable to access disks 131
  - unable to determine if a file system is mounted 108
  - unable to start 73
  - underlying disk subsystem failures 127
  - understanding Persistent Reserve 138
  - unmount failure 104
  - unused underlying multipath device 141
  - usage errors 111, 114
  - using mmpmon 119
  - value to large failure 143
  - value to large failure while creating a file 143
  - varyon problems 137
  - volume group 137
  - volume group on each node 137
  - Windows file system 147
  - Windows issues 92, 93
  - working with Samba 123
- GPFS cluster
  - problems adding nodes 87
  - recovery from loss of GPFS cluster configuration data files 77
- GPFS cluster data
  - backup 78
  - locked 77
- GPFS cluster data files storage 77
- GPFS command
  - failed 89
  - return code 89
  - unsuccessful 89
- GPFS commands
  - unsuccessful 89
- GPFS configuration data 109
- GPFS daemon 75, 79, 80, 95, 105
  - crash 83
  - fails to start 80
  - went down 20, 83
  - will not start 79
- GPFS daemon went down 83
- GPFS failure
  - network failure 84
- GPFS GUI logs 41
- GPFS is not using the underlying multipath device 141
- GPFS kernel extension 79
- GPFS local node failure 102
- GPFS log 1, 2, 80, 81, 83, 95, 99, 100, 101, 102, 103, 104, 105, 167
- GPFS messages 173
- GPFS modules
  - cannot be loaded 79
  - unable to load on Linux 79
- GPFS problems 73, 95, 127
- GPFS startup time 2
- GPFS trace facility 34
- GPFS Windows SMB2 protocol (CIFS serving) 93
- gpfs.snap 26

- gpfs.snap command 167
  - data always gathered for a master snapshot 25
  - data always gathered on AIX 24
  - data always gathered on all platforms 23
  - data always gathered on Linux 25
  - data always gathered on Windows 25
  - using 23
- grep command 19
- Group Services
  - verifying quorum 81
- GROUP\_ID attribute 55, 56
- GUI
  - logs 41
- GUI logs 41

## H

- hard loop ID 75
- HDFS
  - transparency log 8
- hints and tips for GPFS problems 145
- Home and .ssh directory ownership and permissions 92

## I

- I/O failure
  - remote file system 100
- I/O hang 85
- I/O operations slow 21
- IBM Spectrum Scale 46
  - /tmp/mmfs directory 146
  - automount fails to mount on Linux 98
  - abnormal termination in mmpmon 120
  - active file management 148
  - AIX 99
  - AIX platform 24
  - application calls 78
  - application program errors 91, 92
  - authentication issues 93
  - authentication on Linux 28
  - authorization issues 76
  - automated 63
  - automount fails to mount on AIX 99
  - automount failure 99
  - automount failure in Linux 98
  - Automount file system 98
  - Automount file system will not mount 98
  - CES tracing
    - debug data collection 13
  - checking Persistent Reserve 139
  - cipherList option has not been set properly 103
  - clearing a leftover Persistent Reserve reservation 139
  - client nodes 103
  - cluster configuration
    - issues 77, 78
    - recovery 77
  - cluster crash 74
  - cluster data
    - backup 78
  - cluster name 101
  - cluster overload
    - detection 68
  - cluster state information 43, 44, 45, 46
  - command 43
  - commands 43, 44, 45, 46
  - connectivity problems 76

- IBM Spectrum Scale (*continued*)
  - contact node address 101
  - contact nodes down 102
  - core dumps 38
  - corrupted data integrity 124
  - creating a file 143
  - data always gathered 23
  - data gathered 24, 25, 27, 28
    - Object on Linux 27
  - data gathered for CES on Linux 28
  - data gathered for core dumps on Linux 30
  - data gathered for performance 29
  - data gathered for protocols on Linux 26, 27, 28, 29, 30
  - data gathered for SMB on Linux 26
  - data integrity may be corrupted 124
  - deadlock breakup
    - on demand 67
  - deadlock detection 63
  - deadlocks 63, 65, 66, 67, 68
    - automated data collection 65
  - determining the health of integrated SMB server 122
  - disaster recovery issues 88
  - discrepancy between GPFS configuration data and the on-disk data for a file system 109
  - disk accessing commands fail to complete 138
  - disk connectivity failure 135
  - disk failure 135
  - disk information commands 49, 50, 51, 58, 59, 61
  - disk media failure 133, 134
  - disk media failures 132
  - disk recovery 135
  - displaying NSD information 128
  - dumps 1
  - encryption issues 143
  - encryption rules 143
  - error creating internal storage 147
  - error encountered while creating and using NSD disks 127
  - error log 20, 21
  - error message for file system 97, 98
  - error message "Function not implemented" 100
  - error numbers 98
  - error numbers for GPFS application calls 109
  - error numbers specific to GPFS application calls 98, 124, 135
  - Error numbers specific to GPFS application calls 107
  - error numbers specific to GPFS application calls when data integrity may be corrupted 124
  - error numbers when a file system mount is unsuccessful 98
  - errors associated with filesets 109
  - errors associated with policies 109
  - errors associated with storage pools 109
  - errors encountered 115
  - errors encountered while restoring a snapshot 118
  - errors encountered with filesets 114
  - errors encountered with policies 112
  - errors encountered with storage pools 115
  - failure group considerations 106
  - failures using the mmbackup command 116
  - file placement optimizer 148
  - file system commands 49, 50, 51, 58, 59, 61
  - file system does not mount 100
  - file system fails to mount 95
  - file system fails to unmount 104
  - file system forced unmount 105
  - file system is forced to unmount 107

## IBM Spectrum Scale *(continued)*

- file system is known to have adequate free space 110
- file system is mounted 108
- file system manager appointment fails 109
- file system manager failures 109
- file system mount problems 97, 98
- file system mount status 108
- file system mounting on wrong drive 147
- file systems manager failure 108
- filesets usage errors 113
- GPFS cluster security configurations 101
- GPFS commands unsuccessful 90
- GPFS daemon does not start 82
- GPFS daemon issues 79, 80, 81, 82, 83
- GPFS declared NSD is down 130
- GPFS disk issues 85, 127
- GPFS down on contact nodes 102
- GPFS error message 97
- GPFS error messages 108, 117
- GPFS error messages for disk media failures 135
- GPFS error messages for file system forced unmount problems 107
- GPFS error messages for file system mount status 108
- GPFS error messages for mmbackup errors 116
- GPFS failure
  - network issues 84
- GPFS file system issues 95
- GPFS has declared NSDs built on top of AIX logical volume as down 136
- GPFS is not running on the local node 102
- GPFS modules
  - unable to load on Linux 79
- gpfs.snap 23, 24, 25
  - gpfs.snap command 25
  - Linux platform 25
- gpfs.snap command
  - usage 23
- guarding against disk failures 132
- GUI logs 41
- hang in mmpmon 120
- HDFS transparency log 8
- hints and tips for problem determination 145
- hosts file issue 74
- incorrect output from mmpmon 120, 151
- installation and configuration issues 73, 74, 77, 78, 79, 80, 81, 82, 83, 85, 87, 88, 89, 90, 91, 92
- key rewrap 144
- log 2
- logical volume 145
- logical volumes are properly defined for GPFS use 136
- logs 1
- lsof command 50
- manually disabling Persistent Reserve 140
- manually enabling Persistent Reserve 140
- master log file 2
- master snapshot 25
- message 6027-648 147
- message referring to an existing NSD 130
- message requeuing in AFM 124
- message severity tags 171
- messages 173
- mmafmctl Device getstate 43
- mmapplypolicy -L 0 command 52
- mmapplypolicy -L 1 command 52
- mmapplypolicy -L 2 command 53
- mmapplypolicy -L 3 command 54
- mmapplypolicy -L 4 command 55

## IBM Spectrum Scale *(continued)*

- mmapplypolicy -L 5 command 55
- mmapplypolicy -L 6 command 56
- mmapplypolicy -L command 52, 53, 54, 55, 56
- mmapplypolicy command 51
- mmdumpperfdata command 31
- mmfileid command 59
- MMFS\_DISKFAIL 20
- MMFS\_ENVIRON
  - error log 20
- MMFS\_FSSTRUCT error log 20
- MMFS\_GENERIC error log 20
- MMFS\_LONGDISKIO 21
- mmfsadm command 33
- mmlscluster command 44
- mmlsconfig command 45
- mmlsmount command 50
- mmrefresh command 45
- mmremotecluster command 101
- mmsdrrestore command 46
- mmwindisk command 58
- mount 98, 100
- mount failure 103
- mount failure as the client nodes joined before NSD servers 103
- mount failure for a file system 143
- mounting cluster does not have direct access to the disks 102
- multiple file system manager failures 108
- negative values occur in the 'predicted pool utilizations', 111
- newly mounted windows file system is not displayed 147
- NFS client 121
- NFS on Linux 27
- NFS problems 121
- NFS V4 issues 121
- no replication 134
- NO\_SPACE error 110
- NSD and underlying disk subsystem failures 127
- NSD creation fails 130
- NSD disk does not have an NSD server specified 102
- NSD server 103
- Object logs 6
- offline mmfscck command failure 147
- old NFS inode data 121
- operating system error logs 19, 20, 21, 22
- operating system logs 19, 20, 21, 22
- other problem determination tools 71
- partial disk failure 136
- performance issues 86
- permission denied error message 103
- permission denied failure 144
- Persistent Reserve errors 138
- physical disk association 145
- policies 111
- problem determination 145
- problems while working with Samba 123
- problems with locating a snapshot 116
- problems with non-IBM disks 138
- protocol service logs 3, 6, 8, 11, 13
- quorum loss 85
- quorum nodes 145
- quorum nodes in cluster 145
- remote cluster name 101
- remote cluster name does not match with the cluster name 101
- remote command issues 75, 76

- IBM Spectrum Scale (*continued*)
  - remote file system 100, 101
  - remote file system does not mount 100, 101
  - remote file system does not mount due to differing GPFS cluster security configurations 101
  - remote file system I/O fails with "Function not implemented" error 100
  - remote file system I/O failure 100
  - remote mounts fail with the "permission denied" error 103
  - remote node expelled from cluster 88
  - replicated metadata 134
  - replicated metadata and data 133
  - requeuing of messages in AFM 124
  - security issues 75
  - set up 38
  - setup issues while using mmpmon 119
  - SHA digest 61
  - snapshot directory name conflict 118
  - snapshot problems 116
  - snapshot status errors 117
  - snapshot usage errors 117
  - some files are 'ill-placed' 111
  - stale inode data 121
  - storage pools usage errors 114
  - strict replication 134
  - system load 146
  - timeout executing function error message 148
  - trace facility 34
  - trace reports 34
  - traces 1
  - tracing the mmpmon command 120
  - TSM error messages 116
  - UID mapping 100
  - unable to access disks 131
  - unable to determine if a file system is mounted 108
  - unable to resolve contact node address 101
  - understanding Persistent Reserve 138
  - unused underlying multipath device by GPFS 141
  - usage errors 111
  - value to large failure 143
  - volume group on each node 137
  - volume group varyon problems 137
  - Windows 25
  - Windows issues 92, 93
- IBM Spectrum Scale information units ix
- IBM Spectrum Scalecommand
  - mmafmctl Device getstate 43
- IBM Spectrum Scalemmdiag command 43
- ill-placed files 111, 115
- ILM
  - problems 109
- inode data
  - stale 121
- inode limit 22
- installation and configuration issues 73
- installation problems 73

## J

- junctions
  - deleting 113

## K

- KB\_ALLOCATED attribute 55, 56
- kdb 71

- KDB kernel debugger 71
- kernel module
  - mmfslinux 79
- kernel panic 85
- kernel threads
  - at time of system hang or panic 71
- key rewrap failure 144

## L

- Linux kernel
  - configuration considerations 74
  - crash dump facility 71
- Linux on z Systems 74, 168
- logical volume 136
  - location 145
- Logical Volume Manager (LVM) 132
- logical volumes 136
- logs 1
  - NFS, SMB, and Object logs 3
  - protocol service logs 3
- logsIBM Spectrum Scale
  - NFS logs 4
  - SMB service logs 3
- long waiters
  - increasing the number of inodes 86
- lspp command 167
- lslv command 145
- lsdf command 50, 105
- lspv command 137
- lsvg command 136
- lxtrace command 33, 34

## M

- manually enabling or disabling Persistent Reserve 140
- master GPFS log file 2
- master log file
  - creating 2
- maxblocksize parameter 97
- MAXNUMMPP 116
- memory shortage 20, 75
- message 6027-648 147
- message severity tags 171
- messages 173
  - 6027-1941 74
- metadata
  - replicated 133, 134
- MIGRATE rule 51, 54
- migration
  - file system will not mount 96
  - new commands will not run 90
- mmaddisk command 109, 114, 133, 136, 138
- mmaddnode command 87, 146
- mmafmctl command 124
- mmafmctl Device getstate command 43
- mmapplypolicy -L 0 52
- mmapplypolicy -L 1 52
- mmapplypolicy -L 2 53
- mmapplypolicy -L 3 54
- mmapplypolicy -L 4 55
- mmapplypolicy -L 5 55
- mmapplypolicy -L 6 56
- mmapplypolicy command 51, 111, 112, 115, 144
- mmauth command 61, 101
- mmbackup command 116

- mmccr command
  - failure 148
- mmchcluster command 75
- mmchconfig command 45, 80, 88, 103
- mmchdisk command 96, 106, 109, 114, 115, 127, 130, 131, 133, 135
- mmcheckquota command 21, 57, 92, 106
- mmchfs command 22, 78, 86, 90, 96, 98, 106, 123
- mmchnode command 146
- mmchnsd command 127
- mmchpolicy
  - issues with adding encryption policy 143
- mmcommon 98, 99
- mmcommon breakDeadlock 67
- mmcommon recoverfs command 109
- mmcommon showLocks command 77
- mmcrcluster command 45, 75, 80, 87, 146
- mmcrfs command 90, 123, 127, 138
- mmcrnsd command 127, 130
- mmcrsnapshot command 117, 118
- mmdefedquota command fails 147
- mmdeldisk command 109, 114, 133, 136
- mmdelfileset command 113
- mmdelfs command 134, 135
- mmdelnode command 87, 90
- mmdelnsd command 130, 134
- mmdelsnapshot command 117
- mmdf command 86, 110, 136
- mmdiag command 43
- mmdsh command 76
- mmdumpperfdata 31
- mmedquota command fails 147
- mmexpelnode command 46
- mmfileid command 59, 124, 133
- MMFS\_ABNORMAL\_SHUTDOWN
  - error logs 20
- MMFS\_DISKFAIL
  - error logs 20
- MMFS\_ENVIRON
  - error logs 20
- MMFS\_FSSTRUCT
  - error logs 20
- MMFS\_GENERIC
  - error logs 20
- MMFS\_LONGDISKIO
  - error logs 21
- MMFS\_QUOTA
  - error log 21
  - error logs 21, 57
- MMFS\_SYSTEM\_UNMOUNT
  - error logs 22
- MMFS\_SYSTEM\_WARNING
  - error logs 22
- mmfs.log 2, 80, 81, 83, 95, 99, 100, 101, 102, 103, 104, 105, 167
- mmfsadm command 33, 37, 81, 87, 124, 133
- mmfsck command 49, 95, 96, 114, 124, 134, 136
  - failure 147
- mmfsd 79, 80, 95, 105
  - will not start 79
- mmfslinux
  - kernel module 79
- mmgetstate command 43, 81, 89
- mmlock directory 77
- mmlsattr command 113
- mmlscluster command 44, 87, 101, 145
- mmlsconfig command 34, 45, 98
- mmlsdisk command 90, 95, 96, 106, 109, 127, 130, 132, 135, 168
- mmlsfileset command 113
- mmlsfs command 97, 133, 134, 167
- mmlsmgr command 33, 96
- mmlsmount command 50, 80, 91, 95, 105, 106, 127
- mmlsnsd command 57, 128, 136
- mmlspolicy command 112
- mmlsquota command 91, 92
- mmlsnapshot command 116, 117, 118
- mmmound command 49, 95, 106, 138
- mmpmon
  - abend 120
  - altering input file 119
  - concurrent usage 119
  - counters wrap 120
  - dump 120
  - hang 120
  - incorrect input 119
  - incorrect output 120
  - restrictions 119
  - setup problems 119
  - unsupported features 119
- mmpmon command 71, 119
  - trace 120
- mmquotaoff command 92
- mmquotaon command 92
- mmrefresh command 45, 96, 98
- mmremotecluster command 61, 101, 102
- mmremotefs command 98, 101
- mmrepquota command 92
- mmrestorefs command 117, 118, 119
- mmrestripefile command 112, 115
- mmrestripefs command 115, 133, 136
- mmrpldisk command 109, 114, 138
- mmsdrbackup 78
- mmsdrfs 78
- mmsdrrestore command 46
- mmsshutdown command 44, 46, 80, 81, 83, 98, 99
- mmsnapdir command 116, 118, 119
- mmstartup command 80, 98, 99
- mmtracectl command
  - generating GPFS trace reports 34
- mmumount command 105, 106, 136
- mmunlinkfileset command 113
- mmwindisk command 58
- mode of AFM fileset, changing 148
- MODIFICATION\_TIME attribute 55, 56
- module is incompatible 79
- mount
  - problems 103
- mount command 95, 96, 98, 134, 138
- mount failure 143
- mounting cluster 102
- Mounting file system
  - error messages 97
- Multi-Media LAN Server 1

## N

- network failure 84
- network problems 20
- NFS 26, 121
  - problems 121
- NFS client
  - with stale inode data 121

- NFS V4
  - problems 121
- no replication 134
- NO SUCH DIRECTORY error code 83
- NO SUCH FILE error code 83
- NO\_SPACE
  - error 110
- node
  - crash 169
  - hang 169
  - rejoin 104
- node crash 74
- node failure 85
- node reinstall 74
- nodes
  - cannot be added to GPFS cluster 87
- non-quorum node 145
- NSD 136
  - creating 130
  - deleting 130
  - displaying information of 128
  - extended information 129
  - failure 127
- NSD build 136
- NSD disks 102
  - creating 127
  - using 127
- NSD failure 127
- NSD server 102, 103, 104
- nsdServerWaitTimeForMount
  - changing 104
- nsdServerWaitTimeWindowOnMount
  - changing 104

## O

- Object 26
  - health 6
  - logs 6
- OpenSSH connection delays
  - Windows 93
- orphaned file 114

## P

- partitioning information, viewing 58
- performance 26, 75
- permission denied
  - remote mounts failure 103
- permission denied failure (key rewrap) 144
- Persistent Reserve
  - checking 139
  - clearing a leftover reservation 139
  - errors 138
  - manually enabling or disabling 140
  - understanding 138
- ping command 76
- PMR 169
- policies
  - DEFAULT clause 112
  - deleting referenced objects 112
  - errors 112
  - file placement 112
  - incorrect file placement 112
  - LIMIT clause 111
  - long runtime 112

- policies (*continued*)
  - MIGRATE rule 111
  - problems 109
  - rule evaluation 112
  - usage errors 111
  - verifying 51
- policy file
  - detecting errors 52
  - size limit 112
  - totals 52
- policy rules
  - runtime problems 113
- POOL\_NAME attribute 55, 56
- possible GPFS problems 73, 95, 127
- predicted pool utilization
  - incorrect 111
- primary NSD server 103
- problem
  - locating a snapshot 116
  - not directly related to snapshot 116
  - snapshot 116
    - snapshot directory name 118
    - snapshot restore 118
    - snapshot status 117
    - snapshot usage 117
    - snapshot usage errors 117
- problem determination
  - cluster state information 43
  - remote file system I/O fails with the "Function not implemented" error message when UID mapping is enabled 100
  - tools 49
  - tracing 34
- Problem Management Record 169
- problems
  - configuration 73
  - installation 73
  - mmbackup 116
- problems running as administrator, Windows 93
- protocol (CIFS serving), Windows SMB2 93
- protocol authentication log 9
- Protocols 93

## Q

- quorum 81, 145
  - disk 85
  - loss 85
- quorum node 145
- quota
  - cannot write to quota file 106
  - denied 91
  - error number 78
- quota files 57
- quota problems 21

## R

- RAID controller 132
- rcp command 75
- read-only mode mount 49
- recovery
  - cluster configuration data 77
- recovery log 85
- recreation of GPFS storage file
  - mmchcluster -p LATEST 77

- Reliability 151
- remote command problems 75
- remote file copy command
  - default 75
- remote file system
  - mount 101
- remote file system I/O fails with "Function not implemented"
  - error 100
- remote mounts fail with permission denied 103
- remote node
  - expelled 88
- remote node expelled 88
- remote shell
  - default 75
- removing the setuid bit 83
- replicated
  - metadata 134
- replicated data 133
- replicated metadata 133
- replication 114
  - of data 132
- replication, none 134
- reporting a problem to IBM 33
- resetting of setuid/setgits at AFM home 148
- restricted mode mount 49
- rpm command 167
- rsh
  - problems using 75
- rsh command 75, 89

## S

- Samba
  - client failure 123
- scp command 76
- Secure Hash Algorithm digest 61
- Serviceability 151
- serving (CIFS), Windows SMB2 protocol 93
- set up
  - core dumps 38
- setuid bit, removing 83
- setuid/setgid bits at AFM home, resetting of 148
- severity tags
  - messages 171
- SHA digest 61, 101
- shared segments 81
  - problems 82
- SMB 26
- SMB on Linux 26
- SMB server 122
- SMB service
  - log 3
  - logs 4
- SMB2 protocol (CIFS serving), Windows 93
- snapshot
  - directory name conflict 118
  - error messages 116, 117
  - invalid state 117
  - restoring 118
  - status error 117
  - usage error 117
  - valid 116
- snapshot problems 116
- ssh command 76
- steps to follow
  - GPFS daemon does not come up 80

- storage pools
  - deleting 112, 115
  - errors 115
  - failure groups 114
  - problems 109
  - slow access time 115
  - usage errors 114
- strict replication 134
- subnets attribute 88
- support for troubleshooting
  - contacting IBM support center 167
    - how to contact IBM support center 169
  - information to be collected before contacting IBM support center 167
- syslog facility
  - Linux 19
- syslogd 100
- system load 146
- system snapshots 23
- system storage pool 111, 114
- System z 74, 168

## T

- threads
  - tuning 75
  - waiting 87
- Tivoli Storage Manager server 116
- trace
  - active file management 35
  - allocation manager 35
  - basic classes 35
  - behaviorals 37
  - byte range locks 35
  - call to routines in SharkMsg.h 36
  - checksum services 35
  - cleanup routines 35
  - cluster security 37
  - concise vnop description 37
  - daemon routine entry/exit 35
  - daemon specific code 37
  - data shipping 35
  - defragmentation 35
  - dentry operations 35
  - disk lease 35
  - disk space allocation 35
  - DMAPI 35
  - error logging 35
  - events exporter 35
  - file operations 35
  - file system 35
  - generic kernel vfs information 36
  - inode allocation 36
  - interprocess locking 36
  - kernel operations 36
  - kernel routine entry/exit 36
  - low-level vfs locking 36
  - mailbox message handling 36
  - malloc/free in shared segment 36
  - miscellaneous tracing and debugging 37
  - mmpmon 36
  - mnode operations 36
  - mutexes and condition variables 36
  - network shared disk 36
  - online multinode fsck 36
  - operations in Thread class 37
  - page allocator 36



- trace (*continued*)
  - parallel inode tracing 36
  - performance monitors 36
  - physical disk I/O 35
  - physical I/O 36
  - pinning to real memory 36
  - quota management 36
  - rdma 36
  - recovery log 36
  - SANergy 36
  - scsi services 37
  - shared segments 37
  - SMB locks 37
  - SP message handling 37
  - super operations 37
  - tasking system 37
  - token manager 37
  - ts commands 35
  - vdisk 37
  - vdisk debugger 37
  - vdisk hospital 37
  - vnode layer 37
- trace classes 35
- trace facility 34, 35
  - mmfsadm command 33
- trace level 37
- trace reports, generating 34
- traversing a directory that has not been cached 148
- troubleshooting
  - CES 11
    - disaster recovery issues 88
    - setup problems 88
  - GUI logs 41
- troubleshooting errors 92
- troubleshooting Windows errors 92
- TSM client 116
- TSM server 116
  - MAXNUMMP 116
- tuning 75

## U

- UID mapping 100
- umount command 105, 106
- unable to start GPFS 81
- underlying multipath device 141
- understanding, Persistent Reserve 138
- unsuccessful GPFS commands 89
- usage errors
  - policies 111
- useNSDserver attribute 135
- USER\_ID attribute 55, 56
- using the gpfs.snap command 23

## V

- v 75
- value too large failure 143
- varyon problems 137
- varyonvg command 138
- viewing disks and partitioning information 58
- volume group 137

## W

- Windows 92
  - data always gathered 25
  - file system mounted on the wrong drive letter 147
  - gpfs.snap 25
  - Home and .ssh directory ownership and permissions 92
  - mounted file systems, Windows 147
  - OpenSSH connection delays 93
  - problem seeing newly mounted file systems 147
  - problem seeing newly mounted Windows file systems 147
  - problems running as administrator 93
  - Windows 147
- Windows issues 92
- Windows SMB2 protocol (CIFS serving) 93

## Z

- z Systems 74, 168







Product Number: 5725-Q01  
5641-GPF  
5725-S28

Printed in USA

GA76-0443-06

