

CXFS<sup>™</sup> Software Installation and  
Administration Guide

007-4016-002

## CONTRIBUTORS

Written by Lori Johnson

Edited by Rick Thompson

Illustrated by Chris Wengelski

Engineering contributions by François Barbou des Places, Ken Beck, Dean Jansa, Kent Koeninger, Ted Kline, Steve Lord, LaNet Merrill, John Relph, Elaine Robinson, Dean Roehrich, Wesley Smith, Paddy Sreenivasan

## COPYRIGHT

© 1999 Silicon Graphics, Inc. All Rights Reserved. This document or parts thereof may not be reproduced in any form unless permitted by contract or by written permission of Silicon Graphics, Inc.

## LIMITED AND RESTRICTED RIGHTS LEGEND

Use, duplication, or disclosure by the Government is subject to restrictions as set forth in the Rights in Data clause at FAR 52.227-14 and/or in similar or successor clauses in the FAR, or in the DOD, DOE or NASA FAR Supplements. Unpublished rights reserved under the Copyright Laws of the United States. Contractor/manufacturer is Silicon Graphics, Inc., 1600 Amphitheatre Pkwy., Mountain View, CA 94043-1351.

## TRADEMARKS

Silicon Graphics, IRIS, and IRIX are registered trademarks and SGI, CXFS, IRIS FailSafe, IRIX Interactive Desktop, Origin, and the Silicon Graphics logo are trademarks of Silicon Graphics, Inc.

Adaptec is a trademark of Adaptec, Inc. BROCADE is a trademark of Brocade Communication Systems, Inc. CIPRICO is a trademark of Ciprico Inc. CLARiiON is a trademark of Data General Corporation. Emulex is a trademark of Emulex Corporation. FLEXlm is a trademark of GLOBETROTTER Software, Inc. Legato NetWorker is a trademark of Legato Systems, Inc. NFS and Sun are trademarks of Sun Microsystems, Inc. Netscape is a trademark of Netscape Communications Corporation. Prisa is a trademark of Prisa Networks, Inc. QLogic is a trademark of QLogic Corporation. UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company Limited. Windows is a trademark of Microsoft Corporation. VERITAS is a trademark of VERITAS Software Corporation.

---

## **New Features in This Guide**

This update contains additional troubleshooting information and instructions for unmounting and remounting filesystems with the CLI. It has been reorganized to make the tasks of installation and configuration clearer.



---

## Record of Revision

<b>Version</b>	<b>Description</b>
001	September 1999 Supports the CXFS 1.1 product in the IRIX 6.5.6 release.
002	October 1999 Supports the CXFS 1.1 product in the IRIX 6.5.6 release.



---

# Contents

<b>About This Guide</b>	<b>xvii</b>
Related Publications	xvii
Obtaining Publications	xvii
Conventions	xvii
Reader Comments	xviii
<b>1. Introduction to CXFS</b>	<b>1</b>
Cluster Environment	1
Cluster Membership and Quorum	4
Node Weights	5
Metadata Client/Server Model	6
System View	7
Comparison of Networked and CXFS Filesystems	8
Networked Filesystems	8
CXFS Filesystems	8
XFS Features Present in CXFS	10
When to Use CXFS	11
When to Use XFS	11
Hardware and Software Requirements and Support	12
CXFS Cluster Manager Tools	13
CXFS Cluster Manager GUI Overview	13
<b>CXFS Manager</b>	<b>14</b>
<b>CXFS Cluster View</b>	<b>16</b>
Starting the GUI	19

Launching the <b>CXFS Cluster View</b>	19
Launching the <b>CXFS Manager</b>	20
Viewing Cluster Item Details	21
Performing Tasks	21
Using Tasksets	22
CXFS Cluster Manager CLI Overview	22
Getting Help	23
Completing Actions and Cancelling	23
Invoking a Shell from within the CLI	23
<b>2. Installing CXFS Software and Preparing the System</b>	<b>25</b>
Installing Required Software	25
Configuring System Files	29
Configuring <code>/etc/services</code>	29
Configuring <code>/etc/config/cad.options</code>	29
Configuring <code>/etc/config/fs2d.options</code>	30
<i>(Optional)</i> Configuring <code>/.rhosts</code>	33
<i>(Optional)</i> Configuring for Automatic Restart	34
Configuring Network Interfaces	34
Configuring the Serial Ports	37
Rebooting the System	37
<b>3. Initial Configuration of the Cluster</b>	<b>39</b>
Configuring with the GUI	39
Configuring with the CLI	42
<b>4. Reference to GUI Tasks</b>	<b>51</b>
Initial Guided Configuration Tasksets	52

Set up a New Cluster with the GUI . . . . .	52
Set Up a New Filesystem with the GUI . . . . .	54
Node Tasks . . . . .	55
Define a Node with the GUI . . . . .	55
Add/Remove Nodes in the Cluster with the GUI . . . . .	61
Reset a Node with the GUI . . . . .	61
Modify a Node with the GUI . . . . .	62
Delete a Node with the GUI . . . . .	63
Display a Node with the GUI . . . . .	63
Cluster Tasks . . . . .	64
Define a Cluster with the GUI . . . . .	64
Modify the Cluster Administrator Notification with the GUI . . . . .	65
Delete a Cluster with the GUI . . . . .	65
Display a Cluster with the GUI . . . . .	66
CXFS Cluster Services Tasks . . . . .	66
Start CXFS Cluster Services with the GUI . . . . .	66
Stop CXFS Cluster Services with the GUI . . . . .	67
Set Log Configuration with the GUI . . . . .	68
Display Log Group Definitions with the GUI . . . . .	69
Configure Log Groups with the GUI . . . . .	69
Filesystem Tasks . . . . .	70
Define a Filesystem Metadata Server with the GUI . . . . .	70
Mount a Filesystem with the GUI . . . . .	72
Unmount a Filesystem with the GUI . . . . .	72
Modify a Filesystem with the GUI . . . . .	72
Delete a Filesystem with the GUI . . . . .	73
Diagnostic Task . . . . .	73
Test Connectivity with the GUI . . . . .	74

<b>5. Reference to CLI Tasks</b>	<b>75</b>
Set Configuration Defaults with the CLI	76
Node Tasks	77
Define a Node with the CLI	77
Modify a Node with the CLI	78
Delete a Node with the CLI	79
Display a Node with the CLI	82
Reset a Node with the CLI	83
Cluster and Filesystem Tasks	84
Define a Cluster with the CLI	84
Modify a Cluster and Define and Mount Filesystems with the CLI	85
Unmount a Filesystem with the CLI	88
Remount a Defined Filesystem with the CLI	88
Delete a Cluster with the CLI	89
Display a Cluster with the CLI	90
CXFS Cluster Services Tasks	91
Start CXFS Cluster Services with the CLI	91
Stop CXFS Cluster Services with the CLI	92
Set Log Configuration with the CLI	93
Display Log Group Definitions with the CLI	93
Configure Log Groups with the CLI	93
Modify Log Groups with the CLI	94
Diagnostic Tasks	95
Test Network Connectivity with the CLI	95
<b>6. Administration and Maintenance</b>	<b>97</b>
Discovering the Metadata Server for a Filesystem	97
Volume Management	101

Disk Management . . . . .	.101
Backups . . . . .	.101
NFS . . . . .	.102
Quotas . . . . .	.102
RealTime . . . . .	.102
SAMBA . . . . .	.102
Filesystem Maintenance . . . . .	.102
Mounting Filesystems . . . . .	.103
Unmounting Filesystems . . . . .	.104
Growing Filesystems . . . . .	.104
Dump and Restore . . . . .	.104
Administrative Shutdown and CXFS Cluster Services . . . . .	.105
Cluster Database Backup and Restore . . . . .	.105
Configuration State Checker Flags . . . . .	.105
<b>7. Monitoring Cluster and Node Status . . . . .</b>	<b>.107</b>
Monitoring System Status with the GUI . . . . .	.107
Node Status . . . . .	.108
Monitoring Node Status with the GUI . . . . .	.108
Querying Node Status with the CLI . . . . .	.109
Pinging the System Controller with the CLI . . . . .	.109
Monitoring Reset Serial Line with the CLI . . . . .	.109
Cluster Status . . . . .	.109
Querying Cluster Status with the GUI . . . . .	.110
Querying Cluster Status with the CLI . . . . .	.110
<b>8. Troubleshooting . . . . .</b>	<b>.111</b>
Determining Where Problems Lie . . . . .	.111

Getting a Dump of the Cluster Configuration Database . . . . .	.113
Avoiding GUI Problems . . . . .	.113
Restarting CXFS Cluster Services . . . . .	.113
Corruption of the Root XFS Filesystem . . . . .	.114
Rebooting . . . . .	.114
Performing Temporary Maintenance on a Node . . . . .	.114
Reboot without Rejoining the Cluster . . . . .	.115
Clearing the Cluster Configuration Database . . . . .	.115
Avoiding Problems with Log Files . . . . .	.116
Error Messages . . . . .	.116
License Error . . . . .	.116
I/O Errors . . . . .	.116
IP Address Error . . . . .	.117
SYSLOG Errors . . . . .	.117
cli Errors . . . . .	.120
clconfd Errors . . . . .	.120
crsd Errors . . . . .	.123
cmond Errors . . . . .	.124
fs2d Errors . . . . .	.125
General Messages . . . . .	.126
Log File Errors . . . . .	.127
cad Messages . . . . .	.128
cli Messages . . . . .	.130
crsd Errors . . . . .	.131
fs2d Errors . . . . .	.132
<b>Index . . . . .</b>	<b>.135</b>

---

## Figures

<b>Figure 1-1</b>	Example of a Four-Node Cluster using Reset . . . . .	3
<b>Figure 1-2</b>	Example of a Two-Node Cluster using Reset . . . . .	4
<b>Figure 1-3</b>	<b>CXFS Manager Overview</b> screen . . . . .	15
<b>Figure 1-4</b>	<b>CXFS Cluster View</b> Screen . . . . .	16
<b>Figure 1-5</b>	Detailed Cluster View . . . . .	18
<b>Figure 3-1</b>	CXFS Manager . . . . .	40
<b>Figure 4-1</b>	Example Node Definition . . . . .	59
<b>Figure 4-2</b>	Example System Controller Settings . . . . .	60
<b>Figure 6-1</b>	Selecting the Filesystem Name . . . . .	99
<b>Figure 6-2</b>	Displaying the Metadata Server for a Filesystem . . . . .	100



---

## Tables

<b>Table 4-1</b>	Default Log Files in /var/cluster/ha/log . . . . .	69
<b>Table 8-1</b>	SYSLOG Error Message Format . . . . .	.119
<b>Table 8-2</b>	Log Error Message Format . . . . .	.128



---

## About This Guide

This publication documents CXFS 1.1 running on a storage area network (SAN). It assumes that you are already familiar with the XFS file system documented in *IRIX Admin: Disks and Filesystems* and that you have access to the *XVM Volume Manager Administrator's Guide*.

## Related Publications

The following documents contain additional information:

- *FDDIXPress Administration Guide*
- *IRIX Installation Instructions*
- *IRISconsole Administrator's Guide*
- *IRIX Admin: Disks and Filesystems*
- *NIS Administrator's Guide*
- *Personal System Administration Guide*
- *XVM Volume Manager Administrator's Guide*

## Obtaining Publications

To order SGI documentation, go to the SGI Technical Publications Library at <http://techpubs.sgi.com>. Find the title that you want and choose `order` to get the ordering information page for that document.

## Conventions

The following conventions are used throughout this document:

Convention	Meaning
<code>command</code>	This fixed-space font denotes literal items such as commands, files, routines, path names, signals, messages, and programming language structures.
<i>variable</i>	Italic typeface denotes variable entries and words or concepts being defined.
<b>user input</b>	This bold, fixed-space font denotes literal items that the user enters in interactive sessions. Output is shown in nonbold, fixed-space font.
[ ]	Brackets enclose optional portions of a command or directive line.
...	Ellipses indicate that a preceding element can be repeated.

## Reader Comments

If you have comments about the technical accuracy, content, or organization of this document, please tell us. Be sure to include the title and part number of the document with your comments.

You can contact us in any of the following ways:

- Send e-mail to the following address:  
`techpubs@sgi.com`
- Send a fax to the attention of "Technical Publications" at fax number +1 650 932 0801.
- Use the Feedback option on the Technical Publications Library World Wide Web page:<http://techpubs.sgi.com>
- Call the Technical Publications Group, through the Technical Assistance Center, using the following number: 1 800 800 4SGI
- Send mail to the following address:  
Technical Publications  
SGI  
1600 Amphitheatre Pkwy., M/S 535

Mountain View, California 94043-1351

We value your comments and will respond to them promptly.



## Introduction to CXFS

CXFS is a shared XFS filesystem that allows groups of computers to coherently share large amounts of data while maintaining high performance. It runs on storage area network (SAN) disks, such as Fibre Channel, in a cluster environment.

This introduction discusses the following:

- "Cluster Environment", page 1
- "Comparison of Networked and CXFS Filesystems", page 8
- "XFS Features Present in CXFS", page 10
- "When to Use CXFS", page 11
- "When to Use XFS", page 11
- "Hardware and Software Requirements and Support", page 12
- "CXFS Cluster Manager Tools", page 13

## Cluster Environment

A *cluster* is a collection of 2 to 16 cluster nodes coupled to each other by a private TCP/IP network (one that is accessible only by administrators and not by users).

A *cluster node* is a single IRIX image. Usually, a node is an individual computer.

---

**Note:** The term *node* is used in this guide for brevity; it does not have the same meaning as a node in an Origin system.

---

The nodes are connected to a storage area network (SAN) consisting of a number of disks.

To prevent the possibility of data corruption in certain rarely seen error situations, SGI recommends that you use the hardware required for remote reset.

For example, it is possible that one node in the cluster could be unable to communicate with other nodes in the cluster (due to a software or hardware failure) but still be able to access shared disks, despite the fact that the cluster does not see

this node as an active member. In this case, the reset hardware will allow one of the other nodes to forcibly prevent the failing node from accessing the disk at the instant the error is detected and prior to recovery from the node's departure from the cluster, ensuring no further activity from this node.

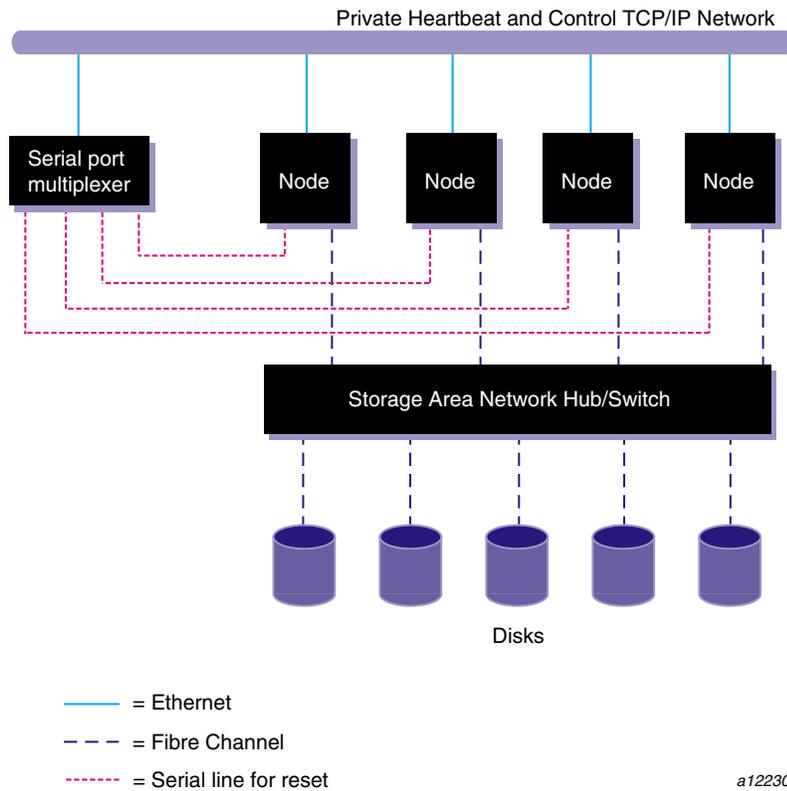
The remote reset connections take the following forms:

- Clusters of two nodes can be directly connected with serial lines.
- Clusters of three or more nodes should be connected with an Ethernet multiplexer or SCSI multiplexer. Each node is defined to have an *owner host*, which is the node that has the ability to reset it.

The nodes that have been defined using the CXFS Cluster Manager, and are therefore available for clustering, are known as the *pool*. Each node in the pool stores a copy of the cluster database. All nodes in the cluster are by definition part of the pool, but all nodes in the pool are not necessarily part of the cluster.

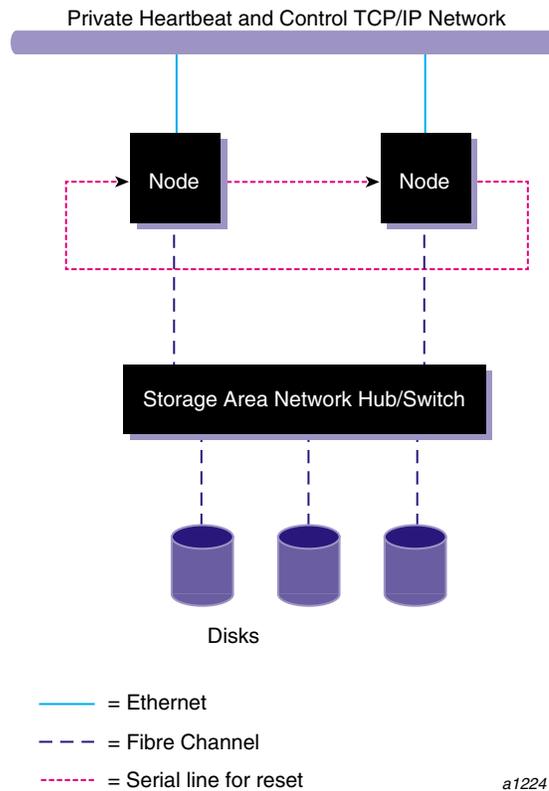
A cluster is identified by a simple name; this name must be unique within the pool. There can be only one cluster within a pool.

Figure 1-1 shows an example of the CXFS hardware components for a four-node cluster using the reset capability and an Ethernet serial port multiplexer. (The reset capability is optional.) A SCSI multiplexer can be used in place of the Ethernet multiplexer with a similar configuration.



**Figure 1-1** Example of a Four-Node Cluster using Reset

Figure 1-2 shows a two-node cluster using direct serial lines for reset. (The reset capability is optional.)



a12241

**Figure 1-2** Example of a Two-Node Cluster using Reset

With CXFS, each node has direct access to the logical volumes on the SAN, which enhances performance. Each node is more tightly coupled to the logical volumes than in a network file system (NFS) solution. To ensure data integrity, each filesystem has a metadata server; see "Metadata Client/Server Model", page 6. CXFS software provides the infrastructure that allows the metadata server to decide cluster membership and access questions.

## Cluster Membership and Quorum

The *cluster membership* is the group of nodes that are actively sharing resources in the cluster, which may be a subset of the nodes defined in a cluster. During the boot

process, a node applies for membership in the cluster. Once accepted by the existing membership, the node can share the resources of the cluster.

CXFS prevents problems arising because of possible network partitioning by using a *cluster membership quorum*, which is calculated based on the combined weight of nodes attempting to participate in the membership compared to the total weight of all nodes defined in the cluster. The membership exists as long as a quorum of nodes are included. The quorum is based on the weight of each node, as assigned during configuration. For the initial cluster membership, a quorum requires a majority of the total defined weight for the cluster. For an existing membership, a quorum requires half the total weight of all nodes defined in the cluster. For more information, see "Node Weights".

The membership monitors itself with normal messaging and heartbeating. If a failure occurs, the offending nodes are removed from the membership and are reset to prevent further access to shared resources by these nodes. A node that discovers it has been eliminated from the membership (due to a communication failure) may also forcibly shut itself down if it is not first reset. When a failed node is rebooted, it can rejoin the membership.

The number of nodes possible in the membership can be changed to either expand the cluster to include new nodes or to remove nodes that have left the membership. Removal of a node that is down and will be unavailable for an extended period reduces the weight required to form a quorum.

For example:

- Nodes in the pool, each with a weight of 1: A, B, C, D, E, F
- Nodes defined for the cluster: A, B, C, D, E
- Number of nodes needed for a quorum of the cluster: three nodes
- Nodes that are currently down: B
- Nodes that are therefore in the active cluster membership for the cluster: A, C, D, E

## Node Weights

Nodes have a default weight of 1. When all nodes have a weight of 1, a quorum calculation essentially becomes a majority of the number of nodes.

In cases where a cluster consists of one or more CXFS metadata servers and multiple CXFS client-only nodes, you may want to configure the node weights such that the

quorum consists of only the metadata servers, regardless of the state of the clients. In this case, you should define a weight of 1 for the metadata servers and 0 for the clients. If clients are weighted 0, the clients will be unable to form a membership by themselves. In this scenario, a membership will always require a quorum of metadata servers.

---

**Note:** Node weight values other than 0 or 1 are not recommended.

---

## Metadata Client/Server Model

*Metadata* is information that describes a file, such as the file's name, size, location, and permissions. Metadata tends to be small, usually about 512 bytes per file in XFS. This differs from the *data*, which is the contents of the file. The data may be many Mbytes or Gbytes in size.

While the data moves directly between the hosts and disks (peer-to-disk), the metadata is coordinated through a metadata server.

For each CXFS filesystem, you must designate one of the nodes in the cluster to act as a central clearing-house, known as a *metadata server*. The other nodes in the cluster are called *metadata clients*.

---

**Note:** Do not confuse *metadata server* and *metadata client* with the traditional data-path client/server model used by networked filesystems. Only the metadata passes through the metadata server; the data is passed directly to and from disk on the metadata client.

---

The metadata server must perform cluster-coordination functions such as the following:

- Metadata logging
- File locking
- Buffer coherency

All CXFS requests for metadata are routed over a TCP/IP network and through a metadata server, and all changes to metadata are sent to the metadata server. The metadata server uses the advanced XFS journal features to log the metadata changes. Because the size of the metadata is typically small, the bandwidth of a fast Ethernet local area network (LAN) is generally sufficient for the metadata traffic.

The operations to the CXFS metadata server are typically infrequent compared with the data operations directly to the disks. For example, opening a file causes a request for the file information from the metadata server. After the file is open, a process can usually read and write the file many times without additional metadata requests. When the file size or other metadata attributes for the file change, this triggers a metadata operation.

The following rules apply:

- Any node in the cluster can act as a metadata server.
- A single node in the cluster can act as the metadata server for multiple filesystems at once.
- There can be multiple nodes acting as metadata servers, each with different sets of filesystems. However, a given filesystem has a single metadata server on a single node.
- Although you can configure multiple nodes to be the metadata server for a given filesystem, only the first of these nodes to mount the filesystem will actually become the metadata server. Although the list of metadata servers for a given filesystem is ordered, it is impossible to predict which node will become the metadata server because of network latencies and other unpredictable delays.
- Once a node has become the metadata server for a filesystem, the only way to change to a different node is to unmount the filesystem everywhere and remount with a different metadata server on a different node.
- If the metadata server for a filesystem goes down while there are active clients, all of the clients will be forced out of the filesystem.

## System View

CXFS provides a single-system view of the filesystems; each host in the SAN has equally direct access to the shared disks and common path names to the files. CXFS lets you scale the shared-file system performance as needed by adding disk channels and storage to increase the direct host-to-disk bandwidth. The CXFS shared-file performance is not limited by LAN speeds or a bottleneck of data passing through a centralized file server. It combines the speed of near-local disk access with the flexibility, scalability, and reliability of clustering.

## Comparison of Networked and CXFS Filesystems

Networked filesystems and CXFS filesystems perform many of the same functions, but with important performance and functional differences noted here.

### Networked Filesystems

Accessing remote files over local area networks (LANs) can be significantly slower than accessing local files. The network hardware and software introduces delays that tend to significantly lower the *transaction rates* (I/O per second) and the *bandwidth* (Mbytes per second). These delays are difficult to avoid in the client-server architecture of LAN-based networked filesystems. The delays stem from the limits of the LAN bandwidth and latency and the shared path through the data server.

LAN bandwidths force an upper limit for the speed of most existing shared filesystems. For example, 10- to 1000-Mbit Ethernets have limits of 1 to 100 Mbytes per second, respectively. This is one to several orders of magnitude slower than the bandwidth possible across multiple disk channels to local or shared disks. The layers of network protocols and server software also tend to limit the bandwidth rates.

A shared file server can be a bottleneck for performance when multiple clients wait their turns for data, which must pass through the centralized file server. For example, network file system (NFS) and Samba servers read data from disks attached to the server, copy the data into UDP/IP or TCP/IP packets, and then send it over a LAN to a client system. When many clients access the sever simultaneously, the server's responsiveness degrades.

### CXFS Filesystems

CXFS is a clustered XFS filesystem that allows for logical file sharing, as with networked filesystems, but with significant performance and functionality advantages. CXFS runs on top of a storage area network (SAN), where each computer system in the cluster has direct high-speed data channels to a shared set of disks.

CXFS has the following unique features:

- A *peer-to-disk* model for the data access. The shared files are treated as local files by all of the computer systems in the cluster. Each system can read and write the disks at near-local disk speeds; the data passes directly from the disks to the system requesting the I/O, without passing through a data server or over a local

area network (LAN). For the data path, each system is a peer on the SAN; each can have equally fast direct data paths to the shared disks.

Therefore, adding disk channels and storage to the SAN can scale the bandwidth. On large IRIX systems, the bandwidth can scale to Gbytes and even tens of Gbytes per second. Compare this with a networked filesystem with the data typically flowing over a 1- to 100-Mbyte-per-second LAN.

This peer-to-disk data path also removes the file-server data-path bottleneck found in most LAN-based shared filesystems.

- Each host system can buffer the shared disk much as it would for locally attached disks. CXFS maintains the coherency of these distributed buffers, preserving the advanced buffering techniques of the XFS filesystem.
- A flat, single-system view of the filesystem; it is identical from all hosts sharing the file system and is not dependent on any particular host. The path name is a normal POSIX path name; for example, `/u/username/directory`.

The path does not vary if the metadata server moves from one node to another, if the metadata server name is changed, or if a metadata server is added or replaced. This simplifies storage management for administrators and users. Multiple processes on one symmetric multiprocessor (SMP) host and processes distributed across multiple hosts have the same view of the filesystem, with performance similar on each host.

This differs from typical networked filesystems, which tend to include the name of the fileserver in the path name. This difference reflects the simplicity of the SAN architecture with its *direct-to-disk* I/O compared with the extra hierarchy of the LAN filesystem that goes through a named server to get to the disks.

- A full IRIX filesystem interface, including POSIX, System V, and BSD interfaces. This includes filesystem semantics such as mandatory and advisory record locks. No special record-locking library is required.

CXFS has the following restrictions:

- Some filesystem semantics are not appropriate and not supported in shared filesystems. For example, the root filesystem is not an appropriate shared filesystem. Root filesystems belong to a particular host, with system files configured for each particular host's characteristics.
- All processes using a named pipe must be on the same node.

## XFS Features Present in CXFS

CXFS is based on the XFS file system. XFS features that are also present in CXFS include the following:

- Reliability and fast (subsecond) recovery of a log-based filesystem
- 64-bit scalability to 9 million terabytes (9 exabytes) per file
- Speed (high bandwidths, high transaction rates, and fast metadata operations)
- Real-time filesystems
- Dynamically allocated metadata space
- Quotas. (The quota mount options must be the same on all mounts of the filesystem. You can administer quotas anywhere in the cluster just as if this were a regular XFS filesystem.)
- Filesystem reorganizer (defragmenter), which must be run from the CXFS metadata server. See the `fsr_xfs(1M)` man page.

CXFS preserves these underlying XFS features while distributing the I/O directly between the disks and the hosts. The efficient XFS I/O path uses asynchronous buffering techniques to avoid unnecessary physical I/Os by delaying writes as long as possible. This allows the filesystem to allocate the data space efficiently and often contiguously. The data tends to be allocated in large contiguous chunks, which yields sustained high bandwidths.

The XFS directory structure is based on B-trees, which allow XFS to maintain good response times, even as the number of files in a directory grows to tens or hundreds of thousands of files.

The primary difference between XFS and CXFS lies in the configuration tasks.

For more information about XFS features, see *IRIX Admin: Disks and Filesystems*.

The following XFS features are not supported in CXFS:

- Guaranteed-rate I/O
- Swap to a file
- Nesting of mount points; that is, you cannot mount a filesystem on top of a CXFS filesystem

## When to Use CXFS

For the following operations, CXFS exhibits fast file performance, similar to what would be observed with local filesystems. With multiple Fibre Channel connections and multiple Redundant Arrays of Independent Disks (RAID), the bandwidths for these operations can be many hundreds of Mbytes per second or even tens of Gbytes per second:

- Reads from and writes to a file that is opened by only one process
- Reads from and writes to a file where all processes with that file open reside on the same host
- Reads from a file where multiple processes on multiple hosts read the same file
- Reads from and writes to a file using direct-access I/O for multiple processes on multiple hosts

For most filesystem loads, the scenarios above represent the bulk of the file accesses. Thus, CXFS tends to deliver fast local file performance. CXFS is also useful when the amount of data I/O is larger than the amount of metadata I/O. CXFS is faster than NFS because the data does not go through the network.

## When to Use XFS

Some operations can be slower in CXFS than in local XFS filesystems. For example, metadata operations can take longer to complete through the metadata server than on local filesystems. Metadata transaction examples include the following:

- Opening and closing a file
- Changing file size (usually extending a file)
- Creating and deleting files
- Searching a directory

In addition, multiple processes on multiple hosts that are reading and writing the same file using buffered I/O can be slower than when using a local filesystem. This performance difference comes from maintaining coherency among the distributed file buffers; a write into a shared, buffered file will invalidate data (pertaining to that file) that is buffered in other hosts.

Applications that need frequent and high-performance shared, distributed reads and writes can use direct-access I/O in place of buffered I/O for fast local-file performance. For example, distributed parallel databases would tend to use direct-access I/O.

## Hardware and Software Requirements and Support

---

**Note:** You can install either CXFS or IRIS FailSafe 2.0 on a system with applicable licensing. Both products cannot reside on the same system. If you first install IRIS FailSafe, you must deinstall it before you can install CXFS (and vice versa).

---

CXFS requires the following:

- FLEXlm license key
- A private TCP/IP network connected to each CXFS metadata server and client. (A private network is one that is accessible only by administrators and not by users.)
- Origin 2000 or Origin 200 platforms
- One of the following:
  - QLogic HBAs on BROCADE switches with CLARiiON RAID or CIPRICO RAID disks
  - Prisa HBAs on SGI 2800/2400 switches with CLARiiON RAID or CIPRICO RAID disks
  - QLogic HBAs on Emulex hubs with CLARiiON RAID or CIPRICO RAID disks
  - QLogic HBAs with a CLARiiON dual-port RAID disk controller (internal hub)
  - Adaptec host bus adapters (HBAs) with a CLARiiON dual-port RAID disk controller (internal hub)

CXFS works with the following:

- IRISconsole
- Ethernet multiplexer or SCSI multiplexer (used for the reset capability in clusters of three or more nodes)

## CXFS Cluster Manager Tools

You can perform cluster configuration tasks using either the CXFS Cluster Manager graphical user interface (GUI) or command line interface (CLI). Although these tools use the same underlying software to configure and monitor a cluster, the GUI provides the following additional features, which are particularly important in a production system:

- Online help is provided with the **Help** button. You can also click any blue text to get more information about that concept or input field.
- The cluster state is shown visually for instant recognition of status and problems.
- The state is updated dynamically for continuous system monitoring.
- All inputs are checked for correct syntax before attempting to change the cluster database information. In every task, the cluster configuration will not update until you click **OK**.
- Tasks and tasksets take you step-by-step through configuration and management operations, making actual changes to the cluster database as you complete a task.
- The graphical tools can be run securely and remotely on any computer that has a Java virtual machine, including Windows computers and laptops.

The CLI is more limited in its functions. It enables you to configure and administer cluster system using a command line interface only on an IRIX system. It provides a minimum of help or formatted output and does not provide dynamic status except when queried. However, an experienced administrator may find the CLI to be convenient when performing basic configuration tasks or isolated single tasks in a production environment, or when running scripts to automate some cluster administration tasks.

## CXFS Cluster Manager GUI Overview

The CXFS Cluster Manager GUI allows you to administer the entire cluster from a single point. When CXFS daemons have been started in a cluster, you must be sure to connect to a node that is running all the CXFS daemons to obtain the correct cluster status. When CXFS daemons have not yet been started in a cluster, you can connect to any node in the pool.

To ensure that the required privileges are available for performing all of the tasks, you should log in to the GUI as `root`. However, some or all privileges can be

granted to any user by the system administrator using the Privilege Manager, part of the IRIX Interactive Desktop System Administration (*sysadmdesktop*) product. For more information, see the *Personal System Administration Guide*.

The GUI consists of the **CXFS Manager** and the **CXFS Cluster View**.

## CXFS Manager

The **CXFS Manager** provides access to the tasks and tasksets that help you set up and administer your cluster:

- *Tasks* let you set up and monitor individual components of a CXFS cluster.
- *Guided configuration tasksets* consist of a group of tasks collected together to accomplish a larger goal. For example, **Set Up a New Cluster** steps you through the process for creating a new cluster and allows you to launch the necessary individual tasks by simply clicking their titles.

The **CXFS Manager** has the following categories:

- **Overview**, which describes the categories available
- **Guided Configuration**, which contains the tasksets to set up your cluster, define filesystems, and modify an existing cluster
- **Nodes & Cluster**, which defines and manages the nodes and cluster and allows you to start CXFS cluster services (and set the configuration such that they will be automatically restarted upon reboot), or stop CXFS cluster services (and turn off the automatic restart function)
- **Filesystems**, which defines and manage filesystems in the cluster
- **Diagnostics**, which tests the cluster and nodes for configuration problems
- **Find a Task**, which lets you use keywords to search for a specific task

Figure 1-3 shows the **CXFS Manager Overview** screen.

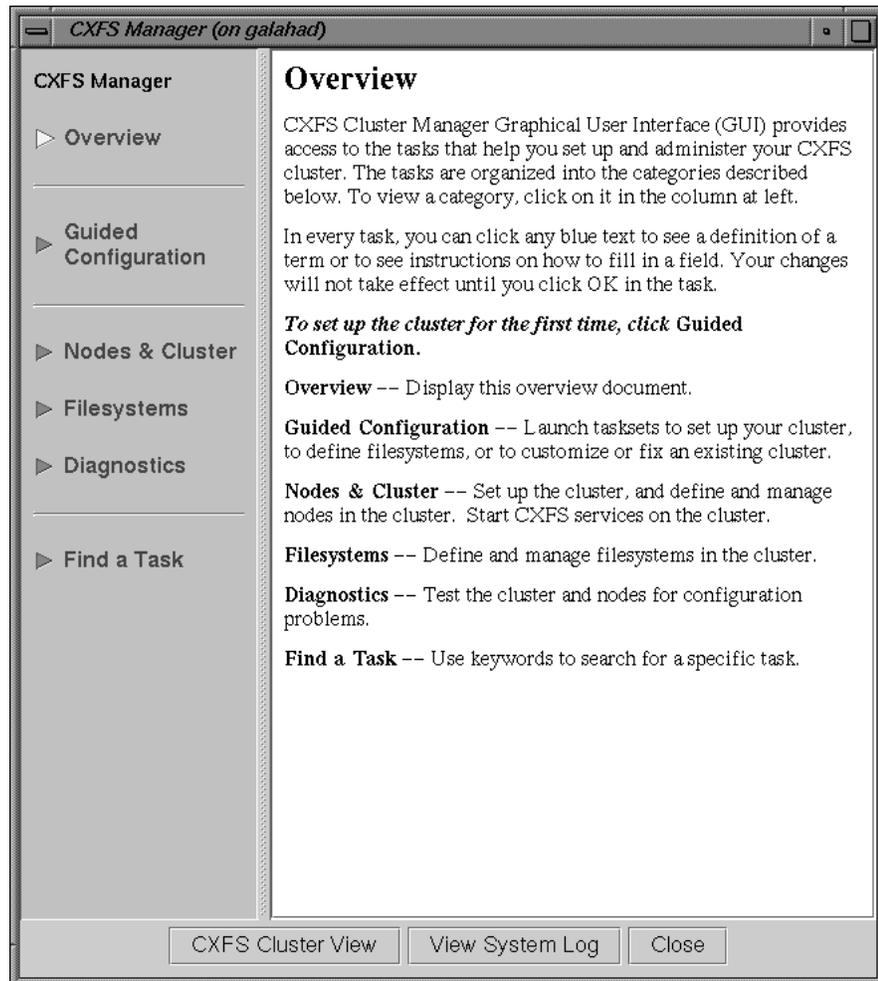


Figure 1-3 CXFS Manager Overview screen

## CXFS Cluster View

The **CXFS Cluster View** provides the following capabilities:

- Shows when a task, such as defining a node, has been completed. It is important that you wait for a task to be completed before continuing; if you move too quickly, errors will occur.
- Shows the relationships among the cluster, nodes, and filesystems.
- Gives access to every item's configuration and status details.
- Shows health of the cluster.
- Gives access to the **CXFS Manager** and to the **SYSLOG**.
- Gives access to help information.

From the **CXFS Cluster View**, you can click on a cluster, or node, or filesystem name to display key information. Figure 1-4 shows an example of a **CXFS Cluster View** screen.



**Figure 1-4** CXFS Cluster View Screen

Figure 1-5 shows an example of a cluster view when the following has been selected from the **View** menu:

**Options**  
**Show Details**

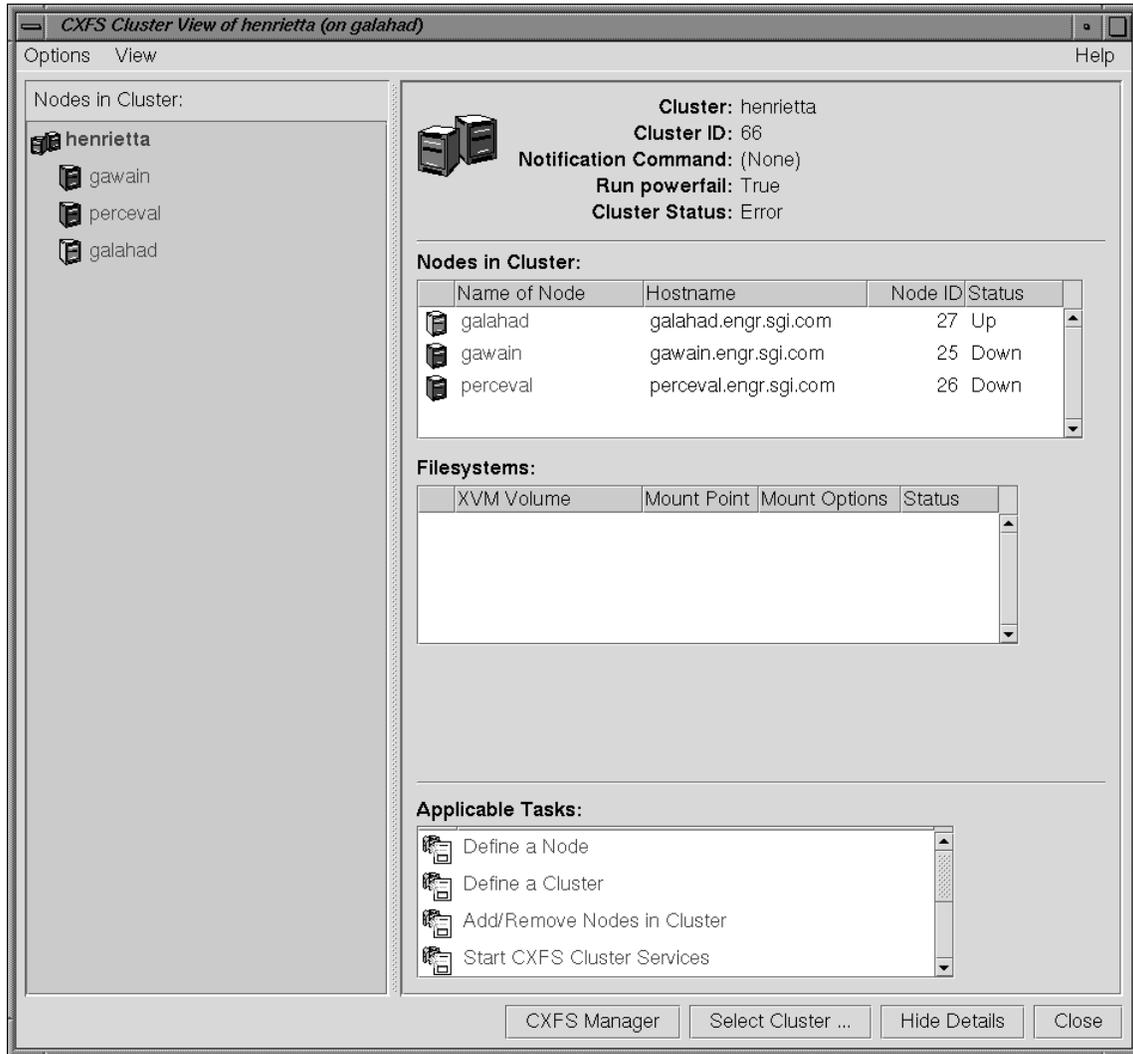


Figure 1-5 Detailed Cluster View

## Starting the GUI

You can start the GUI by launching either the **CXFS Cluster View** or the **CXFS Manager**. However, you should always use the **CXFS Cluster View** to see your progress and to avoid making changes too quickly. After defining a node, you should wait for it to appear in the **CXFS Cluster View** before adding another node. After defining a cluster, you should wait for it to appear before you add nodes to it. If you make changes too quickly, errors can occur.

### Launching the CXFS Cluster View

To launch the **CXFS Cluster View**, use one of the following methods:

- Choose the following from the Toolchest:

**CXFS**  
**CXFS Cluster View**

---

**Note:** You must restart the Toolchest after installing CXFS in order to see the **CXFS** entry on the Toolchest display. Enter the following commands to restart the Toolchest:

```
# killall toolchest  
# /usr/bin/X11/toolchest &
```

In order for this to take effect, `sysadm_cxfs.sw.desktop` must be installed on the client system.

---

A dialog box will appear prompting you to log in to a host.

- Enter the following command line:

```
# /usr/sbin/cxdetail
```

This command line invokes the **CXFS Cluster View**. If no cluster has been defined, the **CXFS Manager** will come up as well.

---

**Note:** If you invoke `cxdetail` before you have defined a cluster, you will get an error message. If you have not yet defined a cluster, you can ignore this error.

---

You also can open the **CXFS Cluster View** window by clicking on the **CXFS Cluster View** button at the bottom of the **CXFS Manager** window. This is the preferred

method if you will have both the **CXFS Manager** and the **CXFS Cluster View** windows open at the same time, because it reuses the existing Java process to open the second window instead of starting a new one, which saves memory usage on the client.

### Launching the CXFS Manager

To launch the **CXFS Manager**, use one of the following methods:

- Choose the following from the Toolchest:

**CXFS**  
**CXFS Manager**

---

**Note:** You must restart the Toolchest after installing CXFS in order to see the **CXFS** entry on the Toolchest display. Enter the following commands to restart the Toolchest:

```
# killall toolchest  
# /usr/bin/X11/toolchest &
```

In order for this to take effect, `sysadm_cxfs.sw.desktop` must be installed on the client system.

---

- Enter the following command line:

```
# /usr/sbin/cxtask
```

- In your Web browser, enter `http://server/CXFSManager` (where *server* is the name of node in the pool or cluster that you want to administer) and press Enter. At the resulting Web page, click on the disk volume icon.

You can use this method of launching the **CXFS Manager** if you want to administer the GUI from a non-IRIX system. If you are running the GUI on an IRIX system, the preferred method is to use the Toolchest menu item or `/usr/sbin/cxtask`.

This method of launching the **CXFS Manager** works only if you have installed the Java Plug-in, exited all Java processes, restarted your browser, and enabled Java. It also requires that the web subsystems are installed on the server. If there is a long delay before the disk volume icon appears, you can click on the **non plug-in** link, but operational glitches may occur as the result of running in the browser-specific Java.

### Viewing Cluster Item Details

To view the details on any cluster item, use the following procedure:

1. Open the **CXFS Cluster View** window
2. Click on the name or icon of any item

The configuration and status details will appear in a separate window. To see the details in the same window, select **Options**. When you then click on the **Show Details** option, the status details will appear in the right side of the window.

### Performing Tasks

To perform an individual task, do the following:

1. Click on the name of a category in the lefthand column of the **CXFS Manager** window.

A list of individual tasks and taskset topics appears in the righthand column.

2. Click on the title of a task in the righthand column.

The task window appears.

---

**Note:** You can click any blue text to get more information about that concept or input field.

---

3. Enter information in the appropriate fields and click **OK** to complete the task. (Some tasks consist of more than one window; in these cases, click **Next** to go to the next window, complete the information there, and then click **OK**.)

---

**Note:** In every task, the cluster configuration will not update until you click **OK**.

---

A dialog box appears confirming the successful completion of the task and displaying additional tasks that you can launch.

4. Continue launching tasks as needed.

## Using Tasksets

The GUI also provides tasksets to guide you through the steps necessary to complete a goal that encompasses several different tasks. Follow these steps to access the CXFS tasksets:

1. Click on **Guided Configuration** in the lefthand column of the **CXFS Manager** window.

A list of tasksets appears in the right hand column.

2. Click on a taskset in the righthand column.

A window appears and lists the series of tasks necessary to accomplish the desired goal.

3. Follow the steps shown, launching tasks by clicking them.

As you click a task, its task window appears. After you complete all of the tasks listed, you can close the taskset window by double-clicking the upper-left corner of its window or clicking **Close** if there is a **Close** button on the window.

## CXFS Cluster Manager CLI Overview

To use the CXFS Cluster Manager command line interface (CLI), you must be logged in as `root`. Then enter either of the following:

```
# /usr/cluster/bin/cluster_mgr -p
```

or

```
# /usr/cluster/bin/cmgr -p
```

---

**Note:** You **must** use the `-p` option.

---

After you have entered this command, you will see the following message and the cluster manager CLI command prompt (`cmgr>`):

```
Welcome to SGI CXFS Cluster Manager Command-Line Interface
```

```
cmgr>
```

## Getting Help

After the command prompt displays, you can enter subcommands. At any time, you can enter `?` or `help` to bring up the CLI help display.

## Completing Actions and Cancelling

When you are creating or modifying a component of a cluster, you can enter either of the following commands:

- `cancel`, which aborts the current mode and discards any changes you have made
- `done`, which commits the current definitions or modifications and returns to the `cmgr` prompt

## Invoking a Shell from within the CLI

To invoke a shell from within the CLI, enter the following

```
cmgr> sh
```

To exit the shell and to return to the `cmgr>` prompt, enter the following:

```
# exit
```



## Installing CXFS Software and Preparing the System

This chapter cover the following steps:

1. "Installing Required Software", page 25
2. "Configuring System Files", page 29
3. "(Optional) Configuring for Automatic Restart", page 34
4. "Configuring Network Interfaces", page 34
5. "Configuring the Serial Ports", page 37
6. "Rebooting the System", page 37

After completing these steps, see Chapter 3, "Initial Configuration of the Cluster", page 39.

### Installing Required Software

---

**Note:** You can install either CXFS or IRIS FailSafe 2.0 on a system with applicable licensing. Both products cannot reside on the same system. If you first install IRIS FailSafe, you must deinstall it before you can install CXFS (and vice versa).

---

Installing the CXFS base CD requires approximately 30.3 Mbytes of space. To install the required software, do the following:

1. On each node in the pool, upgrade to IRIX 6.5.6f according to the *IRIX Installation Instructions*.

To verify that a given node has been upgraded, use the following command to display the currently installed system:

```
# uname -aR
```

2. (For sites with a serial port multiplexer) On each node in the pool, install the version of the multiplexer driver that is appropriate to the operating system. Use the CD that accompanies the multiplexer. Reboot the system after installation.

3. On each node in the pool, do the following:

- a. Install the CXFS license key. When you order a product that requires a license key, the key will be sent to you automatically through electronic mail by the order desk along with instructions for installing it. If you do not have this information, contact SGI or your local support provider. For more information about installing software licenses, see the *IRIX Installation Instructions* booklet.

---

**Note:** Although it is possible to load and configure the product before the license is installed, you cannot start CXFS cluster services without first installing the license.

---

- b. Insert CD-ROM #2 into the CD drive.
- c. Instruct Inst to read the already inserted CD-ROM as follows:

```
Inst> from /CDROM/dist
```

- d. Press <ENTER> to read the CD-ROM:

```
Install software from : [/CDROM/dist] <ENTER>
```

- e. Install the CXFS software:

```
Inst> install *cxfs*
Inst> install sysadm_base*
```

The following subsystems will be installed:

```
cxfs_cluster_admin.man.man
cxfs_cluster_admin.sw.base
cxfs_cluster_control.sw.base
cxfs_cluster_control.sw.cli
cxfs_cluster_services.books.CXFS_AG
cxfs_cluster_services.man.man
cxfs_cluster_services.sw.base
cxfs_cluster_services.sw.cli
cxfs.books.XVM_AG
cxfs.sw.cxfs
cxfs.sw.xvm
sysadm_base.sw.client
sysadm_base.sw.dso
```

```
sysadm_base.sw.server  
sysadm_cxfs.sw.client  
sysadm_cxfs.sw.server  
sysadm_cxfs.sw.web
```

---

**Note:** If you want to run the CXFS Cluster Manager graphical user interface (GUI) as a login other than `root`, you may also want to install `sysadmdesktop`. This action provides commands that allow you to give users privileges, including the privileges required to run the CXFS commands. If you install `sysadmdesktop`, you will install the following subsystems:

```
sysadmdesktop.man.base  
sysadmdesktop.man.relnotes  
sysadmdesktop.sw.base  
sysadmdesktop.sw.data  
sysadmdesktop.sw.sysadm
```

---

4. If you want to use a Web-based version of the GUI, the `sysadm_base.sw.client`, `sysadm_cxfs.sw.client`, and `sysadm_cxfs.sw.web` subsystems must be installed on the pool nodes that you will connect to (by means of a Java-enabled web browser running on any platform) for performing administrative operations. These subsystems are part of the default software installed already in step 3d. You must also have one of the following installed:

- `nss_fasttrack.sw.server` (from the 6.5.5 applications CD-ROM)
- `nss_enterprise.sw.server` (from the Netscape CD-ROM)

If one of these subsystems is not already installed, you must load the appropriate CD-ROM and install the subsystem.

5. If the workstation from which you want to perform CXFS administration (the *administrative workstation*, which can be a node in the cluster or outside of the cluster) runs the GUI client from an IRIX desktop, install the following subsystems:

```
Inst> keep *  
Inst> install cxfs_cluster_services.books.CXFS_AG  
Inst> install java_eoe.sw (version 3.1.1)  
Inst> install java_eoe.sw32  
Inst> install sysadm_base.man
```

```
Inst> install sysadm_base.sw.client
Inst> install sysadm_cxfs.man
Inst> install sysadm_cxfs.sw.client
Inst> install sysadm_cxfs.sw.desktop
Inst> go
```

6. If the administrative workstation launches the GUI client from a Web browser that supports Java, install the `java_plugin` from the CXFS CD.

---

**Note:** This is not the recommended method to launch the GUI on IRIX. Running the GUI client from an IRIX desktop, as in step 5 above, is preferred.

---

If you try to install all subsystems in `java_plugin`, `inst` reports incompatible subsystems (`java_plugin.sw.swing101`, `java_plugin.sw.swing102`, and `java_plugin.sw.swing103`). Do not install these three subsystems because the GUI does not use them.

If the Java plug-in is not installed when the GUI is run from a browser, the browser is redirected to the following URL:

*<http://java.sun.com/products/plugin/1.1/plugin-install.html>*

After installing the Java plug-in, you must close all browser windows and restart the browser.

For a non-IRIX workstation, download the Java Plug-in from the above URL. If the Java plug-in is not installed when the GUI is run from a browser, the browser is redirected to this site.

7. Exit from `inst`:

```
Inst> quit
```

The requickstarting process may take a few minutes to complete.

After you have installed the system and you quit the `inst` interface, you are prompted to reboot the system to apply the changes. However, you will reboot in the step documented by "Rebooting the System", page 37.

## Configuring System Files

When you install the CXFS software, there are some system file considerations you must take into account. This section describes the required and optional changes you make to the following files for each node in the pool:

- `/etc/services`
- `/etc/config/cad.options`
- `/etc/config/fs2d.options`
- (Optional) `/.rhosts`

### Configuring `/etc/services`

Edit the `/etc/services` file so that it contains entries for `sgi-cad` and `sgi-crsd` before installing the `cluster_admin` product on each node in the pool. The port numbers assigned for these processes must be the same in all nodes in the pool.

---

**Note:** `sgi-cad` requires a TCP port.

---

The following shows an example of `/etc/services` entries for `sgi-cad` and `sgi-crsd`:

```
sgi-crsd      7500/udp      # Cluster reset services daemon
sgi-cad      9000/tcp      # Cluster Admin daemon
```

### Configuring `/etc/config/cad.options`

The `/etc/config/cad.options` file contains the list of parameters that the cluster administration daemon (CAD) reads when the process is started. The CAD provides cluster information to the GUI.

The following options can be set in the `cad.options` file:

```
--append_log      Append CAD logging information to the CAD log file
                  instead of overwriting it.

--log_file filename  CAD log file name. Alternately, this can be specified as
                  -lf filename.
```

`-vvvv`                      Verbosity level. The number of `v` characters indicates the level of logging. Setting `-v` logs the fewest messages. Setting `-vvvv` logs the highest number of messages.

The default file has the following options:

```
-lf /var/cluster/ha/log/cad_log --append_log
```

The following example shows an `/etc/config/cad.options` file that uses a medium-level of verbosity:

```
-vv -lf /var/cluster/ha/log/cad_nodename --append_log
```

The default log file is `/var/cluster/ha/log/cad_log`. Error and warning messages are appended to the log file if log file is already present. The contents of this file cannot be modified using the CLI or the GUI.

---

**Note:** If you make a change to the `cad.options` file at any time other than the initial configuration time, you must restart the CAD processes in order for these changes to take effect. You can do this by rebooting the nodes or by entering the following command:

```
# /etc/init.d/cluster restart
```

If you do not have the CXFS license properly installed, you will see the following error:

```
CXFS not properly licensed for this host. Run
    '/usr/cluster/bin/cxfslicense -d'
for detailed failure information.
```

For information about licensing, see "Installing Required Software", page 25.

---

## Configuring `/etc/config/fs2d.options`

The `/etc/config/fs2d.options` file contains the list of parameters that the `fs2d` daemon reads when the process is started. The `fs2d` daemon manages the distribution of the cluster configuration database (CDB) across the nodes in the pool.

The following options can be set in the `fs2d.options` file:

<code>-logevents</code> <i>event name</i>	Log selected events. The following event names may be used: <code>all</code> , <code>internal</code> , <code>args</code> , <code>attach</code> , <code>chandle</code> , <code>node</code> , <code>tree</code> , <code>lock</code> , <code>datacon</code> , <code>trap</code> , <code>notify</code> , <code>access</code> , <code>storage</code> . The default is <code>all</code> .
<code>-logdest</code> <i>log destination</i>	Set log destination. The following log destinations may be used: <code>all</code> , <code>stdout</code> , <code>stderr</code> , <code>syslog</code> , <code>logfile</code> . If multiple destinations are specified, the log messages are written to all of them. If <code>logfile</code> is specified, it has no effect unless the <code>-logfile</code> option is also specified. The default is <code>-logdest stderr</code> , but logging is then disabled if <code>fs2d</code> runs as a daemon, since <code>stdout</code> and <code>stderr</code> are closed when <code>fs2d</code> is running as a daemon. The default is <code>logfile</code> .
<code>-logfile</code> <i>filename</i>	Set log file name. The default is <code>/var/cluster/ha/log/fs2d_log</code>
<code>-logfilemax</code> <i>maximum size</i>	Set log file maximum size (in bytes). If the file exceeds the maximum size, any preexisting <code>filename.old</code> will be deleted, the current file will be renamed to <code>filename.old</code> , and a new file will be created. A single message will not be split across files. If <code>-logfile</code> is set, the default is 10000000.
<code>-loglevel</code> <i>loglevel</i>	Set log level. The following log levels may be used: <code>always</code> , <code>critical</code> , <code>error</code> , <code>warning</code> , <code>info</code> , <code>moreinfo</code> , <code>freq</code> , <code>morefreq</code> , <code>trace</code> , <code>busy</code> . The default is <code>info</code> .
<code>-trace</code> <i>trace_class</i>	Trace selected events. The following trace classes may be used: <code>all</code> , <code>rpcs</code> , <code>updates</code> , <code>transactions</code> , <code>monitor</code> . If you specify this option, you must also specify <code>-tracefile</code> and/or <code>-tracelog</code> . No tracing is done, even if it is requested for one or more classes of events, unless either or both of <code>-tracefile</code> or <code>-tracelog</code> is specified. The default is <code>transactions</code> .
<code>-tracefile</code> <i>filename</i>	Set trace file name. There is no default.
<code>-tracefilemax</code> <i>maximum_size</i>	Set trace file maximum size (in bytes). If the file exceeds the maximum size, any preexisting <code>filename.old</code> will be deleted, the current file will be renamed to <code>filename.old</code> , and a new file will be created.

- [no]tracelog [Do not] trace to log destination. When this option is set, tracing messages are directed to the log destination or destinations. If there is also a trace file, the tracing messages are written there as well. The default is -tracelog.
- [no]parent\_timer [Do not] exit when parent exits. The default is -noparent\_timer.
- [no]daemonize [Do not] run as a daemon. The default is -daemonize.
- l Do not run as a daemon.
- h Print usage message.
- o help Print usage message.

If you use the default values for these options, the system will be configured so that all log messages of level `info` or less, and all trace messages for transaction events to file `/var/cluster/ha/log/fs2d_log`. When the file size reaches 10 Mbytes, this file will be moved to its namesake with the `.old` extension, and logging will roll over to a new file of the same name. A single message will not be split across files.

The following example shows an `/etc/config/fs2d.options` file that directs all `fs2d` logging information to `/var/adm/SYSLOG`, and all `fs2d` tracing information to `/var/cluster/ha/log/fs2d_ops1`. All log events are being logged, and the following trace events are being logged: `rpcs`, `updates`, and `transactions`. When the size of the tracefile `/var/cluster/ha/log/fs2d_ops1` exceeds 100000000, this file is renamed to `/var/cluster/ha/log/fs2d_ops1.old` and a new file `/var/cluster/ha/log/fs2d_ops1` is created. A single message is not split across files. (Line breaks have been added for readability.)

```
-logevents all -loglevel trace -logdest syslog
-trace rpcs -trace updates -trace transactions
-tracefile /var/cluster/ha/log/fs2d_ops1 -tracefilemax 100000000
```

The following example shows an `/etc/config/fs2d.options` file that directs all log and trace messages into one file, `/var/cluster/ha/log/fs2d_chaos6`, for which a maximum size of 100000000 is specified. `-tracelog` directs the tracing to the log file. (Line breaks have been added for readability.)

```
-logevents all -loglevel trace -trace rpcs -trace updates
-trace transactions -tracelog
-logfile /var/cluster/ha/log/fs2d_chaos6
-logfilemax 100000000 -logdest logfile.
```

---

**Note:** If you make a change to the `fs2d.options` file at any time other than the initial configuration time, you must restart the `fs2d` processes in order for those changes to take effect. You can do this by rebooting the nodes or by entering the following command:

```
# /etc/init.d/cluster restart
```

If you do not have the CXFS license properly installed, you will see the following error:

```
CXFS not properly licensed for this host. Run
      '/usr/cluster/bin/cxfslicense -d'
for detailed failure information.
```

For information about licensing, see "Installing Required Software", page 25

---

### **(Optional) Configuring `/.rhosts`**

If you want to use the connectivity diagnostics provided with CXFS, ensure that the `/.rhosts` file allows all the hosts in the cluster to have access to each other to run remote commands such as `rcp(1)` or `rsh(1)`.

For example, suppose you have a cluster with three nodes: `cxfs6`, `cxfs7`, and `cxfs8`. The `/.rhosts` file on each node will be as follows (prompt denotes node name):

```
cxfs6# cat /.rhosts
cxfs7 root
cxfs7-priv root
cxfs8 root
cxfs8-priv root
```

```
cxfs7# cat /.rhosts
cxfs6 root
cxfs6-priv root
cxfs8 root
cxfs8-priv root
```

```
cxfs8# cat /.rhosts
```

```
cxfs6 root
cxfs6-priv root
cxfs7 root
cxfs7-priv root
```

### **(Optional) Configuring for Automatic Restart**

If you want nodes to restart automatically when they are reset or the node is powered on, you must set the boot parameter `AutoLoad` variable to `yes` as follows during the hardware installation of nodes:

```
# nvram AutoLoad yes
```

This setting is recommended, but is not required for CXFS.

You can check the setting of this variable with the following command:

```
# nvram AutoLoad
```

### **Configuring Network Interfaces**

To configure network interfaces, do the following:

1. Ensure that name services are available.

If you are not using DNS or NIS, you should add every IP address and hostname for the nodes to `/etc/hosts` on one node. For example:

```
190.0.2.1 server1-company.com server1
190.0.2.3 stocks
190.0.3.1 priv-server1
190.0.2.2 server2-company.com server2
190.0.2.4 bonds
190.0.3.2 priv-server2
```

You should then add all of these IP addresses to `/etc/hosts` on the other nodes in the cluster.

See the `hosts(4)`, `named(1M)`, `dns(7P)`, and `nis(7P)` man pages and the *IRIX Admin: Networking and Mail* guide and *NIS Administrator's Guide*.

2. On one node, add that node's interfaces and their IP addresses to the `/etc/config/netif.options` file.

For the example:

```
if1name=ec0
if1addr=$HOSTNAME
```

`$HOSTNAME` is an alias for an IP address that appears in `/etc/hosts`.

If there are additional interfaces, their interface names and IP addresses appear on lines like the following:

```
if2name=
if2addr=
```

In the example, the control network name and IP address are as follows:

```
if3name=ec3
if3addr=priv-$HOSTNAME
```

The control network IP address in this example, `priv-$HOSTNAME`, is an alias for an IP address that appears in `/etc/hosts`.

3. If there are more than eight interfaces on the node, change the value of `if_num` to the number of interfaces. For fewer than eight interfaces, the line is as follows:

```
if_num=8
```

4. Repeat steps 1–3 for the other nodes.
5. Edit the file `/etc/config/routed.options` on each node so that the routes are not shown over the control network (routing is turned off). See the `routed(1M)` man page for a list of options.

For example:

```
-q -h -Prdisc_interval=45
```

The options do the following:

- Turn off routing
- Cause host or point-to-point routes to not be advertised (provided there is a network route going the same direction)

- Set the nominal interval with which Router Discovery Advertisements are transmitted to 45 seconds (and their lifetime to 135 seconds)
6. Configure an e-mail alias on each node that sends the CXFS e-mail notifications of cluster transitions to a user outside the CXFS cluster and to a user on the other nodes in the cluster. For example, suppose there are two nodes called `cxfs1` and `cxfs2`. On `cxfs1`, add the following line to the `/usr/lib/aliases` file:

```
cxfs_admin:operations@console.xyz.com,admin_user@cxfs2.xyz.com
```

On `cxfs2`, add the following line to the `/usr/lib/aliases` file:

```
cxfs_admin:operations@console.xyz.com,admin_user@cxfs1.xyz.com
```

The alias you choose (`cxfs_admin` in this case) is the value you will use for the mail destination address when you configure your system. In this example, `operations` is the user outside the cluster and `admin_user` is a user on each node.

---

**Note:** You must run the `newaliases(1M)` command after editing `/usr/lib/aliases`.

---

7. If the nodes use NIS (that is, if `yp` is set to `on` in `chkconfig`) or the BIND domain name server (DNS), then using local name resolution is recommended. You should modify the `/etc/nsswitch.conf` file so that it reads as follows:

```
hosts:                files nis dns
```

---

**Note:** Exclusive use of NIS or DNS for IP address lookup for the cluster nodes has been shown to reduce availability in situations where the NIS or DNS service becomes unreliable.

---

8. If FDDI is being used, finish configuring and verifying the new FDDI station, as explained in the FDDIXPress release notes and the *FDDIXPress Administration Guide*.

## Configuring the Serial Ports

If a node is configured to reset another machine, you must turn off the `getty` process for the tty ports to which the reset serial cables are connected. To do this, perform the following steps on each node:

1. Determine which port is used for the reset serial line.
2. Open the file `/etc/inittab` for editing.
3. Find the line for the port by looking at the comments on the right for the port number from step 1.
4. Change the third field of this line to `off`. For example:

```
t2:23:off:/sbin/getty -N ttyd2 co_9600          # port 2
```

5. Save the file.
6. Enter the following commands to make the change take effect:

```
# killall getty
# init q
```

---

**Note:** If you configure a multinode cluster with the reset daemon running on an IRISconsole system, do not configure the reset port into the IRISconsole, because it may conflict with the reset daemon that the CXFS system is running.

---

## Rebooting the System

Execute the following command on each node to reboot it:

```
# reboot
```

The shutdown process then runs `autoconfig(1M)` to generate the kernel with your changes.



## Initial Configuration of the Cluster

This chapter provides a summary of the steps required to initially configure a cluster using either the graphical user interface (GUI) or the command line interface (CLI). It assumes that you have already performed the steps in Chapter 2, "Installing CXFS Software and Preparing the System", page 25. It points to detailed descriptions in the task reference chapters and in *XVM Volume Manager Administrator's Guide*.

---

**Note:** CXFS requires a license. Although you can load and configure the product before the license is installed, you cannot run CXFS cluster services without first installing the license. See "Installing Required Software", page 25.

---

### Configuring with the GUI

To initially configure the cluster with GUI, do the following:

1. Start the **CXFS Manager** and **CXFS Cluster View** by entering the following:

```
# /usr/sbin/cxdetail
```

You must supply the name of the node you wish to connect to and the `root` password.

---

**Note:** Because you have not yet defined a cluster, you will be notified that one is not available. This is an expected message at this point.

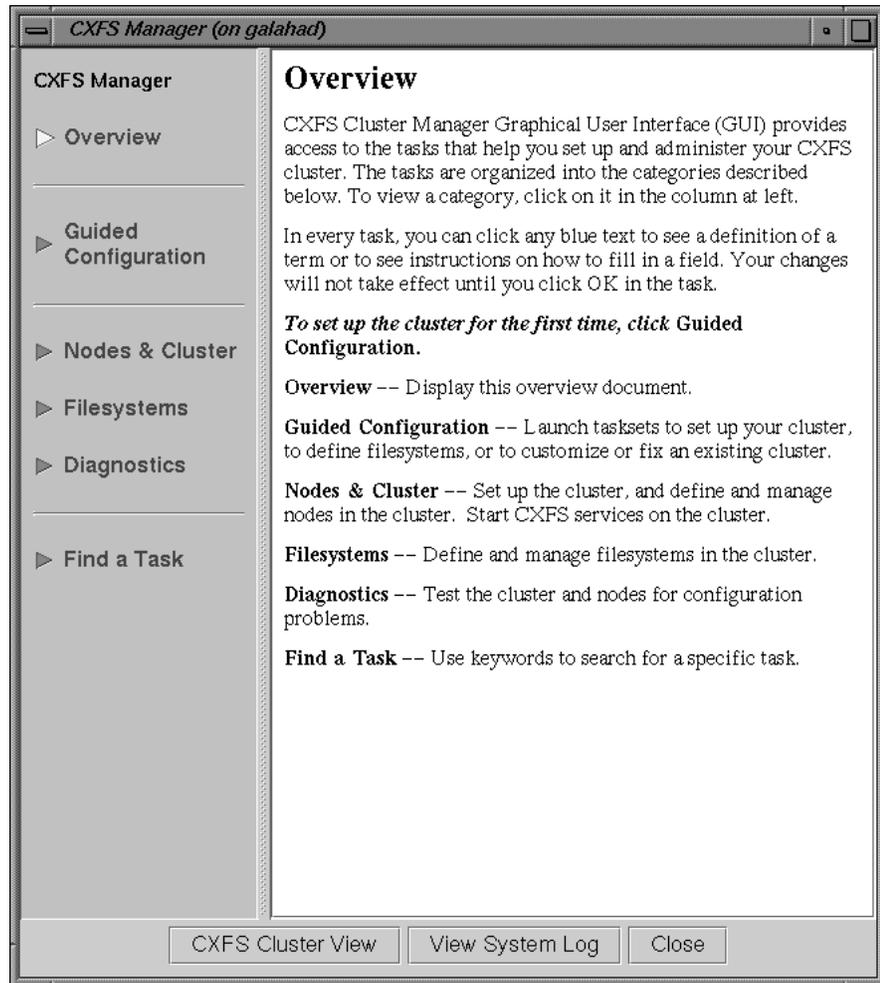
---

You should use the **CXFS Cluster View** to see your progress and to avoid adding nodes too quickly. After defining a node, you should wait for it to appear in the **CXFS Cluster View** before adding another node. After defining a cluster, you should wait for it to appear before you add nodes to it. If you make changes too quickly, errors can occur.

From the **View** menu, you can choose elements within the cluster to examine. To view details of the cluster, click on the cluster name or icon. Status and configuration information will appear in a new window.

There are multiple ways to start these tools. For alternatives, see "Launching the **CXFS Manager**", page 20, and "Launching the **CXFS Cluster View**", page 19.

Figure 3-1 shows the **CXFS Manager** initial screen.



**Figure 3-1** CXFS Manager

2. Select the guided configuration taskset to set up a new cluster:
  - a. On the left side of the **CXFS Manager**, click on the **Guided Configuration** category.

- b. On the right side of the display, click on **Set Up a New Cluster** to launch the taskset.
3. Within the taskset, do the following to define the new cluster:
  - a. Define the node to which you connected when you invoked the GUI by clicking on the first **Define a Node** task (step 1). See "Define a Node with the GUI", page 55.
  - b. *(Optional)* After the first node appears in the **CXFS Cluster View**, define another node by clicking on the second **Define a Node** task (step 2). See "Define a Node with the GUI", page 55.

---

**Note:** Do not add another node until this node appears in the **CXFS Cluster View**. If you add nodes too quickly (before the database can include the node), errors will occur.

---

- c. Repeat step b for each node in the cluster.
- d. Create the cluster definition by clicking on **Define a Cluster**. See "Define a Cluster with the GUI", page 64.
- e. After the cluster icon appears in the **CXFS Cluster View**, add the nodes to the new cluster by clicking on **Add/Remove Nodes in Cluster**. See "Add/Remove Nodes in the Cluster with the GUI", page 61.  
  
Move to the second screen of tasks by clicking on **Next**.
- f. *(Optional)* Verify that the nodes are physically connected by clicking on **Test Connectivity**. See "Test Connectivity with the GUI", page 74.
- g. Click on **Start CXFS Cluster Services**. See "Start CXFS Cluster Services with the GUI", page 66. This action starts CXFS cluster services and sets the configuration so that CXFS cluster services will be restarted automatically whenever the system reboots.

---

**Note:** If you stop CXFS cluster services using either the GUI or the CLI, the automatic restart capability is turned off. You must reboot the nodes and then start CXFS cluster services again to reinstate the automatic restart capability.

---

- h. Click on **Close** at the bottom of the screen to exit the taskset.

4. Use the guided configuration taskset to set up the required filesystems. To launch the taskset, click on **Set Up a New Filesystem** on the right side of the **CXFS Manager**.

Within the taskset, do the following to define the new filesystem:

- a. Click on **Start CXFS Cluster Services** if the services have not been started already. (The current status is displayed beneath the task link.) See "Start CXFS Cluster Services with the GUI", page 66.
- b. Obtain a shell window for one of the nodes in the cluster and use the `fx(1M)` command to create volume headers on the disk drives you will use in the volume. (You do not use the GUI for this step.) For information, see the `fx(1M)` man page and the chapter about performing disk administration procedures in *IRIX Admin: Disks and Filesystems*.
- c. In the shell window, create the XVM logical volumes. For information, see the *XVM Volume Manager Administrator's Guide*.
- d. In the shell window, make the CXFS filesystems using the `mkfs(1M)` command. See the *IRIX Admin: Disks and Filesystems* guide.
- e. Returning to the GUI, click on **Define a Filesystem**. See "Define a Filesystem Metadata Server with the GUI", page 70.
- f. Click on **Mount a Filesystem**. See "Mount a Filesystem with the GUI", page 72.

Repeat these steps for each filesystem.

## Configuring with the CLI

To initially configure the cluster with the CLI, do the following:

1. Define the nodes that are eligible to be part of the cluster. See "Define a Node with the CLI", page 77.

The following example sequence defines three nodes.

To define the first node, named `cxfs6`, enter the following:

---

**Note:** To use the default value for a prompt, press the Enter key. The Enter key is not shown in the examples in this chapter.

---

```
# /usr/cluster/bin/cmgr -p
Welcome to SGI CXFS Cluster Manager Command-Line Interface

cmgr> define node cxfs6
Enter commands, you may enter "done" or "cancel" at any time to exit

Hostname[optional] ? cxfs6
Node ID ? 1001
Do you wish to define system controller info[y/n]:y
Sysctrl Type <chalL|msc|mmsc>? (msc)
Sysctrl Password[optional] ? ( )
Sysctrl Status <enabled|disabled> ? enabled
Sysctrl Owner ? cxfs7
Sysctrl Device ? /dev/ttyd2
Sysctrl Owner Type <tty> ? (tty)
Number of Network Interfaces ? (1)
NIC 1 - IP Address ? 128.162.89.39
NIC 1 - Heartbeat HB (use network for heartbeats) <true|false> ? true
NIC 1 - (use network for control messages) <true|false> ? true
NIC 1 - Priority <1,2,...> ? 1
Node Weight ? (1)

Successfully defined node cxfs6
```

To define the second node, named `cxfs7`, enter the following:

```
cmgr> define node cxfs7
Enter commands, you may enter "done" or "cancel" at any time to exit

Hostname[optional] ? cxfs7
Node ID ? 1002
Do you wish to define system controller info[y/n]:y
Sysctrl Type <chalL|msc|mmsc>? (msc)
Sysctrl Password[optional] ? ( )
Sysctrl Status <enabled|disabled> ? enabled
Sysctrl Owner ? cxfs6
Sysctrl Device ? /dev/ttyd2
```

### 3: Initial Configuration of the Cluster

---

```
Sysctrl Owner Type <tty> ? (tty)
Number of Network Interfaces ? (1)
NIC 1 - IP Address ? 128.162.89.44
NIC 1 - Heartbeat HB (use network for heartbeats) <true|false> ? true
NIC 1 - (use network for control messages) <true|false> ? true
NIC 1 - Priority <1,2,...> ? 1
Node Weight ? (1)
```

Successfully defined node cxf7

To define the third node, named cxf8, enter the following:

```
#cmgr> define node cxf8
Enter commands, you may enter "done" or "cancel" at any time to exit

Hostname[optional] ? cxf8
Node ID ? 1003
Do you wish to define system controller info[y/n]:y
Sysctrl Type <chalL|msc|mmsc>? (msc)
Sysctrl Password[optional] ? ( )
Sysctrl Status <enabled|disabled> ? enabled
Sysctrl Owner ? cxf7
Sysctrl Device ? /dev/ttyd2
Sysctrl Owner Type <tty> ? (tty)
Number of Network Interfaces ? (1)
NIC 1 - IP Address ? 128.162.89.45
NIC 1 - Heartbeat HB (use network for heartbeats) <true|false> ? true
NIC 1 - (use network for control messages) <true|false> ? true
NIC 1 - Priority <1,2,...> ? 1
Node Weight ? (1)
```

Successfully defined node cxf8

You now have three nodes defined in the pool. To verify this, enter the following:

```
cmgr> show nodes in pool

3 Machine(s) defined
    cxf6
    cxf7
    cxf8
```

To show the contents of node `cxfs6`, enter the following:

```
cmgr> show node cxfs6
Logical Machine Name: cxfs6
Hostname: cxfs6
Is FailSafe: false
Is Cellular: true
Nodeid: 1001
Reset type: powerCycle
System Controller: msc
System Controller status: enabled
System Controller owner: cxfs7
System Controller owner device: /dev/ttyd2
System Controller owner type: tty
ControlNet Ipaddr: 128.162.89.39
ControlNet HB: true
ControlNet Control: true
ControlNet Priority: 1
Node Weight: 1
```

2. Define the cluster and add the nodes to it. See "Define a Cluster with the CLI", page 84.

For example, to define a cluster named `mycluster` and add the nodes that are already defined, enter the following:

```
cmgr> define cluster mycluster
Enter commands, you may enter "done" or "cancel" at any time to exit

Cluster Notify Cmd [optional] ?
Cluster Notify Address [optional] ?
Cluster ID ? 55
Number of Cluster FileSystems ? (0)

No nodes in cluster mycluster

Add nodes to or remove nodes from cluster mycluster
Enter "done" when completed or "cancel" to abort

mycluster ? add node cxfs6
mycluster ? add node cxfs7
mycluster ? add node cxfs8
```

```
mycluster ? done  
Successfully defined cluster mycluster
```

For more information, see "Define a Cluster with the CLI", page 84.

To verify the cluster and its contents, enter the following:

```
cmgr> show clusters
```

```
1 Cluster(s) defined  
  mycluster
```

```
cmgr> show cluster mycluster
```

```
Cluster Name: mycluster  
Cluster Is FailSafe: false  
Cluster Is Cellular: true  
Cluster ID: 55
```

```
Cluster mycluster has following 3 machine(s)
```

```
  cxf6  
  cxf7  
  cxf8
```

3. Start CXFS cluster services for the cluster by entering the following:

```
cmgr> start cx_services for cluster clustername
```

For example, for a cluster named *c*, you would enter the following:

```
cmgr> start cx_services for cluster mycluster
```

This action starts CXFS cluster services and sets the configuration so that CXFS cluster services will be restarted automatically whenever the system reboots.

---

**Note:** If you stop CXFS cluster services using either the GUI or the CLI, the automatic restart capability is turned off. You must reboot the nodes and then start CXFS cluster services again to reinstate the automatic restart capability.

---

For more information, see "Display a Cluster with the CLI", page 90.

4. Obtain a shell window for one of the nodes in the cluster and use the `fx(1M)` command to create a volume header on the disk drive. For information, see *XVM Volume Manager Administrator's Guide*.

5. In the shell window, create the XVM logical volumes. For information, see *XVM Volume Manager Administrator's Guide*.
6. In the shell window, use the `mkfs(1M)` command to make the filesystems.
7. Add the filesystems to the cluster by using the `modify cluster` subcommand to `cmgr(1M)`. See "Modify a Cluster and Define and Mount Filesystems with the CLI", page 85.

The following example shows two potential metadata servers for the `test1` filesystem; if `cxfs6` (the preferred server, with rank 0) is not up when the cluster starts or later fails or is removed from the cluster, then `cxfs8` (rank 1) will be used.

---

**Note:** Although the list of metadata servers for a given filesystem is ordered, it is impossible to predict which server will become the server because of network latencies and other unpredictable delays.

---

Do the following, pressing the Enter key to use the default settings:

```
cmgr> modify cluster mycluster
Enter commands, you may enter "done" or "cancel" at any time to exit

Cluster Notify Cmd [optional] ?
Cluster Notify Address [optional] ?
Cluster ID ? (55)
Number of Cluster FileSystems ? (0) 2
FILESYSTEM 0 - Device Name ? /dev/xvm/test1
FILESYSTEM 0 - Mount Point ? /mnts/test1
FILESYSTEM 0 - Mount Options[optional] ?
FILESYSTEM 0 - Status <enabled|disabled> ? enabled
FILESYSTEM 0 - Forced unmount ? <true|false> ? false
Number of Servers ? (0) 2
FILESYSTEM 0 SERVER 0 - Server Node ? cxfs6
FILESYSTEM 0 SERVER 0 - Server Rank ? 0
FILESYSTEM 0 SERVER 1 - Server Node ? cxfs7
FILESYSTEM 0 SERVER 1 - Server Rank ? 1
FILESYSTEM 1 - Device Name ? /dev/xvm/test2
FILESYSTEM 1 - Mount Point ? /mnts/test2
FILESYSTEM 1 - Mount Options[optional] ?
FILESYSTEM 1 - Status <enabled|disabled> ? enabled
FILESYSTEM 1 - Forced unmount ? <true|false> ? false
```

### 3: Initial Configuration of the Cluster

---

```
Number of Servers ? (0) 1
FILESYSTEM 1 SERVER 0 - Server Node ? cxfs8
FILESYSTEM 1 SERVER 0 - Server Rank ? 0
```

Current nodes in cluster mycluster:

```
Node - 1: cxfs6
Node - 2: cxfs7
Node - 3: cxfs8
```

Add nodes to or remove nodes from cluster mycluster  
Enter "done" when completed or "cancel" to abort

```
mycluster ? done
Successfully modified cluster mycluster
```

To see the modified contents of cluster mycluster, enter the following:

```
cmgr> show cluster mycluster
Cluster Name: mycluster
Cluster Is FailSafe: false
Cluster Is Cellular: true
Cluster ID: 55
FileSystem Device Name: /dev/xvm/test1
FileSystem Mount Point: /mnts/test1
FileSystem Mount Options:
FileSystem Status: enabled
FileSystem Force flag: false
    FileSystem Server Node: cxfs7
    FileSystem Server Rank: 1
    FileSystem Server Node: cxfs6
    FileSystem Server Rank: 0
FileSystem Device Name: /dev/xvm/test2
FileSystem Mount Point: /mnts/test2
FileSystem Mount Options:
FileSystem Status: enabled
FileSystem Force flag: false
    FileSystem Server Node: cxfs8
    FileSystem Server Rank: 0

Cluster mycluster has following 3 machine(s)
    cxfs6
    cxfs7
```

`cxfs8`

8. To quit out of `cmgr`, enter the following:

```
cmgr> quit
```



## Reference to GUI Tasks

This chapter tells you how to use the **CXFS Manager** graphical user interface (GUI) to perform the following tasks:

- Initial configuration tasksets:
  - "Set up a New Cluster with the GUI", page 52
  - "Set Up a New Filesystem with the GUI", page 54
- Node tasks:
  - "Define a Node with the GUI", page 55
  - "Add/Remove Nodes in the Cluster with the GUI", page 61
  - "Reset a Node with the GUI ", page 61
  - "Modify a Node with the GUI", page 62
  - "Delete a Node with the GUI", page 63
  - "Display a Node with the GUI", page 63
- Cluster tasks:
  - "Define a Cluster with the GUI", page 64
  - "Modify the Cluster Administrator Notification with the GUI", page 65
  - "Delete a Cluster with the GUI", page 65
  - "Display a Cluster with the GUI", page 66
- CXFS cluster services tasks:
  - "Start CXFS Cluster Services with the GUI", page 66
  - "Stop CXFS Cluster Services with the GUI", page 67
  - "Set Log Configuration with the GUI", page 68
- Filesystem tasks:
  - "Define a Filesystem Metadata Server with the GUI", page 70

- "Mount a Filesystem with the GUI", page 72
- "Unmount a Filesystem with the GUI", page 72
- "Modify a Filesystem with the GUI", page 72
- "Delete a Filesystem with the GUI", page 73
- Diagnostic task:
  - "Test Connectivity with the GUI", page 74

## Initial Guided Configuration Tasksets

There are two initial configuration tasksets:

- "Set up a New Cluster with the GUI", page 52
- "Set Up a New Filesystem with the GUI", page 54

Additional tasksets summarize the steps required to do the following:

- Make changes to a cluster
- Fix or upgrade cluster nodes
- Check status

To access these tasksets, click on **Guided Configuration** in the left side of the **CXFS Manager**.

## Set up a New Cluster with the GUI

The **Set Up a New Cluster** taskset leads you through the steps required to create a new cluster. It encompasses tasks that are detailed elsewhere.

---

**Note:** Within the CXFS tasks, you can click on any blue text to get more information about that concept or input field. In every task, the cluster configuration will not update until you click on **OK**.

---

Do the following to launch the taskset:

1. On the left side of the **CXFS Manager**, click on **Guided Configuration**.

2. On the right side of the display, click on **Set Up a New Cluster**. Within the taskset, do the following to define the new cluster:

---

**Note:** The GUI provides a convenient display of a cluster and its components through the **CXFS Cluster View**. You should use the **CXFS Cluster View** to see your progress and to avoid adding nodes too quickly. After defining a node, you should wait for it to appear in the **CXFS Cluster View** before adding another node. After defining a cluster, you should wait for it to appear before you add nodes to it. If you make changes too quickly, errors can occur. See "Launching the **CXFS Cluster View**", page 19.

---

- a. Click on **Define a Node** to define the node to which you connected. See "Define a Node with the GUI", page 55.
- b. (*Optional*) After the first node icon appears in the Cluster View, click on **Define a Node** to define the other nodes in the cluster. See "Define a Node with the GUI", page 55.

---

**Note:** Do not add another node until this node appears in the **CXFS Cluster View**. If you add nodes too quickly (before the database can include the node), errors will occur.

---

- c. Click on **Define a Cluster** to create the cluster definition. See "Define a Cluster with the GUI", page 64.
- d. After the cluster icon appears in the Cluster View, click on **Add/Remove Nodes in Cluster** to add the nodes to the new cluster. See "Add/Remove Nodes in the Cluster with the GUI", page 61.  
Click on **Next** to move to the second screen of tasks.
- e. (*Optional*) Click on **Test Connectivity** to verify that the nodes are physically connected. See "Test Connectivity with the GUI", page 74.
- f. Click on **Start CXFS Cluster Services**. See "Start CXFS Cluster Services with the GUI", page 66.

---

**Note:** CXFS requires a license. If you install the software without properly installing the license, you cannot start CXFS cluster services. You will see the following message:

```
CXFS not properly licensed for this host. Run
      '/usr/cluster/bin/cxfslicense -d'
      for detailed failure information.
```

For more information about licensing, see "Installing Required Software", page 25.

---

- g. Click on **Close**. Clicking on **Close** exits the taskset; it does not undo the task.

## Set Up a New Filesystem with the GUI

The **Set Up a New Filesystem** taskset leads you through the steps required to create a new filesystem and mount it on all nodes in your cluster. It encompasses tasks that are detailed elsewhere.

---

**Note:** Within the CXFS tasks, you can click on any blue text to get more information about that concept or input field. In every task, the cluster configuration will not update until you click on **OK**.

---

Do the following to launch the taskset:

1. On the left side of the **CXFS Manager**, click on **Guided Configuration**.
2. On the right side of the display, click on **Set Up a New Filesystem**.

Within the taskset, do the following to define the new filesystem:

1. Click on **Start CXFS Cluster Services** if the services have not been started already. (The current status is displayed beneath the task link.) See "Start CXFS Cluster Services with the GUI", page 66.
2. Obtain a shell window for one of the nodes in the cluster and use the `fx(1M)` command to create volume headers on the disk drives you will use in the volume. (You do not use the GUI for this step.) For information, see the `fx(1M)` man page and the chapter about performing disk administration procedures in *IRIX Admin: Disks and Filesystems*.

3. In the shell window, create the XVM logical volumes. For information, see the *XVM Volume Manager Administrator's Guide*.
4. In the shell window, make the CXFS filesystems using the `mkfs(1M)` command. See the *IRIX Admin: Disks and Filesystems* guide.
5. Returning to the GUI, click on **Define a Filesystem**. See "Define a Filesystem Metadata Server with the GUI", page 70.
6. Click on **Mount a Filesystem**. See "Mount a Filesystem with the GUI", page 72.

## Node Tasks

This section tells you how to define, modify, delete, display, and reset a node using the GUI.

---

**Note:** The **Set Up a New Cluster** guided configuration taskset leads you through the tasks required to set up the cluster and nodes. See "Set up a New Cluster with the GUI", page 52.

---

### Define a Node with the GUI

A *cluster node* is a single IRIX image. Usually, a cluster node is an individual computer. The term *node* is used in this guide for brevity; this use of node does not have the same meaning as a node in an Origin system. The *pool* is the entire set of nodes that are available for clustering. The cluster configuration database is distributed to all of the nodes in the pool.

The first node you define must be the node that you have logged into, in order to perform cluster administration.

---

**Note:** Within the CXFS tasks, you can click on any blue text to get more information about that concept or input field. In every task, the cluster configuration will not update until you click on **OK**.

---

To define a node, do the following:

1. On the left side of the **CXFS Manager**, click on **Nodes & Cluster**.

2. On the right side of the display, click on **Define a Node**.

3. Enter the following:

- Hostname of the node you are defining, such as `mynode.company.com` (this can be abbreviated to `mynode` if it is resolved on all nodes). If you do not know the hostname, you can see it in the output of the `ping(1)` command. For example, the following example shows that the full hostname of `lilly` is `lilly.acme.com`:

```
lilly% ping lilly
PING lilly.acme.com (128.162.191.67): 56 data bytes
```

Hostnames cannot begin with an underscore (`_`) or include any whitespace characters.

- Logical name of the node, which can be an abbreviation of the hostname (such as `lilly`) or an entirely different name (such as `nodeA`). Logical names cannot begin with an underscore (`_`) or include any whitespace characters, and can be at most 255 characters.

---

**Note:** If you want to rename a node, you must delete it and then define a new node. You must reuse the node ID number.

---

- Node ID, which must be an integer in the range 1 through 32767 that is unique among the nodes in the pool. You cannot change the node ID number after the node has been defined.

Click on **Next** to move to the set system controller port screen.

- Membership weight (0 or 1), which specifies how much weight to give the node in membership decisions. Use 0 if you do not want the node considered when determining the node membership. The default is 1. Values other than 0 or 1 are not recommended. For more information, see "Node Weights", page 5.
- Networks for incoming cluster messages. Do the following:
  - Enter the IP address or hostname of the top priority network. (The hostname must be resolved in the `/etc/hosts` file or by NIS or DNS.)
  - Select the type of messages that will accepted by the node on this network. You must have networks for heartbeat and control messages; these messages can both be sent on the same network. You can use the **None**

setting if you want to temporarily define a network but do not want it to accept messages.

- Click on **Add** to add the network to the list.

If you later want to modify the network, click on the network in the list to select it, then click on **Modify**.

If you want to delete a network from the list, click on the network in the list to select it, then click on **Delete**.

- You can choose whether or not to use the system controller port to reset the node. If you want CXFS to be able to use the system controller to reset the node, you select the **Set Reset Parameters** checkbox and provide the following information:
  - Model of the node, such as **Origin 2000 Deskside/Origin200** (select from the pull-down list)
  - Reset method, which you select from the pull-down list:
    - **Power Cycle**, which shuts off power to the node and then starts the server up again.
    - **Reset**, which simulates pressing the reset button on the front of the node.
    - **NMI** (non-maskable interrupt), which core-dumps the OS kernel. This option can help you debug a faulty node.
  - Port password, which is the password for the system controller port, **not** the node's `root` password or PROM password. On some machines, the system administrator may not have set this password. If you wish to set or change the system controller port password, consult the hardware manual for your node.
  - If you want to provide reset information now but do not want to allow the reset capability at this time, check the **Temporarily Disable Port**. If this box is checked, CXFS cannot reset the node.
  - Logical name of the owner node, which is the node sending the remote reset command. Serial cables must physically connect the node being defined and the owner node through the system controller port. At run time, the node must be defined in the CXFS pool.

You can select a logical name from the pull-down list or enter the logical name of a node that is not yet defined. However, you must define the node in CXFS before you run the node connectivity diagnostics task.

- TTY device, which is the name of the terminal port (TTY) on the owner host to which the system controller serial cable is connected, such as `/dev/ttyd2`. The other end of the cable connects to this node's system controller port, so the node can be controlled remotely by the other node.

If you do not want to use the reset function at all, click on the **Set System Controller Parameters** box to uncheck (deselect) it.

4. Click on **OK**.

Figure 4-1 and Figure 4-2 show example screens for a node named `lilly`.

**Define a new node.**

This task lets you make a new node eligible to be included in a cluster. As in every task, your changes will not take effect until you click *OK*.

If no nodes are defined, you must first define the node you're logged into. In that case, the hostname will be pre-specified.

**Note:** You can click any blue text to see a definition of that term or to see instructions on how to fill in that field.

Hostname:

Logical Name:

Node ID:

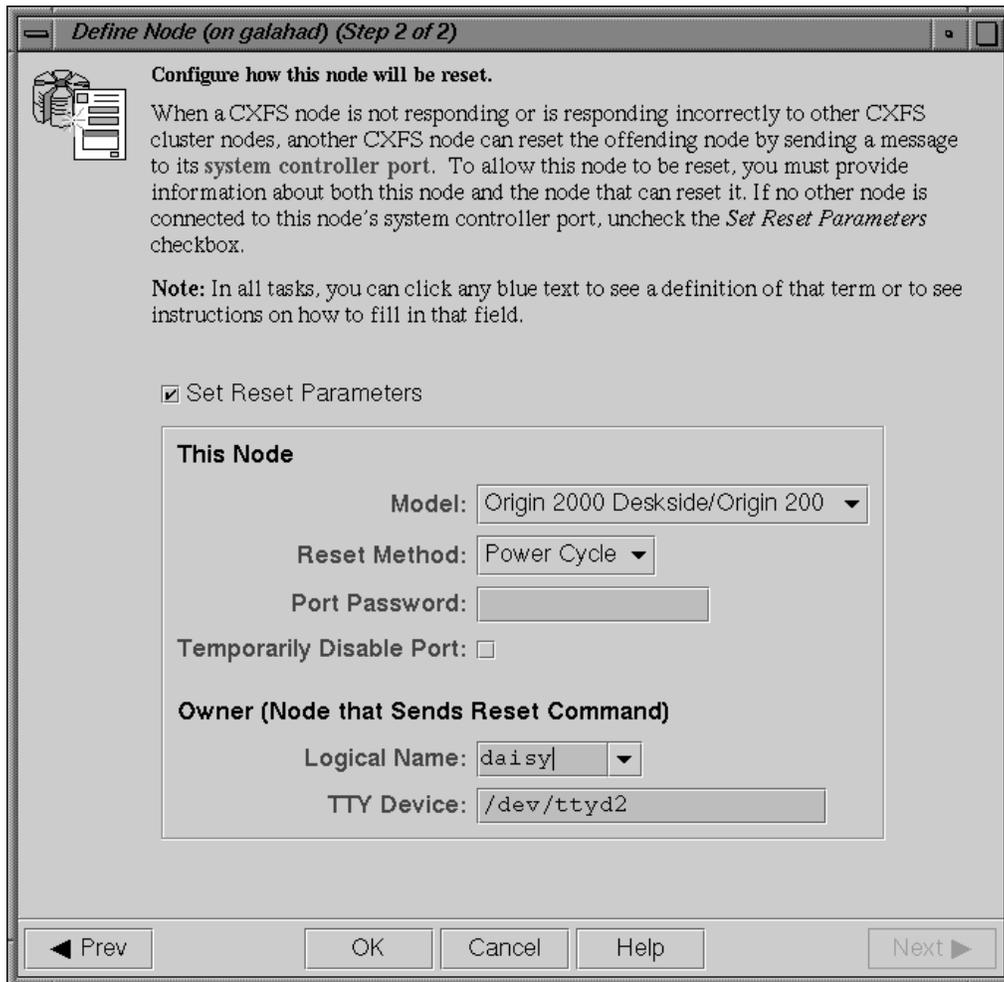
Membership Weight:

**Networks for Incoming Cluster Messages:**

Priority	Networks	Messages to Accept
1.	128.162.191.67	Heartbeat and Control
2.	192.168.168.27	Heartbeat

Navigation buttons:

**Figure 4-1** Example Node Definition



**Figure 4-2** Example System Controller Settings

**Note:** Do not add a second node until the first node icon appears in the **CXFS Cluster View**. For information, see "Launching the **CXFS Cluster View**", page 19.

## Add/Remove Nodes in the Cluster with the GUI

After you have added nodes to the pool and defined the cluster, you can indicate which nodes to include in the cluster.

---

**Note:** Do not add or remove nodes until the cluster icon appears in the **CXFS Cluster View**. For information, see "Launching the **CXFS Cluster View**", page 19.

---

Do the following:

1. On the left side of the **CXFS Manager**, click on **Nodes & Cluster**.
2. On the right side of the display, click on **Add/Remove Nodes in Cluster**.
3. Add or remove the desired nodes:
  - To add a node, select its logical name from the pull-down menu of available nodes and click on **Add**. The node name will appear in the **Nodes to Go into Cluster** list.
  - To delete a node, click on its logical name in the **Nodes to Go into Cluster** screen. (The logical name will be highlighted.) Then click on **Delete**.
4. Click on **OK**.

## Reset a Node with the GUI

You can use the GUI to reset nodes in a cluster. This sends a reset command to the system controller port on the specified node. When the node is reset, other nodes in the cluster will detect the change and remove the node from the active cluster. When the node reboots, it will rejoin the cluster membership.

To reset a node, do the following:

1. On the left side of the **CXFS Manager**, click on **Nodes & Cluster**.
2. On the right side of the display, click on **Reset a Node**.
3. Choose the node to be reset from the pull-down list.
4. Click on **OK**.

## Modify a Node with the GUI

After you have defined a node, you can modify it.

To modify a node, do the following:

1. On the left side of the **CXFS Manager**, click on **Nodes & Cluster**.
2. On the right side of the display, click on **Modify a Node Definition**.
3. Choose the logical name of the node from the pull-down list. After you do this, information for this node will be filled into the various fields.
4. Change the information in the appropriate field as follows:
  - If you want to change the node membership weight, click on in the text entry area and use the backspace key to delete the current value. Then add the correct value. Only use a value of 0 or 1. For more information, see "Node Weights", page 5.
  - If you want to add a network for incoming cluster messages, enter the IP address or hostname into the network text field and click on **Add**.
  - If you want to modify a network that is already in the list, click on the network in the list in order to select it. Then click on **Modify**. This moves the network out of the list and into the text entry area. You can then change it. To add it back into the list, click on **Add**.
  - If you want to delete a network, click on the network in the priority list in order to select it. Then click on **Delete**.
  - If you want to change the priority a network, click on the network in the priority list in order to select it. Then click on the up and down arrows in order to move it to a different position in the list.
  - If you want to change the use of the system controller for resets or change the reset parameters, click on **Next** to move to the next page. Change the values as needed.
5. Click on **OK**.

---

**Note:** If you want to rename a node, you must delete it and then define a new node. You must reuse the node ID number.

---

## Delete a Node with the GUI

You must remove a node from a cluster before you can delete the node from the pool. For information, see "Modify the Cluster Administrator Notification with the GUI", page 65.

To delete a node, perform the following steps:

1. On the left side of the **CXFS Manager**, click on **Nodes & Cluster**.
2. On the right side of the display, click on **Delete a Node**.
3. Select the logical name of the node to be deleted from the pull-down list.
4. Click on **OK**.

## Display a Node with the GUI

After you define nodes, you can display the following:

- Attributes of a node
- Nodes that are members of a specific cluster
- Nodes that have been defined

The GUI provides a convenient graphic display of the defined nodes of a cluster and the attributes of those nodes through the **CXFS Cluster View**. See "Launching the **CXFS Cluster View**", page 19.

To view all nodes defined in the pool, use the following **Cluster View** menu:

**View**  
**Nodes in Pool**

To view all nodes that belong to the cluster, use the following menu:

**View**  
**Nodes in Cluster**

To view filesystems, use the following menu:

**View**  
**Filesystems**

Click on any name or icon to view detailed status and configuration information.

## Cluster Tasks

This section tells how to perform the individual GUI tasks to define, modify, delete, and display a cluster.

---

**Note:** The **Set Up a New Cluster** guided configuration taskset leads you through the tasks required to set up the cluster and nodes. See "Set up a New Cluster with the GUI", page 52.

---

### Define a Cluster with the GUI

A *cluster* is a collection of 2 to 16 cluster nodes coupled to each other by a private (heartbeat) network. A cluster is identified by a simple name. A given node may be a member of only one cluster.

To define a cluster, do the following:

1. On the left side of the **CXFS Manager**, click on **Nodes & Cluster**.
2. On the right side of the display, click on **Define a Cluster**.
3. Enter the following information:
  - Cluster name, which is the logical name of the cluster. The name can have a maximum length of 255 characters.
  - Cluster ID, which is a unique number within your network in the range 1 through 128. The cluster ID is used by the IRIX kernel to make sure that it does not accept cluster information from any other cluster that may be on the network. The kernel does not use the database for communication, so it requires the cluster ID in order to verify cluster communications. This information in the kernel cannot be changed after it has been initialized; therefore, you cannot change a cluster ID after the cluster has been defined.
  - Notification choice:
    - By e-mail, which requires that you specify the e-mail program (/usr/sbin/Mail by default) and the e-mail addresses of those to be identified. To specify multiple addresses, separate them with commas.

CXFS will send e-mail to the addresses whenever the status changes for a node or cluster. If you do not specify an address, notification will not be sent.

- By other command, which requires that you specify the command to be run whenever the status changes for a node or cluster.
  - Never, which specifies that notification is not sent.
4. Click on **OK**.

### **Modify the Cluster Administrator Notification with the GUI**

To change how the cluster administrator is notified of changes in the cluster's state, do the following:

1. On the left side of the **CXFS Manager**, click on **Nodes & Cluster**.
2. On the right side of the display, click on **Modify a Cluster Definition**.
3. Choose the cluster name from the pull-down list.
4. Select the desired notification.
5. Click on **OK**.

To modify the nodes that make up a cluster, see "Add/Remove Nodes in the Cluster with the GUI", page 61.

---

**Note:** If you want to rename a cluster, you must delete it and then define a new cluster. You must use the same cluster ID number.

---

### **Delete a Cluster with the GUI**

You cannot delete a cluster that contains nodes; you must move those nodes out of the cluster first. For information, see "Modify the Cluster Administrator Notification with the GUI", page 65.

To delete a cluster, do the following:

1. On the left side of the **CXFS Manager**, click on **Nodes & Cluster**.
2. On the right side of the display, click on **Delete a Cluster**.

3. Choose the cluster name from the pull-down list.
4. Click on **OK**.

## Display a Cluster with the GUI

The GUI provides a convenient display of a cluster and its components through the **CXFS Cluster View**. See "Launching the **CXFS Cluster View**", page 19.

From the **View** menu, you can choose elements within the cluster to examine. To view details of the cluster, click on the cluster name or icon. Status and configuration information will appear in a new window.

To view this information, select the following menu:

**Options**  
**Show Details**

The status details will appear in the right side of the window. To hide them, use the same menu selection to toggle the **Show Details** setting.

## CXFS Cluster Services Tasks

The following tasks tell you how to start and stop CXFS cluster services and how to set the log configuration.

### Start CXFS Cluster Services with the GUI

---

**Note:** If CXFS cluster services have been previously stopped with the GUI or CLI, you must first reboot the node before starting CXFS cluster services again.

---

To start CXFS cluster services, and set the configuration to automatically restart CXFS cluster services whenever the system is rebooted, do the following:

1. On the left side of the **CXFS Manager**, click on **Nodes & Cluster**.
2. On the right side of the display, click on **Start CXFS Cluster Services**.

3. If you want to start cluster services on only one node, select it from the pull-down list. Otherwise, cluster services will be started on all nodes in the cluster by default.
4. Click on **OK**.

---

**Note:** CXFS requires a license. If you install the software without properly installing the license, you cannot start CXFS cluster services. You will see the following message:

```
CXFS not properly licensed for this host. Run
      '/usr/cluster/bin/cxfslicense -d'
for detailed failure information.
```

For information about licensing, see "Installing Required Software", page 25.

---

To start CXFS cluster services (without changing the setting to restart after a reboot), enter the following:

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

## Stop CXFS Cluster Services with the GUI

When cluster services are stopped on a node, filesystems are automatically unmounted from that node.

To stop cluster services temporarily (that is, allowing them to restart with a reboot if so configured), use the following command line in a shell window outside of the GUI:

```
# /etc/init.d/cluster stop
```

You can stop CXFS on a specified node or cluster, and prevent CXFS cluster services from being restarted by a reboot, by performing the following steps:

---

**Note:** If you stop CXFS cluster services using this method, they will not restart when the node is rebooted.

---

1. On the left side of the **CXFS Manager**, click on **Nodes & Cluster**.
2. On the right side of the display, click on **Stop CXFS Cluster Services**.

3. If you want to stop cluster services on only one node, select it from the pull-down list. Otherwise, cluster services will be stopped on all nodes in the cluster by default.

If you stop cluster services on one node, that node will not longer have access to any filesystems. If that node was acting as a metadata server for a filesystem, no nodes will have access to the filesystem.

4. If you want to forcibly stop CXFS cluster services even if there are errors (which would normally prevent the stop operation), click on the **Force** checkbox.
5. Click on **OK**. It may take a few minutes to complete the process.

After you have stopped CXFS cluster services on a node, the node is no longer an active member of the cluster. You can then shutdown and reboot it using the procedure in "Administrative Shutdown and CXFS Cluster Services", page 105, or use it in a stand-alone fashion. CXFS cluster services will not be restarted when the system reboots.



**Caution:** You should stop CXFS cluster services before using the `shutdown(1M)` or `reboot(1M)` commands. If you execute `shutdown` or `reboot` when cluster services are active, the remaining nodes in the cluster will view it as a node failure and be forced to run recovery against that node.

---

## Set Log Configuration with the GUI

CXFS maintains logs for each of the CXFS daemons. CXFS logs both normal operations and critical errors to `/var/adm/SYSLOG` as well as to individual log files for each log group. You can customize the logs according to the level of logging you wish to maintain.

When you define a log group, you specify the following information:

- The log level
- The file in which to log messages

Table 4-1 shows the default log files.



**Caution:** Do not change the names of the log files. If you change the names, errors can occur.

---

**Table 4-1** Default Log Files in `/var/cluster/ha/log`

Default Log File	Contents
<code>clconfd_nodename</code>	CXFS cluster configuration daemon in <i>nodename</i>
<code>crsd_nodename</code>	Reset daemon in <i>nodename</i>
<code>cli_nodename</code>	Internal administrative commands in node <i>nodename</i> invoked by the GUI and CLI
<code>script_nodename</code>	Scripts in node <i>nodename</i>

### Display Log Group Definitions with the GUI

To display log group definitions, do the following:

1. On the left side of the **CXFS Manager**, click on **Nodes & Cluster**.
2. On the right side of the display, click on **Set Log Configuration**.
3. Choose the log group to display from the menu.

The current log level and log file for that log group will be displayed in the task window, where you can change those settings if you desire.

4. Click on **OK**.

### Configure Log Groups with the GUI

To configure a log group, do the following:

1. On the left side of the **CXFS Manager**, click on **Nodes & Cluster**.
2. On the right side of the display, click on **Set Log Configuration**.
3. Enter the appropriate information:
  - Select the log group from the pull-down list. A *log group* is a set of processes that log to the same log file according to the same logging configuration. Each CXFS daemon creates a log group. Settings apply to all nodes in the pool for the `cli` and `crsd` log groups, and to all nodes in the cluster for the `clconfd` and `diags` log groups.

- Select the log level, which specifies the verbosity of the logging.

The values are as follows:

- **Off** gives no logging
- **Minimal** logs notifications of critical errors and normal operation (these messages are also logged to the **SYSLOG** file)
- **Info** logs **Minimal** notifications plus warnings
- **Default** logs all **Info** messages plus additional notifications
- **Debug 0** through **Debug 9** log increasingly more debug information, including data structures

---

**Note:** The `cmgr(1M)` command uses a set of numbers to indicate these log levels. See "Configure Log Groups with the CLI", page 93.

---

4. Click on **OK**.

## Filesystem Tasks

The following tasks let you configure filesystems as shared XVM volumes. These shared volumes can be directly accessed by all nodes in a CXFS cluster. Each volume is identified by its device name. Each volume must have the same mount point on every node in the cluster.

---

**Note:** The **Set Up a New Filesystem** guided configuration taskset leads you through the steps required to set up a new filesystem. See "Set Up a New Filesystem with the GUI", page 54.

---

### Define a Filesystem Metadata Server with the GUI

---

**Note:** This task assumes that you have created volume headers on your disk drives, created the XVM logical volumes, and made the filesystems. For more information, see "Configuring with the GUI", page 39.

---

To define a filesystem's metadata server, do the following:

1. On the left side of the **CXFS Manager**, click on **Filesystems**.
2. On the right side of the display, click on **Define a Filesystem**.
3. Enter the following information:
  - XVM volume name. This is the device name of an XVM volume that will be shared between all nodes in the CXFS cluster. The name must begin with `/dev/xvm/` and therefore this information is provided for you by default.
  - Mount point. This value is a directory to which the specified XVM volume will be attached. This directory name must begin with a slash (`/`). For more information, see the `mount(1M)` man page.
  - (*Optional*) Mount options. These options are passed to the `mount(1M)` and are used to control access to the specified XVM volume. For a list of the available options, see the `fstab(4)` man page.
  - A list of nodes that will act as metadata servers. To add a node to the list of servers, choose a name from the pull-down node list and click on **Add**.

---

**Note:** The order of servers is significant. The first node listed is the preferred metadata server. Click on a logical name to select it and then click on the arrow buttons to arrange the servers in the order that they should be used.

However, it is impossible to predict which server will actually become the server because of network latencies and other unpredictable delays. The first available node in the list will be used as the metadata server.

---

To delete a node from the list of servers, click on the name in the list to select it and then click on **Delete**.

4. Click on **OK**.

## Mount a Filesystem with the GUI

To mount an existing filesystem on all nodes in the cluster, do the following:

1. On the left side of the **CXFS Manager**, click on **Filesystems**.
2. On the right side of the display, click on **Mount a Filesystem**.
3. Choose a filesystem from the pull-down menu.
4. Click on **OK**.

---

**Note:** If CXFS cluster services are not active, mounting a filesystem will not completely succeed. The filesystem will be marked as ready to be mounted and a warning message will be displayed in the **Mount Filesystem** task. The filesystem will not actually be mounted until you have started CXFS cluster services. For information, see "Start CXFS Cluster Services with the GUI", page 66.

---

## Unmount a Filesystem with the GUI

To unmount a filesystem from all nodes in the cluster, do the following:

1. On the left side of the **CXFS Manager**, click on **Filesystems**.
2. On the right side of the display, click on **Unmount a Filesystem**.
3. Choose a filesystem from the pull-down menu.
4. Click on the **Force** toggle button to force an unmount using the `umount -k` option. The `-k` option attempts to kill processes that have open files or current directories in the appropriate filesystems and then unmount them. For more information, see the `umount(1M)` man page.
5. Click on **OK**.

## Modify a Filesystem with the GUI

To modify an existing filesystem, do the following:

1. On the left side of the **CXFS Manager**, click on **Filesystems**.
2. On the right side of the display, click on **Modify a Filesystem**.

3. Choose a filesystem from the pull-down menu. This displays information for that filesystem in the various fields.
4. Change the information displayed for the selected filesystem as needed:
  - To erase text, backspace over the text or select the text and type over it.
  - To delete a node from the list of servers, click on its name and then click on **Delete**.
  - To add a new node to the list of servers, select it from the pull-down list and click on **Add**.
  - To rearrange the priority of a server, select it by clicking on its name and then click on the arrow buttons as needed.
5. Click on **OK**.

## Delete a Filesystem with the GUI

You cannot delete a filesystem that is currently mounted. To unmount a filesystem, see "Unmount a Filesystem with the GUI", page 72,

To permanently delete an unmounted filesystem, do the following:

1. On the left side of the **CXFS Manager**, click on **Filesystems**.
2. On the right side of the display, click on **Delete a Filesystem**.
3. Choose the name of the filesystem from the pull-down list.
4. Click on **OK**.

## Diagnostic Task

This section tells you how to test connectivity using the GUI. You can test the network and serial connections on the nodes in your cluster by entering the requested inputs. You can test all of the nodes in the cluster at one time, or you can specify an individual node to test.

## Test Connectivity with the GUI

To test connectivity, do the following from the **CXFS Manager**:

---

**Note:** The **Test Node Connectivity** screen requires `rsh(1)` access between hosts. The `.rhosts` file must contain the hosts and local host between which you want to test connectivity.

---

1. On the left side of the **CXFS Manager**, click on **Diagnostics**.
2. Click on **Test Connectivity**.
3. Choose whether to test by network or serial connectivity by clicking on the appropriate radio button.
4. Choose a node to be tested from the pull-down list and add it to the test list by clicking on **Add**.

To delete a node from the list of nodes to be tested, click on the logical name to select it and then click on **Delete**.

5. To start the tests, click on **Start Tests**. To stop the tests, click on **Stop Tests**.
6. To run another test, click on **Clear Output** to clear the status screen and start over with step 3.
7. To exit from the window, click on **Close**.

## Reference to CLI Tasks

The CXFS Cluster Manager command line interface (CLI) consists of the `cmgr(1M)` command. To invoke the CLI, enter one of the following:

```
# /usr/cluster/bin/cmgr -p
# /usr/cluster/bin/cluster_mgr -p
```

This chapter tells you how to use the CLI to perform the following tasks:

- "Set Configuration Defaults with the CLI", page 76
- Node tasks:
  - "Define a Node with the CLI", page 77
  - "Modify a Node with the CLI", page 78
  - "Delete a Node with the CLI", page 79
  - "Display a Node with the CLI", page 82
  - "Reset a Node with the CLI", page 83
- Cluster and Filesystem tasks:
  - "Define a Cluster with the CLI", page 84
  - "Modify a Cluster and Define and Mount Filesystems with the CLI", page 85
  - "Unmount a Filesystem with the CLI", page 88
  - "Remount a Defined Filesystem with the CLI", page 88
  - "Delete a Cluster with the CLI", page 89
  - "Display a Cluster with the CLI", page 90
- CXFS cluster services tasks:
  - "Start CXFS Cluster Services with the CLI", page 91
  - "Stop CXFS Cluster Services with the CLI", page 92
  - "Set Log Configuration with the CLI", page 93

- Diagnostic task:
  - "Test Network Connectivity with the CLI", page 95.

For an overview of the tasks that must be performed to configure a cluster, see Chapter 4, "Reference to GUI Tasks", page 51.

## Set Configuration Defaults with the CLI

You can set a default cluster and node to simplify the configuration process for the current session of the CLI. The default will then be used unless you explicitly specify a name. You can use the following commands to specify default values.

To specify a default cluster, enter the following:

```
cmgr> set cluster clustername
```

To specify a default node, enter the following:

```
cmgr> set node nodename
```

To view the current defaults, enter the following:

```
cmgr> show set defaults
```

For example:

```
cmgr> set cluster mycluster
cmgr> set node cxfs6
cmgr> show set defaults
```

```
Default cluster set to: mycluster
Default node set to: cxfs6
Default cdb set to: /var/cluster/cdb/cdb.db
Default resource_type is not set
Extra prompting is set on
```

## Node Tasks

This section tells you how to define, modify, delete, display, and reset a node using the CLI.

### Define a Node with the CLI

Use the following command to add a logical node definition. Press the Enter key to use default information. (The Enter key is not shown in the examples.) For general information, see "Define a Node with the GUI", page 55.

```
cmgr> define node nodename
Enter commands, you may enter "done" or "cancel" at any time to exit

Hostname[optional] ? hostname
Node ID ? nodeID
Do you wish to define system controller info[y/n]:y|n
Sysctrl Type <chalL|msc|mmsc>? (msc) model_type
Sysctrl Password[optional] ? ( )
Sysctrl Status <enabled|disabled> ? enabled|disabled
Sysctrl Owner ? node_sending_reset_command
Sysctrl Device ? /dev/ttyd2
Sysctrl Owner Type <tty> ? (tty) tty_device
Number of Network Interfaces ? (1) 1|2|3|4|5|6|7|8
NIC 1 - IP Address ? IP_address
NIC 1 - Heartbeat HB (use network for heartbeats) <true|false> ? true|false
NIC 1 - (use network for control messages) <true|false> ? true|false
NIC 1 - Priority <1,2,...> ? priority_number
Node Weight ? (1) 0|1
```

The number of network interfaces must be an integer in the range 1 through 8; 1 is the default. The node weight should be 0 or 1; 1 is the default. For more information, see "Node Weights", page 5, and "Define a Node with the GUI", page 55.

For example:

```
cmgr> define node cxfs6
Enter commands, you may enter "done" or "cancel" at any time to exit

Hostname[optional] ? cxfs6
Node ID ? 1001
Do you wish to define system controller info[y/n]:y
```

```
Sysctrl Type <chalL|msc|mmsc>? (msc)
Sysctrl Password[optional] ? ( )
Sysctrl Status <enabled|disabled> ? enabled
Sysctrl Owner ? cxfs7
Sysctrl Device ? /dev/ttyd2
Sysctrl Owner Type <tty> ? (tty)
Number of Network Interfaces ? (1)
NIC 1 - IP Address ? 128.162.89.39
NIC 1 - Heartbeat HB (use network for heartbeats) <true|false> ? true
NIC 1 - (use network for control messages) <true|false> ? true
NIC 1 - Priority <1,2,...> ? 1
Node Weight ? (1)
```

Successfully defined node cxfs

## Modify a Node with the CLI

To modify an existing node, enter the following:

```
cmgr> modify node nodename
```

After entering this command, you can change any of the information you specified when defining a node except the node ID. See "Define a Node with the CLI", page 77.



**Caution:** Do not change the node ID number after the node has been defined.

---

For example, suppose the definition for a node named *cxfs6* was as follows:

```
cmgr> show node cxfs6
Logical Machine Name: cxfs6
Hostname: cxfs6
Is FailSafe: false
Is Cellular: true
Nodeid: 1001
Reset type: powerCycle
System Controller: msc
System Controller status: enabled
System Controller owner: cxfs7
System Controller owner device: /dev/ttyd2
System Controller owner type: tty
```

```
ControlNet Ipaddr: 128.162.89.39
ControlNet HB: true
ControlNet Control: true
ControlNet Priority: 1
Node Weight: 1
```

To change the system controller to be `cxfs8` and use a node weight of 0, you would enter the following, pressing Enter to keep a current setting (the Enter key is not shown in the examples):

```
cmgr> modify node cxfs6
Enter commands, you may enter "done" or "cancel" at any time to exit

Hostname[optional] ? (cxfs6)
Node ID ? (1001)
Do you wish to modify system controller info[y/n]:y
Sysctrl Type ? (msc)
Sysctrl Password[optional] ? ( )
Sysctrl Status ? (enabled)
Sysctrl Owner ? (cxfs7) cxfs8
Sysctrl Device ? (/dev/ttyd2)
Sysctrl Owner Type ? (tty)
Number of Network Interfaces ? (1)
NIC 1 - IP Address ? (128.162.89.39)
NIC 1 - Heartbeat HB (use network for heartbeats) ? (true)
NIC 1 - (use network for control messages) ? (true)
NIC 1 - Priority <1,2,...> ? (1)
Node Weight ? (1) 0

Successfully modified node cxfs6
```

## Delete a Node with the CLI

To delete a node, enter the following:

```
cmgr> delete node nodename
```

You can delete a node only if the node is not currently part of a cluster. If a cluster currently contains the node, you must first modify that cluster to remove the node from it.

For example, suppose you had a cluster named `mycluster` with the following configuration:

```
cmgr> show cluster mycluster
Cluster Name: mycluster
Cluster Is FailSafe: false
Cluster Is Cellular: true
Cluster ID: 55
FileSystem Device Name: /dev/xvm/test1
FileSystem Mount Point: /mnts/test1
FileSystem Mount Options:
FileSystem Status: enabled
FileSystem Force flag: false
    FileSystem Server Node: cxfs7
    FileSystem Server Rank: 1
    FileSystem Server Node: cxfs6
    FileSystem Server Rank: 0
FileSystem Device Name: /dev/xvm/test2
FileSystem Mount Point: /mnts/test2
FileSystem Mount Options:
FileSystem Status: enabled
FileSystem Force flag: false
    FileSystem Server Node: cxfs8
    FileSystem Server Rank: 0

Cluster mycluster has following 3 machine(s)
    cxfs6
    cxfs7
    cxfs8
```

To delete node `cxfs8`, you would do the following:

```
cmgr> modify cluster mycluster
Enter commands, you may enter "done" or "cancel" at any time to exit

Cluster Notify Cmd [optional] ?
Cluster Notify Address [optional] ?
Cluster ID ? (55)
Number of Cluster FileSystems ? (2)
FILESYSTEM 0 - Device Name ? (/dev/xvm/test1)
FILESYSTEM 0 - Mount Point ? (/mnts/test1)
FILESYSTEM 0 - Mount Options[optional] ?
```

```
FILESYSTEM 0 - Status ? (enabled)
FILESYSTEM 0 - Forced unmount ? ? (false)
Number of Servers ? (2)
FILESYSTEM 0 SERVER 0 - Server Node ? (cxfs7)
FILESYSTEM 0 SERVER 0 - Server Rank ? (1)
FILESYSTEM 0 SERVER 1 - Server Node ? (cxfs6)
FILESYSTEM 0 SERVER 1 - Server Rank ? (0)
FILESYSTEM 1 - Device Name ? (/dev/xvm/test2)
FILESYSTEM 1 - Mount Point ? (/mnts/test2)
FILESYSTEM 1 - Mount Options[optional] ?
FILESYSTEM 1 - Status ? (enabled)
FILESYSTEM 1 - Forced unmount ? ? (false)
Number of Servers ? (1)
FILESYSTEM 1 SERVER 0 - Server Node ? (cxfs8) cxfs6
FILESYSTEM 1 SERVER 0 - Server Rank ? (0)
```

Current nodes in cluster mycluster:

```
Node - 1: cxfs6
Node - 2: cxfs7
Node - 3: cxfs8
```

Add nodes to or remove nodes from cluster mycluster  
Enter "done" when completed or "cancel" to abort

```
mycluster ? remove node cxfs8
mycluster ? done
Successfully modified cluster mycluster
```

```
cmgr> show cluster mycluster
Cluster Name: mycluster
Cluster Is FailSafe: false
Cluster Is Cellular: true
Cluster ID: 55
FileSystem Device Name: /dev/xvm/test1
FileSystem Mount Point: /mnts/test1
FileSystem Mount Options:
FileSystem Status: enabled
FileSystem Force flag: false
    FileSystem Server Node: cxfs6
    FileSystem Server Rank: 0
    FileSystem Server Node: cxfs7
```

```
        FileSystem Server Rank: 1
FileSystem Device Name: /dev/xvm/test2
FileSystem Mount Point: /mnts/test2
FileSystem Mount Options:
FileSystem Status: enabled
FileSystem Force flag: false
        FileSystem Server Node: cxfs6
        FileSystem Server Rank: 0

Cluster mycluster has following 2 machine(s)
        cxfs6
        cxfs7
```

## Display a Node with the CLI

After you have defined a node, you can display the node's parameters with the following command:

```
cmgr> show node nodename
```

For example:

```
cmgr> show node cxfs6
Logical Machine Name: cxfs6
Hostname: cxfs6
Is FailSafe: false
Is Cellular: true
Nodeid: 1001
Reset type: powerCycle
System Controller: msc
System Controller status: enabled
System Controller owner: cxfs8
System Controller owner device: /dev/ttyd2
System Controller owner type: tty
ControlNet Ipaddr: 128.162.89.39
ControlNet HB: true
ControlNet Control: true
ControlNet Priority: 1
Node Weight: 0
```

You can see a list of all of the nodes that have been defined with the following command:

```
cmgr> show nodes in pool
```

For example:

```
cmgr> show nodes in pool
```

```
3 Machine(s) defined
    cxf6
    cxf7
    cxf8
```

You can see a list of all of the nodes that have been defined for a specified cluster with the following command:

```
cmgr> show nodes [in cluster clustername]
```

For example:

```
cmgr> show nodes in cluster mycluster
```

```
Cluster mycluster has following 2 machine(s)
    cxf6
    cxf7
```

If you have specified a default cluster, you do not need to specify a cluster when you use this command. For example:

```
cmgr> set cluster mycluster
```

```
cmgr> show nodes
```

```
Cluster mycluster has following 2 machine(s)
    cxf6
    cxf7
```

## Reset a Node with the CLI

When CXFS is running, you can reset a node with the following command:

```
cmgr> admin reset node nodename
```

This command uses the CXFS daemons to reset the specified node.

You can reset a node in a cluster (even when the CXFS daemons are not running) by using the standalone option of the `admin reset` command:

```
cmgr> admin reset standalone node nodename
```

This command does not go through the CXFS daemons.

## Cluster and Filesystem Tasks

This section tells you how to define, modify, delete, and display a cluster using the CLI. It also tells you how to start and stop CXFS cluster services.

### Define a Cluster with the CLI

When you define a cluster with the CLI, you define a cluster and add nodes to the cluster with the same command. For general information, see "Define a Cluster with the GUI", page 64.

Use the following subcommand to define a cluster:

```
cmgr> define cluster c
Enter commands, you may enter "done" or "cancel" at any time to exit

Cluster Notify Cmd [optional] ?
Cluster Notify Address [optional] ?
Cluster ID ? ID
Number of Cluster FileSystems ? (0) should_be_set_to_0

No nodes in cluster c

Add nodes to or remove nodes from cluster c
Enter "done" when completed or "cancel" to abort

c ? add node node1name
[repeat for each node]
c ? done
Successfully defined cluster c
```

---

**Note:** Normally, you will perform this step when you **modify** the cluster (after performing the steps to create a volume header, create XVM logical volumes, and mount the filesystems). For more information about this sequence of tasks, see "Configuring with the CLI", page 42.

---

For example:

```
cmgr> define cluster mycluster
Enter commands, you may enter "done" or "cancel" at any time to exit

Cluster Notify Cmd [optional] ?
Cluster Notify Address [optional] ?
Cluster ID ? 55
Number of Cluster FileSystems ? (0)

No nodes in cluster mycluster

Add nodes to or remove nodes from cluster mycluster
Enter "done" when completed or "cancel" to abort

mycluster ? add node cxf6
mycluster ? add node cxf7
mycluster ? add node cxf8
mycluster ? done
Successfully defined cluster mycluster
```

The cluster ID is a unique number within your network in the range 1 through 128. The cluster ID is used by the IRIX kernel to make sure that it does not accept cluster information from any other cluster that may be on the network. The kernel does not use the database for communication, so it requires the cluster ID in order to verify cluster communications. This information in the kernel cannot be change after it has been initialized; therefore, you cannot change a cluster ID after the cluster has been defined.

## Modify a Cluster and Define and Mount Filesystems with the CLI

After you have defined a cluster, you can modify its attributes:



**Caution:** Do not enter 0 at the number of filesystems prompt if you have already defined filesystems; if you do, you will lose all of the current information about those filesystems.

---

- To add a new filesystem, enter in a larger filesystem number than is currently defined (that is,  $n+1$ ). Press Enter to keep the defaults for the previous file

systems when displayed; when the new filesystem is displayed, enter in the correct information.

- To permanently remove a filesystem, enter in a smaller number of filesystems at the prompt (that is,  $n-1$ ). If you want to remove the last filesystem in the list, the system is dropped. However, if you want to remove a filesystem that is not the last, you must reenter information.
- To unmount a filesystem, see "Unmount a Filesystem with the CLI", page 88.
- To remount a filestem, see "Remount a Defined Filesystem with the CLI", page 88.

The following example shows the process of modifying a cluster to add two filesystems. There are two potential metadata servers for the `test1` filesystem, `cxfs6` and `cxfs7`. If `cxfs6` (the preferred server, with rank 0) is not up when the cluster starts, then `cxfs7` (rank 1) will be used.

---

**Note:** Although the list of metadata servers for a given filesystem is ordered, it is impossible to predict which server will become the server because of network latencies and other unpredictable delays.

---

Do the following, pressing the Enter key to use the default settings:

```
cmgr> modify cluster mycluster
Enter commands, you may enter "done" or "cancel" at any time to exit

Cluster Notify Cmd [optional] ?
Cluster Notify Address [optional] ?
Cluster ID ? (55)
Number of Cluster FileSystems ? (0) 2
FILESYSTEM 0 - Device Name ? /dev/xvm/test1
FILESYSTEM 0 - Mount Point ? /mnts/test1
FILESYSTEM 0 - Mount Options[optional] ?
FILESYSTEM 0 - Status <enabled|disabled> ? enabled
FILESYSTEM 0 - Forced unmount ? <true|false> ? false
Number of Servers ? (0) 2
FILESYSTEM 0 SERVER 0 - Server Node ? cxfs6
FILESYSTEM 0 SERVER 0 - Server Rank ? 0
FILESYSTEM 0 SERVER 1 - Server Node ? cxfs7
FILESYSTEM 0 SERVER 1 - Server Rank ? 1
FILESYSTEM 1 - Device Name ? /dev/xvm/test2
FILESYSTEM 1 - Mount Point ? /mnts/test2
```

```
FILESYSTEM 1 - Mount Options[optional] ?  
FILESYSTEM 1 - Status <enabled|disabled> ? enabled  
FILESYSTEM 1 - Forced unmount ? <true|false> ? false  
Number of Servers ? (0) 1  
FILESYSTEM 1 SERVER 0 - Server Node ? cxfs8  
FILESYSTEM 1 SERVER 0 - Server Rank ? 0
```

Current nodes in cluster mycluster:

```
Node - 1: cxfs6  
Node - 2: cxfs7  
Node - 3: cxfs8
```

Add nodes to or remove nodes from cluster mycluster  
Enter "done" when completed or "cancel" to abort

```
mycluster ? done  
Successfully modified cluster mycluster
```

To see the modified contents of cluster mycluster, enter the following:

```
cmgr> show cluster mycluster  
Cluster Name: mycluster  
Cluster Is FailSafe: false  
Cluster Is Cellular: true  
Cluster ID: 55  
FileSystem Device Name: /dev/xvm/test1  
FileSystem Mount Point: /mnts/test1  
FileSystem Mount Options:  
FileSystem Status: enabled  
FileSystem Force flag: false  
    FileSystem Server Node: cxfs7  
    FileSystem Server Rank: 1  
    FileSystem Server Node: cxfs6  
    FileSystem Server Rank: 0  
FileSystem Device Name: /dev/xvm/test2  
FileSystem Mount Point: /mnts/test2  
FileSystem Mount Options:  
FileSystem Status: enabled  
FileSystem Force flag: false  
    FileSystem Server Node: cxfs8  
    FileSystem Server Rank: 0
```

```
Cluster mycluster has following 3 machine(s)
  cxfs6
  cxfs7
  cxfs8
```

---

**Note:** If CXFS cluster services are not active, mounting a filesystem will not completely succeed. The filesystem will be marked as ready to be mounted but the filesystem will not actually be mounted until you have started CXFS cluster services. For more information, see "Start CXFS Cluster Services with the CLI", page 91.

---

## Unmount a Filesystem with the CLI

To unmount a filesystem without removing its definition (which will allow you to remount it), change its status to disabled.

The following example unmounts the `test1` filesystem. (Press the Enter key to use the default settings and enter done as soon as you have entered the required argument.)

```
cmgr> modify cluster mycluster
Enter commands, you may enter "done" or "cancel" at any time to exit

Cluster Notify Cmd [optional] ?
Cluster Notify Address [optional] ?
Cluster ID ? (55)
Number of Cluster FileSystems ? (2)
FILESYSTEM 0 - Device Name ? (/dev/xvm/test1)
FILESYSTEM 0 - Mount Point ? (/mnts/test1)
FILESYSTEM 0 - Mount Options[optional] ?
FILESYSTEM 0 - Status <enabled|disabled> ? (enabled) disabled
FILESYSTEM 0 - Forced unmount ? <true|false> ? (false) done

Successfully modified cluster mycluster
```

## Remount a Defined Filesystem with the CLI

To remount a filesystem that has already been defined, change its status to enabled. The following example remounts the `test1` filesystem. (Press the Enter key to use

the default settings and enter done as soon as you have entered the required argument.)

```
cmgr> modify cluster mycluster
Enter commands, you may enter "done" or "cancel" at any time to exit

Cluster Notify Cmd [optional] ?
Cluster Notify Address [optional] ?
Cluster ID ? (55)
Number of Cluster FileSystems ? (2)
FILESYSTEM 0 - Device Name ? (/dev/xvm/test1)
FILESYSTEM 0 - Mount Point ? (/mnts/test1)
FILESYSTEM 0 - Mount Options[optional] ?
FILESYSTEM 0 - Status <enabled|disabled> ? (disabled) enabled
FILESYSTEM 0 - Forced unmount ? <true|false> ? (false) done
Successfully modified cluster mycluster
```

## Delete a Cluster with the CLI

---

**Note:** You cannot delete a cluster that contains nodes; you must first stop CXFS cluster services on the nodes and then redefine the cluster so that it no longer contains the nodes you want to delete.

---

To delete a cluster, enter the following:

```
cmgr> delete cluster clustername
```

For example:

```
cmgr> modify cluster mycluster
Enter commands, you may enter "done" or "cancel" at any time to exit

Cluster Notify Cmd [optional] ?
Cluster Notify Address [optional] ?
Cluster ID ? (55)
Number of Cluster FileSystems ? (2) 0

Current nodes in cluster mycluster:
Node - 1: cxfs6
Node - 2: cxfs7
Node - 3: cxfs8
```

Add nodes to or remove nodes from cluster mycluster  
Enter "done" when completed or "cancel" to abort

```
mycluster ? remove node cxf6
mycluster ? remove node cxf7
mycluster ? remove node cxf8
mycluster ? done
Successfully modified cluster mycluster
```

```
cmgr> delete cluster mycluster
```

```
cmgr> show clusters
```

```
cmgr>
```

## Display a Cluster with the CLI

To display the cluster that has been defined, enter the following:

```
cmgr> show clusters
```

To display the nodes in a cluster, enter the following:

```
cmgr> show cluster clustername
```

For example:

```
cmgr> show clusters
```

```
1 Cluster(s) defined
    mycluster
```

```
cmgr> show cluster mycluster
```

```
Cluster Name: mycluster
Cluster Is FailSafe: false
Cluster Is Cellular: true
Cluster ID: 55
FileSystem Device Name: /dev/xvm/test1
FileSystem Mount Point: /mnts/test1
FileSystem Mount Options:
```

```
FileSystem Status: enabled
FileSystem Force flag: false
    FileSystem Server Node: cxfs6
    FileSystem Server Rank: 0
    FileSystem Server Node: cxfs7
    FileSystem Server Rank: 1
FileSystem Device Name: /dev/xvm/test2
FileSystem Mount Point: /mnts/test2
FileSystem Mount Options:
FileSystem Status: enabled
FileSystem Force flag: false
    FileSystem Server Node: cxfs6
    FileSystem Server Rank: 0
```

```
Cluster mycluster has following 2 machine(s)
    cxfs6
    cxfs7
```

## CXFS Cluster Services Tasks

The following tasks tell you how to start and stop CXFS cluster services and set log levels.

### Start CXFS Cluster Services with the CLI

---

**Note:** If CXFS cluster services have been previously stopped with the GUI or CLI, you must first reboot the node before starting cluster services again.

---

To start CXFS cluster services, and set the configuration to automatically restart CXFS cluster services whenever the system is rebooted, enter the following:

```
cmgr> start cx_services for cluster clustername
```

For example:

```
cmgr> start cx_services for cluster mycluster
```

---

**Note:** CXFS requires a license. If you install the software without properly installing the license, you cannot start CXFS cluster services. You will see the following message:

```
CXFS not properly licensed for this host. Run
      '/usr/cluster/bin/cxfslicense -d'
for detailed failure information.
```

---

To start CXFS cluster services without changing the setting to restart after a reboot, enter the following:

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

## Stop CXFS Cluster Services with the CLI

When cluster services are stopped on a node, filesystems are automatically unmounted from that node.

To stop cluster services temporarily (that is, allowing them to restart with a reboot), use the following command line in a shell window outside of the CLI:

```
# /etc/init.d/cluster stop
```

To stop CXFS cluster services on a specified node or cluster, and prevent CXFS cluster services from being restarted by a reboot, use the following command.

---

**Note:** If you stop CXFS cluster services using this method, they will not restart when the node is rebooted.

---

```
cmgr> stop cx_services [on node nodename] [for cluster clustername] [force]
```

For example:

```
cmgr> stop cx_services on node cxfs6
```

```
cmgr> stop cx_services on cluster mycluster
```

After you have stopped CXFS cluster services in a node, the node is no longer an active member of the cluster. You can then shutdown and reboot it using the procedure in "Administrative Shutdown and CXFS Cluster Services", page 105, or use it in a standalone fashion.



---

**Caution:** You must stop CXFS cluster services before using the `shutdown(1M)` or `reboot(1M)` command. If you execute `shutdown` or `reboot` when cluster services are active, CXFS filesystem metadata servers and XVM servers will not be relocated. The remaining nodes in the cluster will view it as a node failure and be forced to run recovery against that node.

---

## Set Log Configuration with the CLI

For general information about CXFS logs, see "Set Log Configuration with the GUI", page 68.

## Display Log Group Definitions with the CLI

Use the following command to view the log group definitions:

```
cmgr> show log_groups
```

This command shows all of the log groups currently defined, with the log group name, the logging levels, and the log files.

## Configure Log Groups with the CLI

You can configure a log group with the following CLI command:

```
cmgr> define log_group log_group on node nodename [in cluster clustername]
```

The *log\_group* variable can be one of the following:

```
clconfd  
cli  
crsd  
diags
```



---

**Caution:** Do not change the names of the log files. If you change the names, errors can occur.

---

The *log\_level* variable can have one of the following values:

- 0 gives no logging

- 1 logs notifications of critical errors and normal operation (these messages are also logged to the `SYSLOG` file)
- 2 logs Minimal notifications plus warnings
- 5 through 7 log increasingly more detailed notifications
- 10 through 19 log increasingly more debug information, including data structures

For example, to define log group `cli` on node `cxfs6` with a log level of 5:

```
cmgr> sh cli on node cxfs6 in cluster mycluster
```

(Enter "cancel" at any time to abort)

```
Log Level ? (11) 5
```

```
CREATE LOG FILE OPTIONS
```

- 1) Add Log File.
- 2) Remove Log File.
- 3) Show Current Log Files.
- 4) Cancel. (Aborts command)
- 5) Done. (Exits and runs command)

```
Enter option:5
```

```
Successfully defined log group cli
```

### Modify Log Groups with the CLI

Use the following command to modify a log group:

```
cmgr> modify log_group log_group_name on node nodename [in cluster clustername]
```

You modify a log group using the same commands you use to define a log group.

For example, to change the log level of `cli` to be 10, enter the following:

```
cmgr> modify log_group cli on node cxfs6 in cluster mycluster
```

(Enter "cancel" at any time to abort)

```
Log Level ? (2) 10
```

## MODIFY LOG FILE OPTIONS

- 1) Add Log File.
- 2) Remove Log File.
- 3) Show Current Log Files.
- 4) Cancel. (Aborts command)
- 5) Done. (Exits and runs command)

Enter option:5

Successfully modified log group cli

## Diagnostic Tasks

This section tells you how to test connectivity.

### Test Network Connectivity with the CLI

You can use the CLI to test the network connectivity in a cluster. This test checks if the specified nodes can communicate with each other through each configured interface in the nodes. This test will not run if CXFS is running.

Use the following command to test the network connectivity for the nodes in a cluster:

```
cmgr> test connectivity in cluster clustername [on node nodename1 node nodename2 ... ]
```

For example:

```
cmgr> test connectivity in cluster mycluster on node cxfs7
Status: Testing connectivity...
Status: Checking that the control IP_addresses are on the same networks
Status: Pinging address cxfs7 interface ef0 from node cxfs7 [cxfs7]
Notice: overall exit status:success, tests failed:0, total tests executed:1
```

This test yields an error message when it encounters its first error, indicating the node that did not respond. If you receive an error message after executing this test, verify that the network interface has been configured up, using the `ifconfig` command.

For example (line breaks added here for readability):

```
# /usr/etc/ifconfig ef0
ef0: flags=405c43 <UP,BROADCAST,RUNNING,FILTMULTI,MULTICAST,CKSUM,DRVLOCK,IPALIAS>
inet 128.162.89.39 netmask 0xffff0000 broadcast 128.162.255.255
```

The UP in the first line of output indicates that the interface is configured up.

If the network interface is configured up, verify that the network cables are connected properly and run the test again.

## Administration and Maintenance

You can perform offline administration tasks from any node in the cluster. However, when the filesystems are mounted, administration must be done from the server.

The following are the same in CXFS and XFS:

- Disk concepts
- Filesystem concepts
- User interface
- Filesystem creation

For more information about these topics, see *IRIX Admin: Disks and Filesystems*.

The rest of this chapter discusses the following topics:

- "Discovering the Metadata Server for a Filesystem"
- "Volume Management", page 101
- "Disk Management", page 101
- "Filesystem Maintenance", page 102
- "Dump and Restore", page 104

### Discovering the Metadata Server for a Filesystem

To discover the current metadata server for a given filesystem, do the following:

1. Open the CXFS Cluster View. For information about performing this task, see "Launching the **CXFS Cluster View**", page 19.
2. Select the following items from the menu bar:

**View**  
**Filesystems**

**Options**  
**Show Details**

3. Click on the name of the filesystem you wish to view. The name of the metadata server is displayed.

The following figures show this process. In this example, you would click on the `/dev/xvm/cxfs2` filesystem name in Figure 6-1, which pops up a window that shows the details for this filesystem. The metadata server for `/dev/xvm/cxfs2` is `cxfs4`, as shown in Figure 6-2.

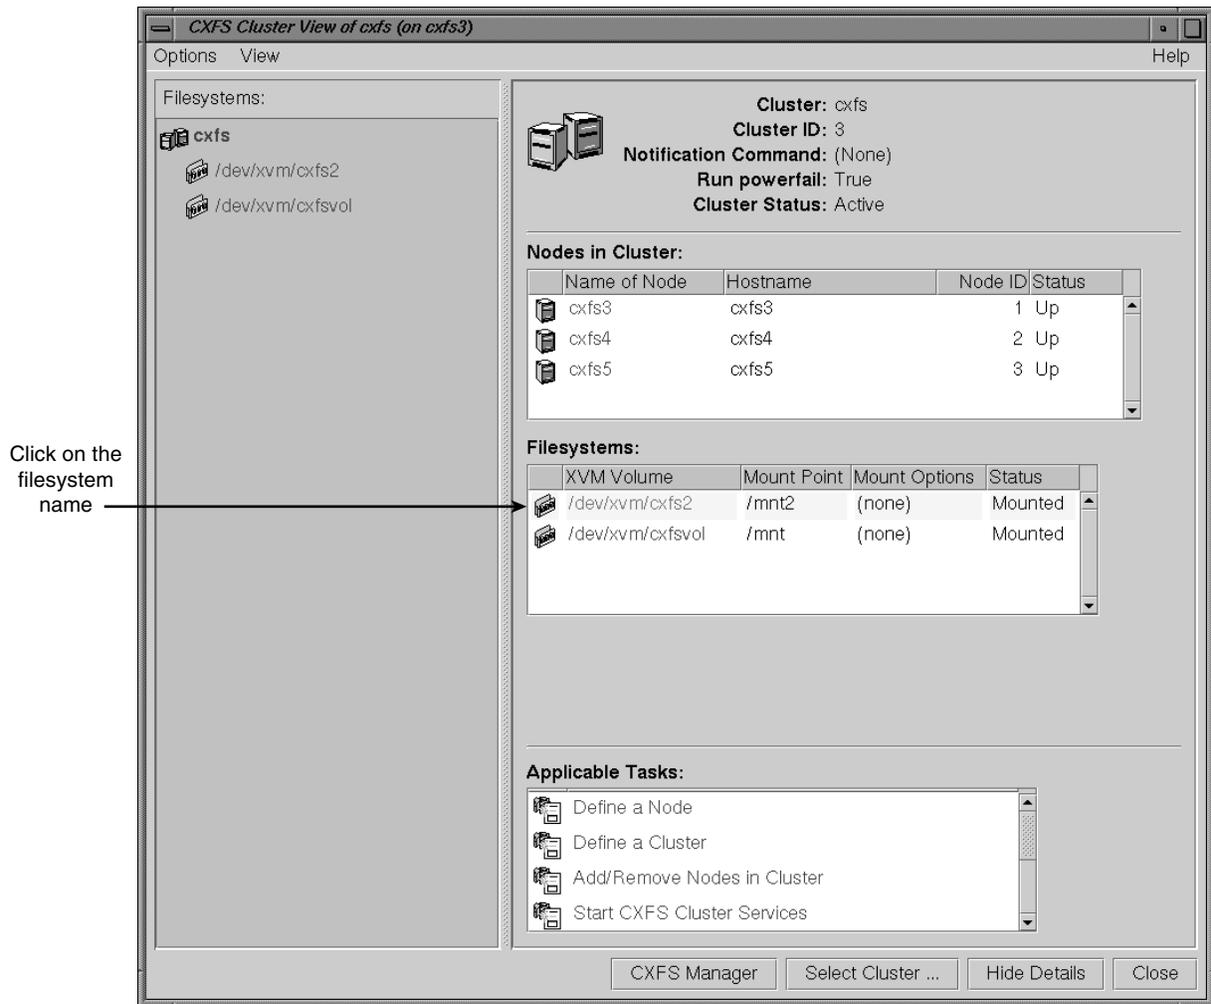


Figure 6-1 Selecting the Filesystem Name

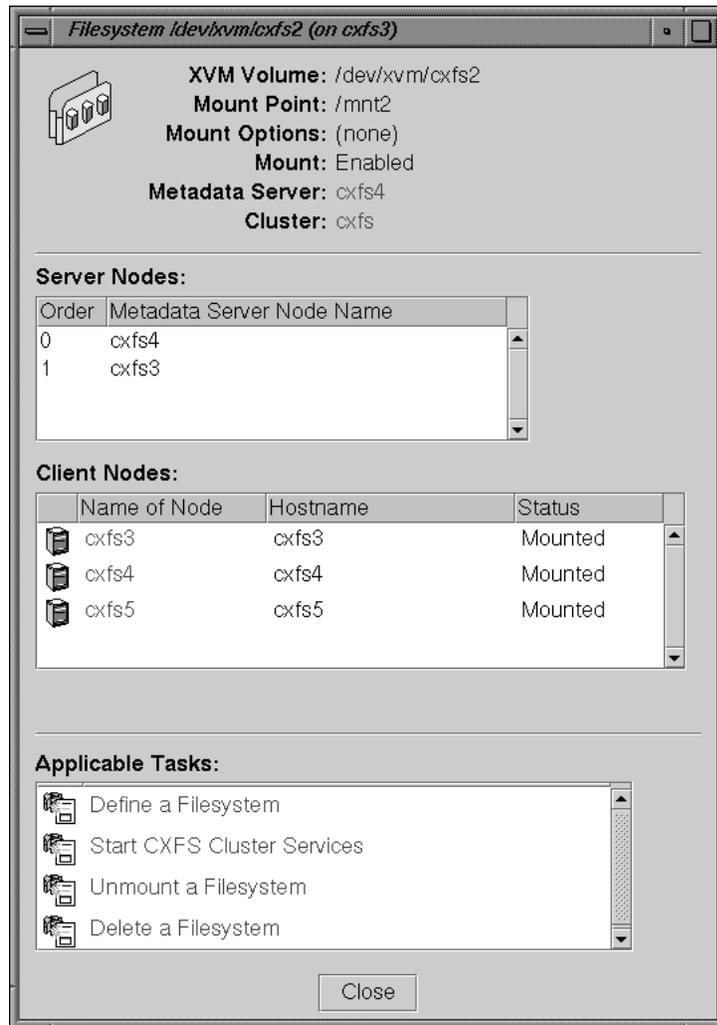


Figure 6-2 Displaying the Metadata Server for a Filesystem

## Volume Management

CXFS uses the XVM volume manager. XVM can combine many disks into high transaction rate, high bandwidth, and highly reliable filesystems. CXFS uses XVM to provide the following:

- Disk striping
- Mirroring
- Concatenation
- Advanced recovery features

---

**Note:** If you try to run an XVM command before starting the CXFS daemons, you will get an error.

---

For more information, see *XVM Volume Manager Administrator's Guide*.

## Disk Management

This section describes the CXFS differences for backups, NFS, Quotas, RealTime, and SAMBA.

### Backups

CXFS enables the use of commercial backup packages such as VERITAS NetBackup and Legato NetWorker for backups that are free from the local area network (LAN), which allows the backup server to consolidate the backup work onto a backup server while the data passes through a storage area network (SAN), rather than through a lower-speed LAN.

For example, a backup package can run on a host on the SAN designated as a backup server. This server can use attached tape drives and channel connections to the SAN disks. It runs the backup application, which views the filesystems through CXFS and transfers the data directly from the disks, through the backup server, to the tape drives.

This allows the backup bandwidth to scale to match the storage size, even for very large filesystems. You can increase the number of disk channels, the size of the backup server, and the number of tape channels to meet the backup-bandwidth requirements.

## **NFS**

You can put an NFS server on top of CXFS so that computer systems that are not part of the cluster can share the file systems. This can be performed on a client or server node.

## **Quotas**

XFS quotas are supported. However, the quota mount options must be the same on all mounts of the filesystem. You can administer quotas from anywhere in the cluster, just as if it was an XFS filesystem.

## **RealTime**

RealTime is supported. A CXFS filesystem can contain a RealTime subvolume. This behaves identically to a RealTime subvolume in XFS.

## **SAMBA**

You can run SAMBA on top of CXFS, allowing NT machines to be clients and have access to the filesystem.

## **Filesystem Maintenance**

Although filesystem information is traditionally stored in `/etc/fstab`, the CXFS filesystems information is relevant to the entire cluster and is therefore stored in the replicated cluster configuration database.

As the administrator, you will supply the CXFS filesystem configuration by using the CXFS Cluster Manager tools. For information about the GUI, see "Filesystem Tasks", page 70; for information about the CLI, see "Cluster and Filesystem Tasks", page 84.

The information is then automatically propagated consistently throughout the entire cluster. The cluster configuration daemon mounts the filesystems on each node according to this information, as soon as it becomes available.

A CXFS filesystem will be automatically mounted on all the nodes in the cluster. You can add a new CXFS filesystem to the configuration when the cluster is active.

Whenever the cluster configuration daemon detects a change in the cluster configuration, it does the equivalent of a `mount -a` command on all the filesystems that are configured. However, this does not allow it to detect that a filesystem has been removed from the cluster definition; therefore, the filesystem will not be automatically unmounted.



---

**Caution:** You must not modify or remove a CXFS filesystem definition while the filesystem is mounted. You must unmount it first and then mount it again after the modifications.

---

## Mounting Filesystems

You supply mounting information with the GUI **Mount a Filesystem** task (which is part of the **Set Up a New Filesystem** guided configuration taskset) or with the CLI `modify` subcommand to `cmgr(1M)`. For information about mounting using the GUI, see "Set Up a New Filesystem with the GUI", page 54, and "Define a Filesystem Metadata Server with the GUI", page 70. For information about defining and mounting a new filestem with the CLI, see "Modify a Cluster and Define and Mount Filesystems with the CLI", page 85; for information about mounting a filesystem that has already been defined but is currently unmounted, see "Remount a Defined Filesystem with the CLI", page 88.

When properly defined and mounted, the CXFS filesystems are automatically mounted on each node by the local cluster configuration daemon, `clconfd`, according to the information collected in the replicated database. After the filesystems configuration has been entered in the database, no user intervention is necessary.



---

**Caution:** Do not attempt to use the `mount(1M)` command to mount a CXFS filesystem. Doing so can result in data loss and/or corruption due to inconsistent use of the filesystem from different machines.

---

CXFS filesystems must be mounted on all nodes in the cluster or none. (Otherwise, the filesystem may be mounted on different nodes in an inconsistent way that may result in data loss and/or corruption.) The GUI and CLI will not let you mount a filesystem on a subset of nodes in the cluster.

Mount points cannot be nested when using CXFS. That is, you cannot have a filesystem within a filesystem, such as `/usr` and `/usr/home`.

## Unmounting Filesystems

To unmount CXFS filesystems, use the GUI **Unmount a Filesystem** task or the CLI `modify` subcommand to `cmgr`. For information, see "Unmount a Filesystem with the GUI", page 72, or "Unmount a Filesystem with the CLI", page 88.

These tasks unmount a filesystem from all nodes in the cluster. Although this action triggers an unmount on all the nodes, some might fail if the filesystem is busy. On metadata servers, the unmount cannot succeed before all the clients have successfully unmounted the filesystem. All nodes will retry the unmount until it succeeds, but there is no centralized report that the filesystem has been unmounted on all nodes. To verify that the filesystem has been unmounted from all nodes, do one of the following:

- Check the `SYSLOG` files on the metadata servers for a message indicating that the filesystem has been unmounted
- Run the GUI or CLI on the metadata server, disable the filesystem from the server and wait until the GUI shows that the filesystem has been fully disabled (it will be an error if it is still mounted on some clients and the GUI will show which clients are left)

## Growing Filesystems

To grow a CXFS filesystem, do the following:

1. Unmount the CXFS filesystem. For information, see "Unmount a Filesystem with the GUI", page 72, or "Unmount a Filesystem with the CLI", page 88.
2. Mount it as an XFS filesystem. See *IRIX Admin: Disks and Filesystems*.
3. Use the `xfsgrowfs(1M)` command to grow it.
4. Unmount the XFS filesystem with the `umount(1M)` command.
5. Mount the CXFS filesystem. See "Mount a Filesystem with the GUI", page 72, or "Remount a Defined Filesystem with the CLI", page 88.

## Dump and Restore

You must perform dump and restore procedures from the server node. The `xfsdump(1M)` and `xfrestore(1M)` command make use of special system calls that

will only function on the CXFS server node. Therefore, you should only use these commands from the server.

The filesystem can have active clients during a dump process.

## Administrative Shutdown and CXFS Cluster Services

To perform an administrative shutdown of a given node, stop CXFS cluster services for the node. For information, see "Stop CXFS Cluster Services with the GUI", page 67, or "Stop CXFS Cluster Services with the CLI", page 92.

After you have stopped CXFS cluster services, the node is no longer an active member of the cluster. You can then safely reboot it or use it in a standalone fashion. If the node will not be used for a length of time, you should remove it from the cluster so that the membership decisions will ignore this machine.



**Caution:** You should stop CXFS cluster services before using the `shutdown(1M)` or `reboot(1M)` command. If you execute `shutdown` or `reboot` when cluster services are active, the remaining nodes in the cluster will view it as a node failure and be forced to run recovery against that node.

---

## Cluster Database Backup and Restore

To perform a backup of the cluster database, use the `cdbBackup(1M)` command. To perform a restore, use the `cdbRestore(1M)`

command. For more information, see the man pages.

## Configuration State Checker Flags

CXFS has the following flags to the `chkconfig(1M)` command:

- `cluster`, which controls the other cluster administration daemons, mainly the replicated cluster configuration database. If it is turned off, the database daemons will not be started at the next reboot and the local copy of the database will not be updated if you make changes to the cluster configuration on the other nodes. This could cause problems later, especially if a majority of nodes are not running the

database daemons. If the database daemons are not running, the cluster configuration database will not be accessible locally and the node will not be configured to join the cluster.

- `cxfs_cluster`, which controls the `clconfd` daemon and whether or not the `cxfs_shutdown` command is used during a system shutdown.

The `cxfs_shutdown` command attempts to withdraw from the cluster gracefully before rebooting. Otherwise, the reboot is seen as a failure and the other nodes have to recover from it.

---

**Note:** These flags only control the processes, not the cluster. Stopping the processes that control the cluster will not stop the cluster, and starting the processes will start the cluster **only** if the CXFS services are marked as activated in the database.

---

## Monitoring Cluster and Node Status

You can view the system status in the following ways:

- Keep continuous watch on the state of a cluster using the **CXFS Cluster View**
- Query the status of an individual node or cluster using either the GUI or the CLI
- View the system log using the GUI or the `tail(1)` command to view the end of the `SYSLOG` file
- Watch log files in `/var/cluster/ha/logs`



**Caution:** Do not change the names of the log files. If you change the names, errors can occur.

---

- Manually test the filesystems with the `ls(1)` command

The following sections describe the procedures for performing some of these tasks.

### Monitoring System Status with the GUI

The easiest way to keep a continuous watch on the state of a cluster is to use the **CXFS Cluster View**. To start it, see "Launching the **CXFS Cluster View**", page 19.

The color legend for component states is as follows:

Grey	Healthy but not online or active
Green	Healthy and active or online
Blinking red	Problems with component

## Node Status

To query the status of a node, you provide the logical name of the node. The node status can be one of the following:

- **UP**, which means that CXFS cluster services are started and the node is part of the cluster membership. For more information about membership, see "Cluster Membership and Quorum", page 4.
- **DOWN**, which means that although CXFS cluster services are started and the node is defined as part of the cluster, the node is not in the current membership.
- **INACTIVE**, which means that the start CXFS cluster services task has not been run.
- **UNKNOWN**, which means that the state cannot be determined because CXFS cluster services are not running on the node performing the query.

State information is exchanged by daemons that run only when CXFS cluster services are started. A given node must be running CXFS cluster services in order to report status on other nodes.

For example, CXFS cluster services must be started on `node1` in order for it to show the status of `node2`. If CXFS cluster services are started on `node1`, then it will accurately report the state of all other nodes in the cluster. However, if `node1`'s CXFS cluster services are not started, it will report the following states:

- **INACTIVE** for its own state, because it can determine that the start CXFS cluster services task has not been run
- **UNKNOWN** as the state of all other nodes, because the daemons required to exchange information with other nodes are not running, and therefore state cannot be determined

## Monitoring Node Status with the GUI

You can use the **CXFS Cluster View** to monitor the status of the nodes. To launch it, see "Launching the **CXFS Cluster View**", page 19.

From the **View** menu, select **Nodes in Cluster** to monitor the health of the nodes in the default cluster.

## Querying Node Status with the CLI

To query node status, use the following CLI command:

```
cmgr> show status of node node_name
```

## Pinging the System Controller with the CLI

When CXFS is running, you can determine whether the system controller on a node is responding by using the following command:

```
cmgr> admin ping node node_name
```

This command uses the CXFS daemons to test whether the system controller is responding.

You can verify reset connectivity on a node in a cluster even when the CXFS daemons are not running by using the `standalone` option of the `admin ping` command:

```
cmgr> admin ping standalone node node_name
```

This command calls the `ping` command directly to test whether the system controller on the indicated node is responding.

## Monitoring Reset Serial Line with the CLI

You can use the CLI to ping the system controller at a node as follows :

```
cmgr> admin ping dev_name device_name of dev_type device_type  
with sysctrl_type system_controller_type
```

## Cluster Status

To query the status of a cluster, you provide the name of the cluster. The cluster status can be one of the following:

- **ACTIVE**, which means the cluster is up and running.
- **INACTIVE**, which means the start CXFS cluster services task has not been run.
- **ERROR**, which means that some nodes are in a **DOWN** state; that is, the cluster **should** be running, but it is not.

- **UNKNOWN**, which means that the state cannot be determined because CXFS cluster services are not running on the node performing the query. For more information, see in "Node Status", page 108.

### Querying Cluster Status with the GUI

You can use the **CXFS Cluster View** to monitor the status of the cluster. To start it, see "Launching the **CXFS Cluster View**", page 19.

### Querying Cluster Status with the CLI

To query node and cluster status, use the following command:

```
cmgr> show status of cluster cluster_name
```

## Troubleshooting

This chapter contains the following topics:

- "Determining Where Problems Lie", page 111
- "Getting a Dump of the Cluster Configuration Database", page 113
- "Avoiding GUI Problems", page 113
- "Restarting CXFS Cluster Services", page 113
- "Corruption of the Root XFS Filesystem", page 114
- "Rebooting", page 114
- "Performing Temporary Maintenance on a Node", page 114
- "Reboot without Rejoining the Cluster", page 115
- "Clearing the Cluster Configuration Database", page 115
- "Avoiding Problems with Log Files", page 116
- "Error Messages", page 116

### Determining Where Problems Lie

You can use the following tools to determine where problems lie:

- Use detailed information panels in the CXFS Cluster Manager GUI to drill down to specific configuration information.
- Run the **Test Connectivity** task in the GUI. See "Test Connectivity with the GUI", page 74.
- Determine how the nodes of the cluster see the current cluster membership by entering the following command on each node:

```
# /usr/cluster/bin/clconf_info
```

This command displays the following fields:

- Node name

- Node ID
- Status (up or down)
- Age (not useful; ignore this field)
- Incarnation (not useful; ignore this field)
- Cell ID, which is a number that is dynamically allocated by the CXFS software when you add a node to a cluster; the user does not define a Cell ID number.

For example:

```
# /usr/cluster/bin/clconf_info
Membership since Fri Sep 10 08:57:36 1999
Node      NodeId   Status   Age    Incarnation   CellId
cxfs6     1001     UP       1      7             0
cxfs7     1002     UP       0      0             1
cxfs8     1003     UP       0      0             2
2 CXFS FileSystems
/dev/xvm/test1 on /mnts/test1 disabled server 0 clients
/dev/xvm/test2 on /mnts/test2 disabled server 0 clients
```

- Check SYSLOG on each node to make sure the CXFS filesystems have been successfully mounted or unmounted. If a mount/unmount fails, the error will be logged and the operation will be retried after a short delay.
- Use the `sar(1)` system activity reporter to show the disks that are active. For example, the following example will show the disks that are active, put the disk name at the end of the line, and poll every second for 10 seconds:

```
# sar DF 1 10
```

For more information, see the `sar(1)` man page.

- Use the `bufview(1)` file system buffer cache activity monitor to view the buffers that are in use. Within `bufview`, you can use the `help` subcommand to learn about available subcommands, such as the `f` subcommand to limit the display to only those with the specified flag. For example, to display the in-use (busy) buffers:

```
# bufview
f
Buffer flags to display bsy
```

For more information, see the `bufview(1)` man page.

- Use the `icrash(1M)` IRIX system crash analysis utility. For more information, see the `icrash(1M)` man page.

## Getting a Dump of the Cluster Configuration Database

When reporting a cluster configuration problem, it is useful to provide a dump of the cluster configuration database. You can extract such a dump with the following command:

```
# /usr/cluster/bin/cdbutil -c 'gettree #' > dumpfile
```

## Avoiding GUI Problems

The GUI provides a convenient display of a cluster and its components through the **CXFS Cluster View**. You should use the **CXFS Cluster View** to see your progress and to avoid adding or removing nodes too quickly. After defining a node, you should wait for it to appear in the **CXFS Cluster View** before adding another node. After defining a cluster, you should wait for it to appear before you add nodes to it. If you make changes too quickly, errors can occur.

You can launch the **CXFS Cluster View** by clicking on **CXFS Cluster View** at the bottom of the **CXFS Manager**.

From the **View** menu, you can choose elements within the cluster to examine. To view details of the cluster, click on the cluster name or icon. Status and configuration information will appear in a new window.

## Restarting CXFS Cluster Services

If CXFS cluster services do not restart after a reboot, it may be that the start flag was turned off by using the stop function of the GUI or CLI. In this case, issuing a `/etc/init.d/cluster start` will not restart the services. You must first reboot the node and then start CXFS cluster services. If you use the GUI or CLI to restart the services, the configuration will be set so that future reboots will also restart CXFS cluster services.

For information, see "Start CXFS Cluster Services with the GUI", page 66, or "Start CXFS Cluster Services with the CLI", page 91.

## Corruption of the Root XFS Filesystem

If the root XFS filesystem is corrupted, do the following:

1. Boot miniroot (for example, with Install System Software in PROM monitor)
2. Select the Admin subcommand menu
3. Obtain a single-user shell, such as by using the `sh(1)` command
4. Execute `xfs_repair(1M)` on the disk with the root XFS filesystem (such as `/dev/dsk/dks0d1s0`)

## Rebooting

Enter the following on every node to reboot:

```
# reboot
```

---

**Note:** If you want CXFS cluster services to restart whenever the node is rebooted, use the GUI or CLI to start CXFS cluster services. For information about the GUI, see "Start CXFS Cluster Services with the GUI", page 66; for information about the CLI, see "Start CXFS Cluster Services with the CLI", page 91.

---

## Performing Temporary Maintenance on a Node

The `cxfs_cluster` flag to `chkconfig(1M)` controls the `clconfd` daemon. If it is turned off, `clconfd` will not be started at the next reboot and the kernel will not be configured to join the cluster. It is useful to turn it off before rebooting the node if you want to temporarily remove the node from the cluster for system or hardware upgrades or for other maintenance work.

Do the following:

```
# /etc/chkconfig cxfs_cluster off
# reboot
```

## Reboot without Rejoining the Cluster

The `cluster` flag to `chkconfig(1M)` controls the other cluster administration daemons, mainly the replicated cluster configuration database. If it is turned off, the database daemons will not be started at the next reboot and the local copy of the database will not be updated if you make changes to the cluster configuration on the other nodes. This could cause problems later, especially if a majority of nodes are not running the database daemons.

If the cluster daemons are causing serious trouble and prevent the machine from booting, you can recover the node by booting in single-user mode, turning off the `cluster` flag, and booting in multiuser mode:

```
# init 1
# /etc/chkconfig cluster off
# init 2
```

## Clearing the Cluster Configuration Database

To clear the cluster configuration database on all the nodes of the cluster, do the following:



---

**Caution:** This procedure deletes all configuration information.

---

1. On **every** node in the cluster, run the following command to stop CXFS cluster services:

```
# /etc/init.d/cluster stop
```

You must complete the step individually on every node in the cluster before proceeding to step 2.

2. Run the following command on every node to clear the database and reconfigure the CXFS flags:

```
# /usr/cluster/bin/cdbreinit /var/cluster/cdb/cdb.db
```

3. Enter the following on every node to reboot:

```
# reboot
```

## Avoiding Problems with Log Files

You should not change the names of the log files. If you change the names of the log files, errors can occur.

However, if the disk is filling with log messages, you can enter the following to reset all log files on the node and ensure that no messages are lost:

```
# /etc/init.d/cluster stop
# mv /var/cluster/ha/log/* /some_other_filesystem_location
# sh -c 'for file in /var/cluster/ha/log/*; do cat /dev/null > $file; done'
# /etc/init.d/cluster start
```

The mv command line above is optional. Using the lines to stop and start cluster services will ensure that no messages are lost; however, they are not required to reset the log files.

## Error Messages

This section describes some of the error messages you may see. In general, the example messages are listed first by type and then in alphabetical order, starting with the message identifier or text.

### License Error

If you see the following message, it means that the CXFS license was not properly installed:

```
CXFS not properly licensed for this host. Run
    /usr/cluster/bin/cxfslicense -d
for detailed failure information.
```

You must install the license before you can start cluster services. See "Installing Required Software", page 25.

### I/O Errors

If a metadata server dies, the clients will get I/O errors such as the following:

```
Sep  2 17:21:39 1A:cxfs0 unix: ALERT: I/O Error Detected.  Shutting down filesystem: /bob2
Sep  2 17:21:39 1A:cxfs0 unix: ALERT: Please umount the filesystem, and rectify the problem(s)
```

```
Sep  2 17:21:39 5A:cxfs0 unix: NOTICE: Signaling end of recovery cell 0
```

You must use the GUI or CLI to unmount the filesystem on each node. (If you use the `umount(1M)` command, CXFS will get confused and try to remount the filesystem.) For information, see "Unmount a Filesystem with the GUI", page 72, or "Modify a Cluster and Define and Mount Filesystems with the CLI", page 85.

You can then use the GUI or CLI to remount the filesystem on another node that will act as the metadata server. For information, see "Mount a Filesystem with the GUI", page 72, or "Modify a Cluster and Define and Mount Filesystems with the CLI", page 85.

If a client dies, the server automatically cleans up the state associated with the client and all other clients continue to access the filesystem.

## IP Address Error

If you have conflicting cluster ID numbers at your site, you will see errors such as the following:

```
WARNING: mtcp ignoring  alive message from 1 with wrong ip addr 128.162.89.34
WARNING: mtcp ignoring  alive message from 0 with wrong ip addr 128.162.89.33
```

A cluster ID number must be unique. To solve this problem, make the cluster ID numbers unique.

This error can occur if you redefine the cluster configuration and start CXFS services while some nodes have stale information from a previous configuration.

To solve the problem, reboot the nodes with stale information. You can determine which nodes have stale information as follows: stale nodes will complain about all of the nodes but the up-to-date nodes will complain only about the stale nodes. The `/var/cluster/ha/log/clconfd_log` file on the stale nodes will also show error messages about `SGI_CMS_CONFIG_ID` failures.

If there are too many error messages to recognize the stale nodes, reboot every node.

## SYSLOG Errors

CXFS logs both normal operations and critical errors to `/var/adm/SYSLOG` as well as to individual log files for each log group.

In general, errors in the `/var/adm/SYSLOG` file take the following form:

*timestamp priority\_&\_facility : nodename process[ID]: <internal\_info> CODE message\_text*

For example:

```
Sep  7 11:12:59 6X:cxf0 cli[5830]: < E clconf 0> CI_IPCERR_NOSEVER, clconf
ipc: ipclnt_connect() failed, file /var/cluster/ha/comm/clconfd-ipc_cxf0
```

Table 8-1 shows the parts of the about SYSLOG message.

**Table 8-1** SYSLOG Error Message Format

Content	Part	Meaning
Sep 7 11:12:59	Timestamp	September 7 at 11:12 AM.
6X	Facility and level	6X indicates an informational message. See <code>syslogd(1M)</code> and the <code>/usr/include/sys/syslog.h</code> file.
cxfs0	Node name	The node whose logical name is <code>cxfs0</code> is the node on which the process is running
cli[5830]	Process[ID]	The process sending the message is <code>cli</code> and its process ID number is 5830
<CI>E clconf 0	Internal information:message source, logging subsystem, and thread ID	The message is from the cluster infrastructure (CI). E indicates that it is an error. The <code>clconf</code> command is the logging subsystem. 0 indicates that it is not multithreaded
CI_IPCERR_NOSEVER, clconf ipc	Internal error code	Information about the type of message; in this case, a message indicating that the server is missing. No error code is printed if it is a normal message.
ipcclnt_connect() failed, file /var/cluster/ha/comm/clconfd- ipc_cxfs0	Message text	A connection failed for the <code>clconfd-ipc_cxfs0</code> file

The following sections present only the message identifiers and text.

### cli Errors

For all `cli` messages, only the last message from the command (which begins with `CLI private command failed`) is meaningful. You can ignore all other `cli` messages.

The following are example errors from the `cli` daemon.

```
CI_ERR_INVALID, CLI private command: failed (Machine (cxfs0)
exists.)
```

You tried to create a new node definition with logical name `cxfs0`; however, that node name already exists in the cluster configuration database. Choose a different name.

```
CI_ERR_INVALID, CLI private command: failed (IP address
(128.162.89.33) specified for control network is cxfs0 is
assigned to control network of machine (cxfs0).)
```

You specified the same IP address for two different control networks of node `cxfs0`. Use a different IP address.

```
CI_FAILURE, CLI private command: failed (Unable to validate
hostname of machine (cxfs0) being modified.)
```

The DNS resolution of the `cxfs0` name failed. To solve this problem, add an entry for `cxfs0` in `/etc/hosts` on all nodes.

```
CI_IPCERR_NOPULSE, CLI private command: failed (Cluster state
is UNKNOWN.)
```

The cluster state is `UNKNOWN` and the command could not complete. This is a transient error. However, if it persists, stop and restart the cluster daemons:

```
# /etc/init.d/cluster stop
# /etc/init.d/cluster start
```

### clconfd Errors

The following errors are sent by the `clconfd` daemon.

```
CI_CONFERR_NOTFOUND, Could not access root node.
```

The cluster configuration database is either non-existent or corrupted, or the database daemons are not responding. Check that the database

does exist; for more information, see "Getting a Dump of the Cluster Configuration Database", page 113.

If you get an error or the dump is empty, re-create the database; for more information, see "Clearing the Cluster Configuration Database", page 115.

If the database exists, restart the cluster administration daemons:

```
# /etc/init.d/cluster stop
# /etc/init.d/cluster start
```

CI\_ERR\_NOTFOUND, Could not get Cellular status for local machine (cxfs1)

The database is corrupted or cannot be accessed. Same actions as above.

CI\_FAILURE, Call to open cdb for logging configuration when it is already open.

This indicates a software problem requiring you to restart the daemons:

```
# /etc/init.d/cluster stop
# /etc/init.d/cluster start
```

CI\_FAILURE, Cell 1 Machine cxfs1: server has no information about a machine that has reset capabilities for this machine

A reset mechanism was not provided for this node. The node will not be automatically reset if it fails. To ensure proper failure handling, use the GUI or the CLI to modify the node's definition and add reset information. See "Define a Node with the GUI", page 55, or "Modify a Node with the CLI", page 78.

CI\_FAILURE, CMD(/sbin/umount -k /dev/xvm/bob1): exited with status 1 (0x1)

An error occurred when trying to unmount the /dev/xvm/bob1 filesystem. Messages from the umount(1M) command are usually issued just before this message and provide more information about the reason for the failure.

```
CI_FAILURE, CMD(/sbin/clmount -o 'server_list=(cxfs0,cxfs1)'  
/dev/xvm/bob2 /bob2): exited with status 1 (0x1)
```

An error occurred when trying to mount the /dev/xvm/bob2 filesystem. Messages from themount(1M) command are usually issued just before this message and provide more information about the reason of the failure.

```
CI_FAILURE, Could not write newincarnation number to CDB, error  
= 9.
```

There was a problem accessing the cluster configuration database. Retry the operation. If the error persists, stop and restart the cluster daemons:

```
# /etc/init.d/cluster stop  
# /etc/init.d/cluster start
```

If the problem persists, clear the database, reboot, and re-create the database. See "Clearing the Cluster Configuration Database", page 115.

```
CI_FAILURE, Exiting, monitoring agent should revive me.
```

The daemon decided requires fresh data. It will be automatically restarted.

```
CI_FAILURE, No node for client (3) of filesystem (/dev/xvm/bob1)  
on (/bob1).
```

(There may be many repetitions of this message.) The filesystem appears to still be mounted on a node that is no longer in the cluster configuration database. If you can identify the node that used to be in the cluster and still has the filesystem mounted, reboot that node. Otherwise, reboot the entire cluster.

```
CI_FAILURE, No node for server (-1) of filesystem  
(/dev/xvm/bob1) on (/bob1).
```

(There may be many repetitions of this message.) The filesystem appears to still be mounted on a node that is no longer in the cluster configuration database. If you can identify the node that used to be in the cluster and still has the filesystem mounted, reboot that node. Otherwise, reboot the entire cluster.

```
CI_FAILURE, Node cxfs0: SGI_CMS_HOST_ID(tcp,128.162.8 >9.33)
error 149 (Operation already in progress)
```

The kernel already had this information; you can ignore this message.

```
CI_FAILURE, Unregistered from crs.
```

The clconfd daemon is no longer connected to the reset daemon and will not be able to handle resets of failed nodes. There is no corrective action.

```
CI_IPCERR_NOSERVER, Crs_register failed,will retry later.
Resetting not possible yet.
```

The clconfd daemon cannot connect to the reset daemon. It will not be able to handle resets of failed nodes. Check the reset daemon's log file (/var/cluster/ha/log/crsd\_) for more error messages.

```
Clconfd is out of membership, will restart after notifying
clients.
```

The clconfd daemon does not have enough information about the current state of the cluster. It will exit and be automatically restarted with fresh data.

```
CMD(/sbin/clmount -o 'server_list=(cxfs0,cxfs1)' /dev/xvm/bob2
/bob2): /dev/xvm/bob2: Invalid argumentSep 9 14:12:43 6X:cxfs0
clconfd[345]: < E clconf 3> CI_FAILURE, CMD(/sbin/clmount -o
'server_list=(cxfs0,cxfs1)' /dev/xvm/bob2 /bob2): exited with
status 1 (0x1)
```

The first message comes from the clmount command (the internal CXFS mount command) and explains the error (an invalid argument was issued). The second message says that the mount failed.

## crsd Errors

The following errors are sent by the crsd daemon.

```
CI_ERR_NOTFOUND, No logging entries found for group crsd, no
logging will take place - Database entry #global#logging#crsd
not found.
```

No crsd logging definition was found in the cluster configuration database. This can happen if you start cluster processes without

creating the database. To create the initial cluster configuration database, enter the following:

```
# /usr/cluster/bin/cbdreinit
```

CI\_ERR\_RETRY, Could not find machine listing.

The crsd daemon could not find the local node in the cluster configuration database. You can ignore this message if the local node definition has not yet been created.

CI\_ERR\_SYS:125, bind() failed.

The sgi-crsd port number in the /etc/services file is not unique, or there is no sgri-crsd entry in the file. For information about adding this entry, see "Configuring /etc/services", page 29.

CI\_FAILURE, Entry for sgi-crsd is missing in /etc/services.

The sgi-crsd entry is missing from the /etc/services file. For information about adding this entry, see "Configuring /etc/services", page 29.

CI\_FAILURE, Initialization failed, exiting.

A sequence of messages will be ended with this message; see the messages prior to this one in order to determine the cause of the failure.

### **cmond Errors**

The following errors are sent by the cmond daemon.

Could not register for notification.cdb\_error = 7

The cluster configuration database was not initialized when the cluster process was started. To create the initial cluster configuration database, enter the following:

```
# /usr/cluster/bin/cbdreinit
```

Process clconfd:343 of group cluster\_cx exited, status = 3.

The clconfd process exited with status 3, meaning that the process will not be restarted by cmond. No corrective action is needed.

```
Process crsd:1790 of group cluster_control exited, status = 127
```

The crsd process exited with an error (non-zero) status. Look at the corresponding daemon logs for error messages.

## fs2d Errors

The following errors are sent by the fs2d daemon.

```
Error 9 writing CDB info attribute for node  
#cluster#elaine#machines#cxfs2#Cellular#status
```

An internal error occurred when writing to the cluster configuration database. Retry the operation. If the error persists, stop and restart the cluster daemons:

```
# /etc/init.d/cluster stop  
# /etc/init.d/cluster start
```

If the problem persists, clear the database, reboot, and re-create the database. See "Clearing the Cluster Configuration Database", page 115.

```
Error 9 writing CDB string value for node  
#cluster#elaine#machines#cxfs2#Cellular#status
```

An internal error occurred when writing to the cluster configuration database. Retry the operation. If the error persists, stop and restart the cluster daemons:

```
# /etc/init.d/cluster stop  
# /etc/init.d/cluster start
```

If the problem persists, clear the database, reboot, and re-create the database. See "Clearing the Cluster Configuration Database", page 115.

```
Failed to update CDB for node  
#cluster#elaine#Cellular#FileSystems#fs1#FSStatus
```

An internal error occurred when writing to the cluster configuration database. Retry the operation. If the error persists, stop and restart the cluster daemons:

```
# /etc/init.d/cluster stop  
# /etc/init.d/cluster start
```

If the problem persists, clear the database, reboot, and re-create the database. See "Clearing the Cluster Configuration Database", page 115.

```
Failed to update CDB for node
#cluster#elaine#machines#cxfs2#Cellular#status
```

An internal error occurred when writing to the cluster configuration database. Retry the operation. If the error persists, stop and restart the cluster daemons:

```
# /etc/init.d/cluster stop
# /etc/init.d/cluster start
```

If the problem persists, clear the database, reboot, and re-create the database. See "Clearing the Cluster Configuration Database", page 115.

```
Machine 101 machine_sync failed with lock_timeout error
```

The fs2d daemon was not able to synchronize the cluster configuration database and the sync process timed out. This operation will be automatically retried by fs2d.

## General Messages

The following are general messages.

```
CI_CONFERR_NOTFOUND,  The cluster configuration database has not been
Logging              initialized. To create the initial cluster configuration
configuration        database, enter the following:
error: could not      # /usr/cluster/bin/cbdreinit
read cluster          database
/var/cluster/cdb/cdb.db,
cdb error = 3.
```

```
NOTICE: Error          Error number 4 (EINTR) on MEMBERSHIP message
reading msg           channel (channel 1; channel 0 is the main channel for
header 4 channel 1   CXFS and XVM data) for connection with node 2. The
cell 2                EINTR indicates that this message channel is purposely
                      being torn down and does not indicate an error in itself.
                      (Any other error number is a real error that will cause
```

	the local node to declare the other node failed.) This is an informative message; no corrective action is required.
NOTICE: Membership delivered. Membership contains 0(21) 1(12) cells	Node 0 and node 1 are in the membership; node 0 has been in the last 21 memberships, node 1 has been in the last 12. This is an informative message; no corrective action is required.
NOTICE: Resetting cells 0x4	The number here is a bitmask of node numbers that a reset is being requested on. In this case, 0x4 equates to node 2. This is an informative message; no corrective action is required.
WARNING: Error receiving messages from cell 2 tcpchannel 1	There has been an error on the membership channel (channel 1; channel 0 is the main message channel for CXFS and XVM data). This may be a result of tearing down the channel or may be an error of the node (node with an ID of 2 in this case). There is no corrective action.

## Log File Errors

CXFS maintains logs for each of the CXFS daemons. For information about customizing these logs, see "Set Log Configuration with the GUI", page 68.

Log file messages take the following form:

*daemon\_log timesamp internal\_process: message\_text*

For example:

```
cad_log:Thu Sep 2 17:25:06.092 cclconf_poll_clconfd: clconf_poll failed with error CI_IPCERR_NOPULSE
```

Table 8-2, page 128, shows the parts in the preceding message.

**Table 8-2** Log Error Message Format

Content	Part	Meaning
cad_log	Timestamp	The message pertains to the cad daemon
Sep 2 17:25:06.092	Timestamp and process ID	September 2 at 5:25 PM, process ID 92.
cclconf_poll_clconfd	Internal process information	Internal process information
clconf_poll failed with error CI_IPCERR_NOPULSE	Message text	The clconfd daemon could not be contacted to get an update on the cluster's status.

**cad Messages**

The following are examples of messages from `/var/cluster/ha/log/cad_log`:

```
ccacdb_cam_open: failed to open connection to CAM server error
4
```

Internal message that can be ignored because the cad operation is automatically retried.

```
ccamail_cam_open: failed to open connection to CAM server error
4
```

Internal message that can be ignored because the cad operation is automatically retried.

```
ccicdb_cam_open: failed to open connection to CAM server error
4
```

Internal message that can be ignored because the cad operation is automatically retried.

```
cclconf_cam_open: failed to open connection to CAM server error
4
```

Internal message that can be ignored because the cad operation is automatically retried.

```
cclconf_poll_clconfd: clconf_poll failed with error  
CI_IPCERR_NOCONN
```

The clconfd daemon is not running or is not responding to external requests. If the error persists, stop and restart the cluster daemons:

```
# /etc/init.d/cluster stop  
# /etc/init.d/cluster start
```

```
cclconf_poll_clconfd: clconf_poll failed with error  
CI_IPCERR_NOPULSE
```

The clconfd daemon could not be contacted to get an update on the cluster's status. If the error persists, stop and restart the cluster daemons:

```
# /etc/init.d/cluster stop  
# /etc/init.d/cluster start
```

```
cclconf_poll_clconfd: clconf_poll failed with error  
CI_CLCONFERR_LONELY
```

The clconfd daemon does not have enough information to provide an accurate status of the cluster. It will automatically restart with fresh data and resume its service.

```
csmr_cam_open: failed to open connection to CAM server error 4
```

Internal message that can be ignored because the cad operation is automatically retried.

```
Could not execute notification cmd. system() failed. Error:  
No child processes
```

No mail message was sent because cad could not fork processes. Stop and restart the cluster daemons:

```
# /etc/init.d/cluster stop  
# /etc/init.d/cluster start
```

```
error 3 sending event notification to client 0x000000021010f078
```

GUI process exited without cleaning up.

```
error 8 sending event notification to client 0x000000031010f138
      GUI process exited without cleaning up.
```

### cli Messages

The following are examples of messages from  
`/var/cluster/ha/log/cli_nodename:`

```
CI_CONFERR_NOTFOUND, No machines found in the CDB.
```

The local node is not defined in the configuration database.

```
CI_ERR_INVALID, Cluster (bob) not defined
```

The cluster called bob is not present in configuration database.

```
CI_ERR_INVALID, CLI private command: failed (Cluster (bob) not
defined)
```

The cluster called bob is not present in configuration database.

```
CI_IPCERR_AGAIN, ipccInt_connect(): file
/var/cluster/ha/comm/clconfd-ipc_cxfs0 lock failed - Permission
denied
```

The CLI was invoked by a login other than root. You should only  
use `cmgr(1M)` when you are logged in as root.

```
CI_IPCERR_NOPULSE, CLI private command: failed (Cluster state
is UNKNOWN.)
```

The cluster state could not be determined. Check if the `clconfd(1M)`  
daemon is running.

```
CI_IPCERR_NOPULSE, ipccInt_pulse_internal(): server failed to
pulse
```

The cluster state could not be determined. Check if the `clconfd(1M)`  
daemon is running.

```
CI_IPCERR_NOSEVER, clconf ipc: ipccInt_connect() failed, file
/var/cluster/ha/comm/clconfd-ipc_cxfs0
```

The local node (`cxfs0`) is not defined in the configuration database.

```
CI_IPCERR_NOSEVER, Connection file  
/var/cluster/ha/comm/clconfd-ipc_cxfs0 not present.
```

The local node (cxfs0) is not defined in the configuration database.

### **crsd Errors**

The following are examples of messages from  
/var/cluster/ha/log/crsd\_nodename:

```
CI_CONFERR_INVALID, Nodeid -1 is invalid.
```

The Node ID was not provided in the node definition or the cluster processes were not running in that node when node definition was created in the configuration database. This is a warning that optional information is not available when expected.

```
I_CONFERR_INVALID, Error from ci_security_init().
```

The Node ID was not provided in the node definition or the cluster processes were not running in that node when node definition was created in the configuration database. This is a warning that optional information is not available when expected.

```
CI_ERR_SYS:125, bind() failed.
```

The Node ID was not provided in the node definition or the cluster processes were not running in that node when node definition was created in the configuration database. This is a warning that optional information is not available when expected.

```
CI_ERR_SYS:125, Initialization failed, exiting.
```

The Node ID was not provided in the node definition or the cluster processes were not running in that node when node definition was created in the configuration database. This is a warning that optional information is not available when expected.

```
CI_ERR_NOTFOUND, Nodeid does not have a value.
```

The Node ID was not provided in the node definition or the cluster processes were not running in that node when node definition was created in the configuration database. This is a warning that optional information is not available when expected.

```
CI_CONFERR_INVALID, Nodeid -1 is invalid.
```

System controller information (optional information) was not provided for node `cxfs2`. Provide system controller information for node `cxfs2` by modifying node definition. This is a warning that optional information is not available when expected.

```
CI_ERR_NOTFOUND, SystemController information for node cxfs2 not found, requests will be ignored.
```

System controller information (optional information) was not provided for node `cxfs2`. Provide system controller information for node `cxfs2` by modifying node definition. This is a warning that optional information is not available when expected. Without this information, the node will not be reset if it fails, which might prevent the cluster from properly recovering from the failure.

```
CI_ERR_NOTFOUND, SystemController information for node cxfs0 not found, requests will be ignored.
```

The owner node specified in the node definition for the node with a node ID of 101 has not been defined. You must define the owner node.

```
CI_CRSERR_NOTFOUND, Reset request 0x10087d48 received for node 101, but its owner node does not exist.
```

The owner node specified in the node definition for the node with a node ID of 101 has not been defined. You must define the owner node.

### **fs2d Errors**

The following are examples of messages from `/var/cluster/ha/log/fs2d_nodename`:

```
Failed to copy global CDB to node cxfs1 (1), error 4
```

There are communication problem between the local node and node `cxfs2`. Check the control networks of the two nodes.

Communication failure send new quorum to machine cxfs2 (102)  
(error 6003)

There are communication problem between the local node and node  
cxfs2. Check the control networks of the two nodes.

Failed to copy CDB transaction to node cxfs2 (1)

There are communication problem between the local node and node  
cxfs2. Check the control networks of the two nodes.



---

## Index

### A

- ACTIVE cluster status, 109
- Active guides, 21
- Adaptec host bus adapters, 12
- add nodes
  - GUI, 61
- administration, 97
- administrative shutdown and CXFS cluster services, 105
- Age, 112
- ALERT error, 117
- AutoLoad boot parameter, 34
- automatic restart of nodes, 34

### B

- backups, 101
- BROCADE switches, 12

### C

- CAD options file, 29
- cdbreinit, 115
- cdbutil, 113
- Cell ID, 112
- Channel traffic, 126
- CIPRICO RAID, 12
- CLARiiON RAID, 12
- Clearing the database, 115
- CLI
  - overview, 13
  - tasks, 75
- cluster definition
  - GUI, 64

- cluster environment, 1
- Cluster Manager CLI
  - invoking a shell, 23
- Cluster Manager tools, 13
- cluster membership, 4
- cluster membership status, 111
- cluster node, 1
  - See node, 55
- cluster services, 66
- cluster status, 109
- Cluster View, 16
- configuration tasks
  - CLI, 75
    - cluster tasks, 84
    - configure log groups, 93
    - defaults, 76
    - define a cluster, 84
    - define a node, 77
    - delete a cluster, 89
    - delete a node, 79
    - display a cluster, 90
    - display a node, 82
    - display log parameters, 93
    - modify a cluster, 86
    - modify a node, 78
    - modify log parameters, 94
    - node tasks, 77
    - reset a node, 83
    - set log parameters, 93
    - start cluster services, 91
    - stop cluster services, 92
    - test network connectivity, 95
  - GUI
    - adding nodes to a cluster, 61
    - configure log groups, 69
    - define a cluster, 64
    - define a node, 55

- delete a cluster, 65
- delete a node, 63
- display a cluster, 66, 113
- display a node, 63
- display log group definitions, 69
- modify a cluster, 65
- modify a node, 62
- reset a node, 61
- set log parameters, 68
- setting up a new cluster, 52
- start cluster services, 66
- stop cluster services, 67
- test connectivity, 74

control network

- defining for node, 56

corruption of the root XFS filesystem, 114

CXFS Cluster Manager GUI

- tasksets, 22

CXFS Cluster View, 16

CXFS Manager

- add nodes, 61
- define a node, 55
- delete a node, 63
- display nodes, 63
- modify a node, 62
- overview, 13
- remove nodes, 61
- reset a node, 61
- set up a new cluster, 52
- set up a new filesystem, 54

**D**

- data management, 101
- Database clearing, 115
- Database dump, 113
- Death of a metadata server, 116
- define a cluster
  - GUI, 64
- define and mount filesystems with the CLI, 86
- defragmenter, 10

- delete a cluster
  - GUI, 65
- delete a node
  - GUI, 63
- diagnostic tasks
  - See "configuration tasks", 68
- Disk buffering, 9
- display a cluster
  - GUI, 66
- display nodes
  - GUI, 63
- DOWN node state, 108
- dump, 104
- Dump of the database, 113

**E**

- Emulex Corporation hubs, 12
- /etc/config/cad.options file, 29
- /etc/config/fs2d.options file, 30
- /etc/services file, 29

**F**

Filesystem

- GUI, 54

filesystem maintenance, 102

filesystem reorganizer, 10

Filesystems, mounting with the CLI, 86

fs2d options file, 30

fsr\_xfs, 10

**G**

- growing filesystems, 104
- guaranteed-rate I/O, 10
- GUI
  - active guides, 21

overview, 13  
tasks  
    See "configuration tasks", 51

## H

hardware supported, 12  
HBA, 12  
heartbeat network, 56  
Host bus adapters, 12  
hostname  
    control network, 56

## I

I/O errors, 116  
INACTIVE cluster status, 109  
INACTIVE node state, 108  
Incarnation, 112  
initial cluster configuration  
    CLI, 42  
    GUI, 39  
installation, 25  
installing CXFS software, 25  
IP address  
    control network, 56  
IRISconsole, 12

## L

LAN, 8  
license error, 116  
local area networks (LANs), 8  
log files, 68  
log groups, 68  
log level, 70

## M

maintenance, 97  
maintenance on nodes and CXFS cluster  
    services, 105  
membership, 4  
membership status, 111  
metadata clients, 6  
Metadata server  
    discovering, 97  
metadata server, 6  
Metadata server death, 116  
modify a node  
    GUI, 62  
mount filesystems with the CLI, 86  
mount-point nesting, 11  
mounting, 103

## N

name restrictions, 56  
Named pipes and CXFS, 9  
nesting of mount points, 10  
network connectivity  
    testing with GUI, 74  
network interface  
    configuration, 34  
networked filesystem comparison, 8  
NFS, 102  
node, 1  
    state, 108  
    status, 108  
node tasks  
    See "configuration tasks", 55, 77  
node weighting, 5  
notification  
    GUI, 65  
NVRAM variables, 34

**O**

Origin200, 12  
Origin2000, 12

**P**

Peer-to-disk model, 9  
Platforms, 12  
pool, 2  
Prisa Networks, Inc., 12

**Q**

QLogic Corporation host bus adapters, 12  
quorum, 5  
Quotas, 102  
quotas, 10

**R**

RAID, 12  
RealTime, 102  
reboot, 114  
reboot the system , 39  
remove nodes  
    GUI, 61  
reorganizer, 10  
reset a node, 61  
resetting nodes, 61  
Resetting the database, 115  
resource  
    status, 109  
Restarting CXFS cluster services, 114  
restarting nodes automatically, 34  
restore, 104  
/.rhosts file, 33

**S**

SAMBA, 102  
serial connections  
    testing with GUI, 74  
serial port configuration, 37  
Server death, 116  
SGI 2800/2400 switches, 12  
shutdown and CXFS cluster services, 105  
software supported, 12  
status  
    cluster, 109  
    node, 108  
    resource, 109  
    system controller, 109  
    system, overview, 107  
stop CXFS cluster services, 67  
stopping CXFS, 67  
subsystems installed, 26  
swap to a file, 10  
system controller  
    status, 109  
system files, 29  
system status, 107  
system view, 8

**T**

tasks, 21  
TCP/IP network, 12  
Tools, 13  
Tree view, 16  
troubleshooting, 111

**U**

UNKNOWN node state, 108  
unmounting, 104  
UP node state, 108

**V**

View of the cluster, 16  
volume management, 101

**W**

weighting, 5

**X**

xf<sub>s</sub>\_fsr (fsr\_xfs), 10  
xf<sub>s</sub>\_repair(1M), 114  
xf<sub>s</sub>dump and xf<sub>s</sub>restore, 104  
XVM, 101