



CXFS™ Administration Guide for
SGI® InfiniteStorage

007-4016-026

COPYRIGHT

© 1999–2007 SGI. All rights reserved; provided portions may be copyright in third parties, as indicated elsewhere herein. No permission is granted to copy, distribute, or create derivative works from the contents of this electronic documentation in any manner, in whole or in part, without the prior written permission of SGI.

The following copyright notice applies to the LZF algorithm:

Copyright (c) 2000-2005 Marc Alexander Lehmann <schmorp@schmorp.de> Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. The name of the author may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

LIMITED RIGHTS LEGEND

The software described in this document is "commercial computer software" provided with restricted rights (except as to included open/free source) as specified in the FAR 52.227-19 and/or the DFAR 227.7202, or successive sections. Use beyond license provisions is a violation of worldwide intellectual property laws, treaties and conventions. This document is provided with limited rights as defined in 52.227-14.

TRADEMARKS AND ATTRIBUTIONS

SGI, Altix, the SGI cube, the SGI logo, IRIS, IRIX, O2, Octane, Onyx, Onyx2, Origin, Silicon Graphics Prism, Silicon Graphics Tezro, and XFS are registered trademarks and CXFS, FailSafe, IRIS FailSafe, FDDIXPress, NUMAlink, Octane2, Performance Co-Pilot, Silicon Graphics Fuel, SGI FailSafe, SGI ProPack, and Trusted IRIX are trademarks of SGI in the United States and/or other countries worldwide.

AIX is a registered trademark of IBM Corporation. Brocade is a trademark of Brocade Communication Systems, Inc. Digi is a trademark of Digi International, Inc. FLEXlm is a trademark of Macrovision Corporation. Java and Sun are registered trademarks and Solaris is a trademark of Sun Microsystems, Inc. HP-UX is a registered trademark of Hewlett-Packard Company. Linux is a registered trademark of Linus Torvalds in several countries. Legato NetWorker is a registered trademark of Legato Systems, Inc. Mac OS is a registered trademark of Apple Computer, Inc. Netscape is a trademark of Netscape Communications Corporation. QLogic is a registered trademark of QLogic Corporation. Red Hat is a registered trademark and RPM is a trademark of Red Hat, Inc. UNIX is a registered trademark of The Open Group in the United States and other countries. VERITAS is a trademark of VERITAS Software Corporation. Windows and Windows NT are registered trademarks of Microsoft Corporation. X/Open is a registered trademark of X/Open Company Ltd. All other trademarks mentioned herein are the property of their respective owners.

New Features in This Guide

Note: Be sure to read the release notes for your platforms and the late-breaking caveats page on Supportfolio to learn about any changes to the installation and configuration procedures.

This version contains the following:

- Support for the SGI InfiniteStorage 220 RAID. See "RAID Hardware" on page 63.
- Support for Intelligent Platform Management Interface (IPMI) reset using a baseboard management controller (BMC). To use IPMI and BMC, see the following:
 - "Create or Modify a Node with `cxfs_admin`" on page 270
 - "IPMI Issues" on page 455
 - "BMC System Controller" on page 539
- As of CXFS 4.2, all server-capable nodes running 4.2 and client-only nodes running 4.2 require server-side licensing. If **all** existing client-only nodes are running a prior supported release, they may continue to use client-side license as part of the rolling upgrade policy until they are upgraded to 4.2. All client-only nodes in the cluster must use the same licensing type — if any client-only node in the cluster is upgraded to 4.2 or if a new 4.2 client-only node is added, then all nodes must use server-side licensing.

Customers with support contracts can exchange their existing client-side licenses for new server-side licenses. For more information, contact SGI customer support.

See:

- Chapter 4, "CXFS License Keys" on page 75
- "CXFS Release Versions and Rolling Upgrades" on page 307
- Support for GPT-labeled LUNs larger than 2 TB. (All nodes that mount a filesystem using LUNs larger than 2 TB must be upgraded to CXFS 4.2 or later.) See Appendix E, "Filesystem Specifications" on page 537.
- A new section "Disk Layout Optimization for Approved Media Customers" on page 161

- If you have multiple clusters using the same public network as the backup CXFS metadata network, use the `-i` option to `cxfs_admin` to identify the cluster name. See "Accessing the Correct Cluster at a Multiple-Cluster Site" on page 269.
- "Precedence of Configuration Options" on page 307
- Support for printing hafdence debug information to the specified file `debugfile` by using the `-d` option in the `/etc/cluster/config/clconfd.options` file. See "clconfd.options on CXFS Administration Nodes" on page 133.
- A new section "SGI ProPack: Using `cxfs-reprobe` on Client-Only Nodes" on page 136
- Information about parameters that must be set for QLogic switches. See "QLogic Fibre Channel Switch" on page 72.
- The ability to use environment variables or the `.cxfs_admin` file to specify defaults for `cxfs_admin`, in addition to the `set` command. See "Setting `cxfs_admin` Defaults" on page 266.
- Documentation for the support of XVM failover version 2 on Windows (first supported in the CXFS 4.1.1 release). See "XVM Failover Version 2 (V2) for MultiOS Clusters" on page 46.
- A new section that describes how to view the current CXFS licenses with the `cxfs_admin` command. See "Show License Information with `cxfs_admin`" on page 286.
- "clconfd.options on CXFS Administration Nodes" on page 133
- Information about the `cmgr` command has been moved to Appendix G, "Reference to `cmgr` Tasks" on page 557. With the exception of performing the following administrative `cmgr` commands, the preferred CXFS configuration tools are `cxfs_admin` and the CXFS graphical user interface (GUI):

```
admin ping
admin reset
admin powerCycle
```

As of the CXFS 5.0 release, this functionality will be provided by the `cxfs_admin` command and the `cmgr` command will not be supported.

Record of Revision

Version	Description
001	September 1999 Supports the CXFS 1.1 product in the IRIX 6.5.6f release.
002	October 1999 Supports the CXFS 1.1 product in the IRIX 6.5.6f release.
003	December 1999 Supports the CXFS product in the IRIX 6.5.7f release.
004	March 2000 Supports the CXFS product in the IRIX 6.5.8f release.
005	June 2000 Supports the CXFS product in the IRIX 6.5.9f release.
006	September 2000 Supports the CXFS product in the IRIX 6.5.10f release.
007	January 2001 Supports the CXFS product in the IRIX 6.5.11f release.
008	March 2001 Supports the CXFS product in the IRIX 6.5.12f release.
009	June 2001 Supports the CXFS product in the IRIX 6.5.13f release.
011	September 2001 Supports the CXFS product in the IRIX 6.5.14f release. (Note, there was no 010 version due to an internal numbering mechanism.)
012	December 2001 Supports the CXFS Version 2 product in IRIX 6.5.15f.
013	March 2002 Supports the CXFS Version 2 product in IRIX 6.5.16f.

014	June 2002 Supports the CXFS Version 2 product in IRIX 6.5.17f.
015	September 2002 Supports the CXFS Version 2 product in IRIX 6.5.18f.
016	December 2002 Supports the CXFS Version 2 product in IRIX 6.5.19f.
017	March 2003 Supports the CXFS Version 2 product in IRIX 6.5.20f.
018	September 2003 Supports the CXFS 3.0 product in IRIX 6.5.22 and CXFS 3.0 for SGI Altix 3000 running SGI ProPack 2.3 for Linux.
019	December 2003 Supports the CXFS 3.1 product in IRIX 6.5.23 and CXFS 3.1 for SGI Altix 3000 running SGI ProPack 2.4 for Linux.
020	March 2004 Supports the CXFS 3.2 product in IRIX 6.5.24 and CXFS 3.2 for SGI Altix 3000 running SGI ProPack 3 for Linux.
021	November 2004 Supports the CXFS 3.2 product in IRIX 6.5.24 and CXFS 3.2 for SGI Altix 3000 running SGI ProPack 3 for Linux.
022	April 2005 Supports the CXFS 3.3 product
023	July 2005 Supports the CXFS 3.4 product
024	May 2006 Supports the CXFS 4.0 product
025	January 2007 Supports the CXFS 4.1 product

026

September 2007
Supports the CXFS 4.2 product

Contents

About This Guide	xliii
Related Publications	xliii
Obtaining Publications	xlvii
Conventions	xlvii
Reader Comments	xlviii
1. Introduction to CXFS	1
What is CXFS?	2
Comparison of XFS and CXFS	3
Supported XFS Features	4
When to Use CXFS	5
Performance Considerations	6
Comparison of Network and CXFS Filesystems	7
Network Filesystems	7
CXFS Filesystems	8
Features	8
Restrictions	9
Cluster Environment	9
Terminology	10
Cluster	10
Node	10
Pool	11
Cluster Database	12
Node Functions	12
Membership	19

Private Network	19
Relocation	20
Recovery	22
CXFS Tiebreaker	23
Cluster Administration Daemons	24
CXFS Services	25
CXFS Control Daemon	26
CXFS Client Daemon	26
Forced CXFS Shutdown	27
Heartbeat Monitoring	27
Isolating Failed Nodes: Failure Policies	28
Fence	30
Reset	34
Shutdown	35
The Cluster Database and CXFS Clients	36
Metadata Server Functions	36
System View	37
CXFS and Highly Available Services	38
Hardware and Software Support	38
Requirements	38
Compatibility	41
Overview of FailSafe Coexecution	41
CXFS Tools Overview	42
Guaranteed-Rate I/O (GRIO) Version 2 and CXFS	44
XVM Failover and CXFS	45
XVM Failover Version 1 (V1) for Clusters with Only IRIX Nodes	45
XVM Failover Version 2 (V2) for MultiOS Clusters	46
Overview of XVM Failover V2	47
Configuring XVM Failover V2 using <code>failover2.conf</code>	48

Generating a failover2.conf File	49
Failover V2 Example Using Two Affinities	50
Failover V2 Example Using Four Affinities	51
Failover V2 Example for Two RAID Controllers	52
XVM Commands Related to Failover V2	52
RAID Units and XVM Failover V2	52
TP9100, RM610/660, and Failover V2	53
TP9300, TP9500, TP9700, and S330 and Failover V1 and V2	53
SGI InfiniteStorage 220 and Failover V2	54
GPT Labels and CXFS	54
Installation and Configuration Overview	56
Client-Only Packages and Commands Installed	56
Administration Packages and Commands Installed	57
Installation Procedure Overview	61
2. SGI RAID for CXFS Clusters	63
RAID Hardware	63
RAID Firmware	64
Number of LUNs Supported	65
RAID Verification	66
3. Switches	67
Brocade Switch	67
Brocade Firmware	67
Verifying the Brocade Switch Firmware Version	69
Verifying the Brocade License	69
Limiting telnet Sessions	70
Brocade 200E/3250/3252/3850/3852/3900/4100/4900/5000 and telnet	70
Brocade 12000/24000/48000 and telnet	70

Changing the Brocade FC Cable Connections	71
QLogic Fibre Channel Switch	72
4. CXFS License Keys	75
Server-Side Licensing Overview	76
Licensing Requirements	76
Server-Side Client License Keys	77
License Key Replication on Server Nodes	78
Cumulative Client License Keys	78
Examples of License Keys Required for Cluster Changes	78
Licensing on IRIX Clusters	82
Gathering the Host Information for IRIX	82
Obtaining the License Keys for IRIX	82
Installing the License Keys on IRIX	83
Licensing on SGI ProPack 5 Clusters	83
Gathering the Host Information for SGI ProPack 5	83
Obtaining the License Keys from SGI for SGI ProPack 5	84
Installing the License Keys on SGI ProPack 5	84
Verifying the License Keys	84
Verifying the License Keys with <code>cxfslicense</code> on Server-Capable Nodes	84
Verifying the LK License Keys with <code>lk_verify</code>	88
Verifying the XVM Mirror Licenses on Client-Only Nodes	91
For More Information About Licensing	92
5. Preinstallation Steps	93
Hostname Resolution and Network Configuration Rules	93
Configuring Network Interfaces	94

Adding a Private Network	94
Configuring IRIX Interfaces	97
Verifying the Private and Public Networks	98
Configuring the Serial Ports for IRIX Administration Nodes	99
6. IRIX CXFS Installation	101
IRIX Administration Software Installation	102
IRIX Client-only Software Installation	107
IRIX Modifications for CXFS Connectivity Diagnostics	111
7. SGI ProPack CXFS Installation	113
SGI ProPack Limitations and Considerations	114
Limitations and Considerations for Any SGI ProPack Node	114
Limitations and Considerations for SGI ProPack Client-Only Nodes	116
SGI ProPack Administration Software Installation	116
SGI ProPack Client-Only Software Installation	120
SGI ProPack Client-Only Installation Overview	120
Installing the Performance Co-Pilot Agent	123
SGI ProPack Installation Verification	124
SGI ProPack Modifications for CXFS Connectivity Diagnostics	124
8. Postinstallation Steps	127
Configuring System Files	127
/etc/exports on All Nodes	127
Administration Node System Files	128
/etc/services on CXFS Administration Nodes	128
cad.options on CXFS Administration Nodes	129
fs2d.options on CXFS Administration Nodes	130
Example 1	132

Example 2	132
clconfd.options on CXFS Administration Nodes	133
Client-only Node System Files	134
IRIX: Configuring for Automatic Restart	135
IRIX: Converting Filesystem Definitions for Upgrades	135
Upgrading from 6.5.12f or Earlier	135
Running with All IRIX Nodes Upgraded to 6.5.14f or Later	136
SGI ProPack: Using <code>cxfs-reprobe</code> on Client-Only Nodes	136
9. Best Practices	139
Configuration Best Practices	139
Fix Network Issues First	140
Use a Private Network	140
Provide Enough Memory	141
Use CXFS Configuration Tools Appropriately	141
Use Netscape and the Brocade Switch GUI on IRIX Appropriately	142
Ensure Cluster Database Membership Quorum Stability	142
Be Consistent in Configuration	143
Do Not Mix Metadata Operating System Flavors	143
Use the Correct Mix of Software Releases	143
Form a Small Functional Cluster First	144
Choose a Metadata Server that is Dedicated to CXFS Work	144
Use an Odd Number of Server-Capable Nodes	145
Make Most Nodes Client-Only	145
Use a Client-Only Tiebreaker	145
Protect Data Integrity on All Nodes	146
System Reset	147
I/O Fencing	147

Avoid Network Partition	148
Minimize the Number of Switches	149
Configure Filesystems Properly	149
Verify the Configuration	151
Use the Recovery Timeout Mechanism	151
Use Proper Storage Management Procedures	151
Administration Best Practices	152
Do Not Run User Jobs on Metadata Servers	153
Do Not Run Backups on a Client Node	153
Use cron Jobs Properly	153
Modify updatedb to Avoid Unnecessary Load	154
Repair Filesystems with Care	154
Defragment Filesystems with Care	155
Use Relocation and Recovery Properly	155
Shut Down Nodes Unobtrusively	155
Remove Unused Cluster Components	156
Use fam Properly	156
Use Trusted IRIX Consistently	157
Upgrade the Software Properly	157
Use Fast Copying for Large CXFS Files	157
Log File Names and Sizes	157
Use System Capacity Wisely	158
Reboot Before Changing Node ID or Cluster ID	158
Restart CXFS on a Node after an Administrative CXFS Stop	158
Restart the Cluster In an Orderly Fashion	159
Disable Reset Capability If You Remove Reset Lines	159
Avoid Performance Problems with Unwritten Extent Tracking and Exclusive Write Tokens	159

Unwritten Extent Tracking	159
Exclusive Write Tokens	160
Use the Appropriate Version of <code>lcrash</code> for SGI ProPack	161
Disk Layout Optimization for Approved Media Customers	161
10. Initial Setup of the Cluster	163
Preliminary Cluster Configuration Steps	164
Verify the License	164
Verify that the Cluster Daemons are Running	164
Gather the Required Information	165
Configure for <code>nsd</code> Use (Optional)	166
Verify that the <code>chkconfig</code> Arguments are On	166
IRIX <code>chkconfig</code> Verification	166
SGI ProPack <code>chkconfig</code> Verification	167
Initial Setup Using One of the Configuration Tools	167
Initial Setup with the CXFS GUI	167
Start the GUI	168
Set Up a New Cluster with the GUI	169
Set Up a New CXFS Filesystem with the GUI	171
Initial Setup with the <code>cxfs_admin</code> Command	172
Configuring a Large Cluster	176
Testing the System	178
Private Network Interface	178
System Reset Connection for CXFS Administration Nodes	178
Testing Serial Connectivity for the L2 on Altix 350 Systems	180
11. Reference to GUI Tasks	183
GUI Overview	183

Starting the GUI	184
Starting the GUI on IRIX	184
Starting the GUI on SGI ProPack	185
Running the Web-based Version	185
Running as a Non-Root User on IRIX	187
Running the GUI from an IRIX Desktop Outside the Cluster	187
Summary of GUI Platforms	187
Logging In	188
Making Changes Safely	189
GUI Windows	189
GUI Features	191
GUI Window Layout	192
File Menu	193
Edit Menu	193
Tasks Menu	193
Help Menu	194
Shortcuts Using Command Buttons	194
View Menu	196
Selecting Items to View or Modify	196
Viewing Component Details	197
Performing Tasks	197
Using Drag-and-Drop	198
Analyzing I/O Performance with Performance Co-Pilot on an IRIX Node	198
Structuring Volume Topologies	199
Configuring Disks	200
Getting More Information	200
Important GUI and xvm Command Differences	200

Key to Icons and States	201
Guided Configuration Tasks	205
Set Up an Existing FailSafe Cluster for CXFS with the GUI	205
Make Changes to Existing Cluster	206
Fix or Upgrade Cluster Nodes	207
Node Tasks with the GUI	207
Define a Node with the GUI	208
Examples of Defining a Node with the GUI	214
Add or Remove Nodes in the Cluster with the GUI	216
Reset a Node with the GUI	217
Modify a Node Definition with the GUI	217
Convert a FailSafe Node for use with CXFS with the GUI	221
Delete a Node with the GUI	222
Test Node Connectivity with the GUI	223
Display a Node with the GUI	223
Cluster Tasks with the GUI	223
Define a Cluster with the GUI	224
Modify a Cluster Definition with the GUI	225
Convert a FailSafe Cluster for use with CXFS with the GUI	226
Delete a Cluster with the GUI	227
Display a Cluster with the GUI	227
Cluster Services Tasks with the GUI	227
Start CXFS Services with the GUI	227
Stop CXFS Services with the GUI	228
Set Tiebreaker Node with the GUI	229
Set Log Configuration with the GUI	230
Display Log Group Definitions with the GUI	230
Configure Log Groups with the GUI	231

Revoke Membership of the Local Node with the GUI	232
Allow Membership of the Local Node with the GUI	232
Switches and I/O Fencing Tasks with the GUI	233
Define a Switch with the GUI	233
Modify a Switch Definition with the GUI	235
Update Switch Port Information with the GUI	236
Delete a Switch Definition with the GUI	236
Raise the I/O Fence for a Node with the GUI	237
Lower the I/O Fence for a Node with the GUI	237
Filesystem Tasks with the GUI	237
Make Filesystems with the GUI	238
Grow a Filesystem with the GUI	240
Define CXFS Filesystems with the GUI	241
Modify a CXFS Filesystem with the GUI	244
Mount CXFS Filesystems with the GUI	245
Unmount CXFS Filesystems with the GUI	246
Mount a Filesystem Locally	246
Unmount a Local Filesystem	247
Delete a CXFS Filesystem with the GUI	247
Remove Filesystem Mount Information	247
Relocate a Metadata Server for a CXFS Filesystem with the GUI	248
Privileges Tasks with the GUI	249
Grant Task Access to a User or Users	249
Granting Access to a Few Tasks	250
Granting Access to Most Tasks	251
Revoke Task Access from a User or Users	252
12. Reference to <code>cxfs_admin</code> Tasks	255

<code>cxfs_admin</code> Overview	255
Command Syntax Overview	256
Getting Help	258
Making Changes Safely	259
Basic and Advanced Mode	260
Using Prompting Mode	262
Command History	263
Waiting for Commands to Complete	264
Entering <code>cxfs_admin</code> Commands on the Command Line	264
Using Script Files	265
Setting <code>cxfs_admin</code> Defaults	266
Setting <code>cxfs_admin</code> Access Permissions	268
Accessing the Correct Cluster at a Multiple-Cluster Site	269
Exiting from <code>cxfs_admin</code>	269
Node Tasks with <code>cxfs_admin</code>	270
Create or Modify a Node with <code>cxfs_admin</code>	270
Delete a Node with <code>cxfs_admin</code>	279
Enable a Node with <code>cxfs_admin</code>	279
Disable a Node with <code>cxfs_admin</code>	279
Show Node Information with <code>cxfs_admin</code>	280
Cluster Tasks with <code>cxfs_admin</code>	282
Create or Modify a Cluster with <code>cxfs_admin</code>	282
Create a Tiebreaker with <code>cxfs_admin</code>	284
Delete a Cluster with <code>cxfs_admin</code>	285
Display a Cluster with <code>cxfs_admin</code>	285
Show License Information with <code>cxfs_admin</code>	286
CXFS Filesystem Tasks with <code>cxfs_admin</code>	287

Create or Modify a CXFS Filesystem with <code>cxfs_admin</code>	287
Mount a CXFS Filesystem with <code>cxfs_admin</code>	292
Unmount a CXFS Filesystem with <code>cxfs_admin</code>	293
Relocate the Metadata Server for a Filesystem with <code>cxfs_admin</code>	293
Delete a CXFS Filesystem with <code>cxfs_admin</code>	294
Show a CXFS Filesystem	294
Network Failover Tasks with <code>cxfs_admin</code>	295
Switch Tasks with <code>cxfs_admin</code>	295
Create a Switch with <code>cxfs_admin</code>	296
Delete a Switch Definition with <code>cxfs_admin</code>	297
Show Switches with <code>cxfs_admin</code>	297
Saving and Recreating the Current Configuration with <code>cxfs_admin</code>	300
13. Administration and Maintenance	305
Precedence of Configuration Options	307
CXFS Release Versions and Rolling Upgrades	307
General Upgrade Procedure	308
Example Upgrade Process	310
Upgrading From 3.4.1 Or Earlier	314
Manual CXFS Start/Stop	315
CXFS and Cluster Administration Initialization Commands	316
Using <code>hafence</code> to Manipulate a Switch	318
CXFS Port Usage	320
CXFS <code>chkconfig</code> Arguments	320
Configuring Real-Time Filesystems For IRIX Nodes	323
Granting Task Execution Privileges to Users	323
Transforming an Existing Node into a Client-Only Node	324

CXFS Mount Scripts	325
Unmounting <code>lofs</code> File Systems	327
Using <code>telnet</code> and I/O Fencing	327
Using <code>fsr</code> and <code>xfss_fsr</code>	328
Using <code>cron</code> in a CXFS Cluster	328
Using Hierarchical Storage Management (HSM) Products	328
Discovering the Active Metadata Server for a Filesystem	329
Discovering the Active Metadata Server with the CXFS GUI	329
Discovering the Active Metadata Server with <code>cxfs_admin</code>	330
Discovering the Active Metadata Server with <code>clconf_info</code>	331
Metadata Server Recovery	332
Shutdown of the Database and CXFS	332
Cluster Database Shutdown	333
Node Status and Cluster Database Shutdown	334
Restart the Cluster Database	335
Normal CXFS Shutdown: Stop CXFS Services	335
Node Status and Stopping CXFS Services	336
When You Should Not Perform Stop CXFS Services	336
Rejoining the Cluster after Stopping CXFS Services	336
Forced CXFS Shutdown: Revoke Membership of Local Node	337
Node Status and Forced CXFS Shutdown	337
Rejoining the Cluster after a Forced CXFS Shutdown	338
Reset Capability and a Forced CXFS Shutdown	338
Avoiding a CXFS Restart at Reboot	338
Log File Management	339
Rotating Log Files on SGI ProPack	340
Rotating All Log Files on IRIX Server-Capable Nodes	340

Rotating Large Log Files on IRIX	341
Volume Management	341
Disk Management	342
Disk Backups	342
NFS	343
Quotas	343
Samba	343
Filesystem Maintenance	344
Mounting Filesystems	345
Unmounting Filesystems	345
Growing Filesystems	346
Dump and Restore	346
Site-Changeable System Tunable Parameters	348
Site-Changeable Static Parameters	349
mtcp_hb_period	349
mtcp_hb_watchdog	350
mtcp_nodelay	350
mtcp_rpc_thread	350
rhelpd_max	351
rhelpd_min	351
Site-Changeable Dynamic Parameters	352
cms_local_fail_action	352
cxfs_client_push_period	352
cxfs_dcvn_timeout	353
cxfs_extents_delta	353
cxfs_punch_hole_restrict	353
cxfs_relocation_ok	354

Contents

cxfs_server_push_period	354
cxfsd_max	355
cxfsd_min	355
mtcp_mesg_validate	356
Restricted System Tunable Parameters	356
Restricted Static Parameters	357
cxfs_extents_block_size	357
cxfs_extents_delta_depth	357
cxfs_shutdown_time	358
mesg_delay_time	358
mtcp_reserve_size	358
Restricted Dynamic Parameters	359
cell_tkm_feature_disable	359
cms_fence_timeout	360
cms_fence_timeout_action	360
cms_reset_error_override	361
cms_reset_timeout	361
cms_reset_timeout_action	361
cms_trace_enable	362
cxfs_recovery_slowdown	362
cxfs_recovery_timeout_panic	363
cxfs_recovery_timeout_period	363
cxfs_recovery_timeout_stalled	363
cxfs_recovery_timeout_start	364
cxfs_token_fault_tolerant	364
cxfs_token_track	365
cxfs_validate_objid	366
cxfs_verify_existence_token	366

<code>cxfsd_sync_force</code>	367
Hardware Changes and I/O Fencing	367
Configuring Private Network Failover	368
Removing and Restoring Cluster Members	369
Removing a Metadata Server from the Cluster	369
Restoring a Metadata Server to the Cluster	371
Removing a Single Client-Only Node from the Cluster	371
Restoring a Single Client-Only Node to the Cluster	373
Stopping CXFS for the Entire Cluster	374
Restarting the Entire Cluster	374
Discovering the WWNs	375
Mapping XVM Volumes to Storage Targets	375
Mapping XVM Volumes to Storage Targets on IRIX	375
Mapping XVM Volumes to Storage Targets on SGI ProPack	376
Switching Between SGIRDAC and SGI AVT Mode for SGI RAID	376
14. Cluster Database Management	379
Performing Cluster Database Backup and Restoration	379
Restoring a Deleted Database from Another Node	379
Using <code>cxfs_admin</code> to Recreate the Cluster Database	381
Using <code>cdbBackup</code> and <code>cdbRestore</code> for the Cluster Database and Logging Information	381
Checking the Cluster Configuration with <code>cxfs-config</code>	382
15. Coexecution with FailSafe	389
Why Run CXFS and FailSafe Together?	389
Coexecution Release Levels	390
Size of the Coexecution Cluster	390
Cluster Type	390

Metadata Server Node Types	392
Separate GUIs	392
Conversion	392
Network Interfaces	393
Metadata Servers and Failover Domain	393
CXFS Resource Type for FailSafe	393
Core-Dump Completion	395
16. Trusted IRIX and CXFS	397
Installation Tips for CXFS and Trusted IRIX	397
Mandatory Access Controls	398
17. Monitoring Status	399
Status in Log Files	400
Cluster Status	402
Check Cluster Status with the GUI	402
Check Cluster Status with <code>clconf_info</code>	403
Check Cluster/Node/Filesystem/Switch Status with <code>cxfs_admin</code>	404
Check Cluster/Node/Filesystem Status with <code>cxfs_info</code>	405
Node Status	406
Monitoring Node Status with the GUI	406
Monitoring Node Status with <code>cxfs_admin</code>	407
Monitoring Node Status with <code>clconf_info</code>	407
XVM Statistics	408
I/O Fencing Status	409
Heartbeat Timeout Status	411
18. Migration from an IRIX Cluster to an SGI ProPack Cluster	415
An Overview of Differences Between IRIX and Linux System Administration	415

Migration Procedure	416
Migration Troubleshooting	423
Filesystems Will Not Mount	423
DMF Filesystems Will Not Mount	424
Do Not Use <code>extlog</code> or <code>rtfs</code> Filesystems	424
19. Troubleshooting	425
Troubleshooting Strategy	425
Know the Troubleshooting Tools	426
Physical Storage Tools	426
Cluster Configuration Tools	429
Cluster Control Tools	430
Networking Tools	431
Cluster/Node Status Tools	431
Performance Monitoring Tools	433
Kernel Status Tools	434
Log Files	437
Gather Cluster Configuration with <code>cxfsdump</code>	437
Identify the Cluster Status	438
Eliminate a Residual Cluster	440
Determine If a Node Is Fenced	440
Locate the Problem	441
Redirect Switch Logs	442
Common Problems	442
Client Membership Loss	443
Node is Permanently Fenced	446
Cannot Access Filesystem	446
Log Files Consume Too Much Disk Space	446

Unable to Define a Node	446
System is Hung	447
Node is Detected but Never Joins Membership	447
Cell ID Count and Membership delivered Messages	447
You Cannot Log In	448
I/O Error in Filesystem	448
Cannot Mount Filesystems	449
GUI Displays Invalid Filesystems	449
Multiple client_timeout Values	449
No HBA WWPNs are Detected	450
XFS Internal Errors in System Log File	452
Multiple Ethernet Interfaces on Altix Systems	452
Clients Unable to Remount Filesystems	453
Forced Filesystem Shutdown Messages and XFS File Corruption	453
GUI Will Not Run	454
IPMI Issues	455
BMC Does Not Respond to a ping Command	455
ipmitool Command Fails	455
Node is Not Reset	457
cxfs_admin Output is Not Current	457
clconfd Is Not Running	457
Understanding Error Messages	458
Normal Messages	458
Relocation Error	461
Controller Disable Messages	461
CMS Error Messages	461
clconfd Daemon Death	462

Out of Logical Swap Space	462
No Cluster Name ID Error	463
Lost CXFS Membership	464
License Key Error	464
IP Address Error	465
System Log File Errors	466
cli Error Messages	468
clconfd Error Messages	468
crsd Error Messages	472
cmond Error Messages	473
cxfslicense Error Message	474
fs2d Error Messages	475
General Messages	476
Log File Error Messages	476
cad Messages	477
cli Messages	478
crsd Errors	479
fs2d Errors	480
cdbreinit Error Messages	481
Messages During Remote Installation	481
cxfs_admin Errors	482
Mount Errors	483
SYSLOG credid Warnings	483
Corrective Actions	484
Restarting CXFS Services	484
Clearing the Cluster Database	485
Rebooting	486

Recovering a Two-Node Cluster	486
Rebooting without Rejoining the Cluster	488
Stopping and Restarting Cluster Administration Daemons	489
Recreating the Cluster Database	489
Verifying Connectivity in a Multicast Environment	490
Perform a Power Cycle on a Node	491
Reset a Node	491
Reporting Problems to SGI	492
Reporting IRIX Problems	492
Reporting SGI ProPack Problems	493
Appendix A. CXFS Software Architecture	495
Daemons	495
Communication Paths	498
Communication Paths in a Coexecution Cluster	503
Flow of Metadata for Reads and Writes	504
Appendix B. Memberships and Quorums	509
Membership Types	510
Cluster Database Membership and Quorum	511
CXFS Kernel Membership, Quorum, and Tiebreaker	512
FailSafe Membership, Quorum, and Tiebreaker	514
Cluster Database Membership Logs	515
Quorum and Tiebreaker Examples	520
Changing CXFS Kernel Membership Quorum Example	520
Coexecution Example	523
CXFS Tiebreaker Node Example	524
Heartbeat Considerations	526

CXFS Recovery Issues in a Cluster with Only Two Server-Capable Nodes	526
Appendix C. IP Filtering for the CXFS Private Network	529
Appendix D. Operating System Path Differences	533
Appendix E. Filesystem Specifications	537
Appendix F. System Reset Configuration	539
BMC System Controller	539
L2 System Controller	545
L1 System Controller	551
Redirecting the Console for Origin 300, Origin 350, Origin 3200C, Onyx 300, Onyx 350, and Onyx 3200C	552
MSC System Controller	553
MMSC System Controller	556
Appendix G. Reference to cmgr Tasks	557
cmgr Overview	558
Making Changes Safely	558
Getting Help	558
Using Prompt Mode	559
Completing Actions and Cancelling	561
Using Script Files	561
Invoking a Shell from within cmgr	563
Entering Subcommands on the Command Line	563
Template Scripts	564
Initial Setup with the cmgr Command	564
Set Configuration Defaults with cmgr	576

Node Tasks with <code>cmgr</code>	577
Define a Node with <code>cmgr</code>	577
Modify a Node with <code>cmgr</code>	587
Example of Partitioning	588
Changing Failure Hierarchy	589
Perform an NMI on a Node with <code>cmgr</code>	591
Convert a Node to CXFS or FailSafe with <code>cmgr</code>	592
Delete a Node with <code>cmgr</code>	593
Display a Node with <code>cmgr</code>	595
Test Node Connectivity with <code>cmgr</code>	596
Test the Serial Connections with <code>cmgr</code>	597
Cluster Tasks with <code>cmgr</code>	597
Define a Cluster with <code>cmgr</code>	597
Modify a Cluster with <code>cmgr</code>	601
Convert a Cluster to CXFS or FailSafe with <code>cmgr</code>	602
Delete a Cluster with <code>cmgr</code>	603
Display a Cluster with <code>cmgr</code>	604
Cluster Services Tasks with <code>cmgr</code>	605
Start CXFS Services with <code>cmgr</code>	605
Stop CXFS Services with <code>cmgr</code>	605
Set the Tiebreaker Node with <code>cmgr</code>	606
Set Log Configuration with <code>cmgr</code>	608
Display Log Group Definitions with <code>cmgr</code>	608
Configure Log Groups with <code>cmgr</code>	608
Modify Log Groups with <code>cmgr</code>	609
Revoke Membership of the Local Node with <code>cmgr</code>	610
Allow Membership of the Local Node with <code>cmgr</code>	610

CXFS Filesystem Tasks with <code>cmgr</code>	611
Define a CXFS Filesystem with <code>cmgr</code>	611
Mount a CXFS Filesystem with <code>cmgr</code>	618
Unmount a CXFS Filesystem with <code>cmgr</code>	619
Modify a CXFS Filesystem with <code>cmgr</code>	619
Relocate the Metadata Server for a Filesystem with <code>cmgr</code>	623
Delete a CXFS Filesystem with <code>cmgr</code>	624
Switches and I/O Fencing Tasks with <code>cmgr</code>	624
Define a Switch with <code>cmgr</code>	624
Modify a Switch Definition with <code>cmgr</code>	626
Raise the I/O Fence for a Node with <code>cmgr</code>	626
Lower the I/O Fence for a Node with <code>cmgr</code>	627
Update Switch Port Information with <code>cmgr</code>	627
Delete a Switch Definition with <code>cmgr</code>	627
Show Switches with <code>cmgr</code>	628
Query Switch Status with <code>cmgr</code>	628
Script Example	629
Creating a <code>cmgr</code> Script Automatically	632
Troubleshooting <code>cmgr</code>	636
Additional <code>cmgr</code> Examples	636
Example of Normal CXFS Shutdown Using <code>cmgr</code>	637
Example of Forced CXFS Shutdown Using <code>cmgr</code>	637
Example of Rejoining the Cluster after a Stopping CXFS Services Using <code>cmgr</code>	638
Example of Rejoining the Cluster after a Forced CXFS Shutdown Using <code>cmgr</code>	638
Example of Configuring Private Network Failover Using <code>cmgr</code>	638
Example of Configuring a Large Cluster Using <code>cmgr</code>	645
Example of Performing a Forced CXFS Shutdown Using <code>cmgr</code>	646

Example of Relocation Error Using `cmgr` 646

Example of Migration from an IRIX Cluster to an SGI ProPack Cluster Using `cmgr` . . . 646

Example of Checking Cluster Status Using `cmgr` 654

Example of Querying Node Status Using `cmgr` 654

Example of Pinging the System Controller Using `cmgr` 654

Example of Monitoring Reset Lines Using `cmgr` 655

Example of I/O Fencing Status Using `cmgr` 655

Example of Using `build_cmgr_script` to Recreate the Cluster Database 655

Appendix H. Migration from `cmgr` to `cxfs_admin` 659

Appendix I. Initial Configuration Checklist 661

Appendix J. Summary of New Features from Previous Releases 663

CXFS Version 1: Original Implementation 663

 IRIX 6.5.6f 663

 IRIX 6.5.6f 663

 IRIX 6.5.7f 663

 IRIX 6.5.8f 664

 IRIX 6.5.9f 664

 IRIX 6.5.10f 664

 IRIX 6.5.11f 665

CXFS Version 2: MultiOS Cluster 665

 IRIX 6.5.12f 665

 IRIX 6.5.13f 666

 IRIX 6.5.14f 667

 IRIX 6.5.15f 668

 IRIX 6.5.16f 669

 IRIX 6.5.17f 670

 IRIX 6.5.18f 671

IRIX 6.5.19f	672
IRIX 6.5.20f	674
CXFS Version 3: IRIX or SGI ProPack (Linux 2.4 Kernel) Servers	674
CXFS 3.0	675
CXFS 3.1	676
CXFS 3.2	676
CXFS 3.3	677
CXFS 3.4	678
CXFS Version 4: IRIX or SGI ProPack (Linux 2.6 Kernel) Servers	678
CXFS 4.0	678
CXFS 4.1	681
Glossary	683
Index	697

Figures

Figure 1-1	Pool and Cluster Concepts	11
Figure 1-2	Installation Differences	14
Figure 1-3	Evenly Distributed Metadata Servers	15
Figure 1-4	Multiple Metadata Servers	16
Figure 1-5	One Metadata Server	17
Figure 1-6	Standby Mode	18
Figure 1-7	Relocation versus Recovery	23
Figure 1-8	I/O Fencing	33
Figure 1-9	Contents of a Coexecution Cluster Database	42
Figure 1-10	Disk Paths	47
Figure 4-1	Server-Side License Keys	79
Figure 4-2	Server-Side License Keys: Adding a New Client-Only Node	80
Figure 10-1	CXFS Manager	169
Figure 11-1	CXFS Manager GUI Showing Details for a Node	190
Figure 11-2	Pop-up Menu that Appears After Clicking the Right Mouse Button	191
Figure 11-3	Example Node Definition	215
Figure 11-4	Example System Reset Settings	216
Figure 11-5	Bit Mask Representation for I/O Fencing	235
Figure 11-6	Task Privileges for a Specific User	251
Figure 13-1	Example Rolling Upgrade Procedure (part 1)	311
Figure 13-2	Example Rolling Upgrade Procedure (part 2)	312
Figure 13-3	Example Rolling Upgrade Procedure (part 3)	313
Figure 13-4	Window Showing the Metadata Server	330

Figure 13-5	Value Bit Interpretation	365
Figure 15-1	Cluster and Node Type Combinations	391
Figure 17-1	pmgxvm chart	409
Figure A-1	Communication within One Administration Node	498
Figure A-2	Daemon Communication within One Administration Node	499
Figure A-3	Communication between Nodes in the Pool	500
Figure A-4	Communication for an Administration Node Not in a Cluster	501
Figure A-5	Communication Among Administration Nodes and Client-Only Nodes	502
Figure A-6	Administrative Communication within One Administration Node under Coexecution	503
Figure A-7	Daemon Communication within One Administration Node under Coexecution	504
Figure A-8	Metadata Flow on a Write	505
Figure A-9	Metadata Flow on a Read on Client B Following a Write on Client A	506
Figure A-10	Metadata Flow on a Read on Client B Following a Read on Client A	507
Figure B-1	One Node is Out of Date: Most Recent Log is Replicated	517
Figure B-2	Unequally Sized Pools are Joined: Log from Larger Pool is Replicated	518
Figure B-3	Equally Sized Pools are Joined: Log from Node with Lowest Node ID is Replicated	520
Figure B-4	Changing Quorum for CXFS Kernel Membership	522
Figure B-5	Example Memberships in a Coexecution Cluster	524
Figure B-6	CXFS Tiebreaker Node	525
Figure F-1	SGI Altix 450 System Control Network	546
Figure F-2	Altix 350 Rear Panel	547
Figure F-3	L2 Rear Panel	547
Figure F-4	IX-brick Rear Panel	548
Figure F-5	Altix 3000 and Origin 3000 Serial Connections	549
Figure F-6	Serial Connection Between SGI Origin 3200 and Origin 3400/3800 Servers	550
Figure F-7	Serial Connection Between Two SGI Origin 3400 or SGI Origin 3800 Servers	551

Figure F-8	Origin 350 Rear Panel	552
Figure F-9	Serial Connection Between Two Origin 200 Deskside Servers	554
Figure F-10	Serial Connection Between Two SGI 2200 Deskside Servers	555
Figure F-11	MMSC Serial Port	556

Tables

Table 1-1	Client-only Commands	57
Table 1-2	Administration Commands	59
Table 2-1	Number of LUNs Supported	65
Table 3-1	Brocade Firmware	68
Table 3-2	QLogic FC Switch Firmware	72
Table 4-1	Examples of License Keys Required for Cluster Changes	81
Table 8-1	<code>fs2d.options</code> File Options	130
Table 8-2	<code>clconfd.options</code> File Options	133
Table 11-1	GUI Platforms	188
Table 11-2	Command Buttons	195
Table 11-3	Key to Icons	201
Table 11-4	Key to States	204
Table 12-1	System Controller Types	278
Table 13-1	CXFS and Cluster Administration Initialization Commands	317
Table 13-2	CXFS <code>chkconfig</code> Arguments	322
Table 19-1	Error Strings	445
Table 19-2	System Log File Error Message Format	467
Table 19-3	Log File Error Message Format	477
Table A-1	CXFS Daemons and Threads	496
Table D-1	IRIX Paths	533
Table D-2	SGI ProPack Paths	534

About This Guide

This publication documents CXFS 4.2 running on a storage area network (SAN). It supports CXFS 4.2. It assumes that you are already familiar with the XFS filesystem and you have access to the *XVM Volume Manager Administrator's Guide*.

You should read through this entire book, especially Chapter 19, "Troubleshooting" on page 425, before attempting to install and configure a CXFS cluster.

Related Publications

The following documents contain additional information:

- *CXFS MultiOS Client-Only Guide for SGI InfiniteStorage*
- *FailSafe Administrator's Guide for SGI InfiniteStorage*
- *SGI InfiniteStorage Cluster Manager for Linux Administrator's Guide*
- *XVM Volume Manager Administrator's Guide*
- Storage area network (SAN) documentation:
 - *EL Serial Port Server Installation Guide* (provided by Digi International)
 - *EL Serial Port Server Installation Guide Errata*
 - *FDDIXPress Administration Guide*
 - *SGI® InfiniteStorage TP9400 and SGI® InfiniteStorage TP9500 and TP9500S RAID User's Guide*
 - *SGI InfiniteStorage TP9300 and TP9300S RAID User's Guide*
 - *SGI Total Performance 9100 Storage System Owner's Guide*
 - *SGI TPSSM Administration Guide*
 - *SGI InfiniteStorage RM610 and RM660 User's Guide*
- *SGI InfiniteStorage 6700 User's Guide*

- IRIX documentation:
 - *IRIX 6.5 Installation Instructions*
 - *IRIX Admin: Disks and Filesystems*
 - *IRIX Admin: Networking and Mail*
 - *Personal System Administration Guide*
 - *Performance Co-Pilot for IRIX Advanced User's and Administrator's Guide*
 - *Performance Co-Pilot Programmer's Guide*
 - *Trusted IRIX Read Me First Notice*
 - *Trusted IRIX/CMW Security Features User's Guide*
- SGI ProPack for Linux, SGI Altix, and SGI Altix XE documentation:
 - The user guide and quick start guide for your SGI Altix or SGI Altix XE system
 - *Guide to Programming Environments and Tools Available on SGI Altix XE System*
 - *NIS Administrator's Guide*
 - *Personal System Administration Guide*
 - *SGI ProPack 5 for Linux Service Pack 3 Start Here*
 - *Performance Co-Pilot for IA-64 Linux User's and Administrator's Guide*
 - *SGI L1 and L2 Controller Software User's Guide*

The following man pages are provided on CXFS server-capable nodes:

Server-Capable Node Man Page	IRIX Subsystem	SGI ProPack RPM
cbeutil(1M)	cluster_admin.man.man	cluster_admin
cdbBackup(1M)	cluster_admin.man.man	cluster_admin
cdbRestore(1M)	cluster_admin.man.man	cluster_admin
cdbconfig(1M)	cluster_admin.man.man	cluster_admin
cdbutil(1M)	cluster_admin.man.man	cluster_admin

Server-Capable Node Man Page	IRIX Subsystem	SGI ProPack RPM
cmond(1M)	cluster_admin.man.man	cluster_admin
fs2d(1M)	cluster_admin.man.man	cluster_admin
cluster_status(1M)	cluster_services.man.man	cluster_services
cluster_services.man.man	cluster_services	
cms_failconf(1M)	cluster_services.man.man	cluster_services
cms_intervene(1M)	cluster_services.man.man	cluster_control
crsd(1M)	cluster_services.man.man	cluster_services
haStatus(1M)	cluster_services.man.man	cluster_services
ha_cilog(1M)	cluster_services.man.man	cluster_services
ha_cmsd(1M)	cluster_services.man.man	cluster_services
ha_exec2(1M)	cluster_services.man.man	cluster_services
ha_gcd(1M)	cluster_services.man.man	cluster_services
ha_ifd(1M)	cluster_services.man.man	cluster_services
ha_ifdadmin(1M)	cluster_services.man.man	cluster_services
ha_macconfig2(1M)	cluster_services.man.man	cluster_services
ha_srmd(1M)	cluster_services.man.man	cluster_services
cxfs_admin(1M)	cxfs_admin.sw.man	cxfs_admin
hafence(1M)	cxfs_cluster.man.man	cxfs_cluster
cxfs-config(1M)	cxfs_util.man.man	cxfs_util
cxfsdp(1)	cxfs_util.man.man	cxfs_util
cxfsdump(1M)	cxfs_util.man.man	cxfs_util
cxfslicense(1M)	cxfs_util.man.man	cxfs_util

About This Guide

Server-Capable Node Man Page	IRIX Subsystem	SGI ProPack RPM
xvm(1M)	eo.e.sw.xvm	N/A
xvm(7M)	eo.e.sw.xvm	N/A
xvm(5)	N/A	cxfs-xvm-cmds
xvm(8)	N/A	cxfs-xvm-cmds
cxfsmgr(1M) [cxdetail(1M) and cxtask(1M) on IRIX] ¹	sysadm_cxfs.man.pages	cxfs-sysadm_cxfs-client
xvmgr(1M)	sysadm_xvm.man.pages	cxfs-sysadm_xvm-client

The following man pages are provided on CXFS client-only nodes:

Client-Only Man Page	IRIX Subsystem	SGI ProPack Subsystem
cxfs_client(1M)	cxfs_client.man.man	cxfs_client
cxfs_info(1M)	cxfs_client.man.man	cxfs_client
cxfs-config(1M)	cxfs_util.man.man	cxfs_util
cxfs_cp(1)	cxfs_util.man.man	cxfs_util
cxfs_dump(1M)	cxfs_util.man.man	cxfs_util

¹ The man page is available from all three command names on IRIX systems for historical purposes.

Obtaining Publications

You can obtain SGI documentation as follows:

- See the SGI Technical Publications Library at <http://docs.sgi.com>. Various formats are available. This library contains the most recent and most comprehensive set of online books, release notes, man pages, and other information.
- If it is installed on your IRIX SGI system, you can use InfoSearch, an online tool that provides a more limited set of online books, release notes, and man pages. On an IRIX system, enter `infosearch` at a command line or select **Help > InfoSearch** from the Toolchest.
- You can view the release notes as follows:
 - On IRIX systems, use either `grelnotes` or `relnotes`
 - On SGI for ProPack Linux systems, see `linux-64/README_CXFS_LINUX64_4.2.0.txt` on the CD
- You can view man pages by typing `man title` at a command line.

Conventions

This guide uses the following terminology abbreviations:

- *Solaris* refers to Solaris 8 or Solaris 9
- *Windows* refers to Microsoft Windows 2000, Microsoft Windows 2003, and Microsoft Windows XP
- *SGI ProPack* refers to SGI ProPack 5 for Linux running the `default` kernel on SGI Altix systems and the `smp` kernel of SGI Altix XE systems.

The following conventions are used throughout this document:

Convention	Meaning
command	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.
user input	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.
GUI element	This bold font denotes the names of graphical user interface (GUI) elements, such as windows, screens, dialog boxes, menus, toolbars, icons, buttons, boxes, and fields.
<TAB>	Represents pressing the specified key in an interactive session

This guide uses *Windows* to refer to both Microsoft Windows 2000 and Microsoft Windows XP nodes when the information applies equally to both. Information that applies to only one of these types of nodes is identified.

Reader Comments

If you have comments about the technical accuracy, content, or organization of this publication, contact SGI. Be sure to include the title and document number of the publication with your comments. (Online, the document number is located in the front matter of the publication. In printed publications, the document number is located at the bottom of each page.)

You can contact SGI in any of the following ways:

- Send e-mail to the following address:

techpubs@sgi.com

- Contact your customer service representative and ask that an incident be filed in the SGI incident tracking system.
- Send mail to the following address:

SGI
Technical Publications
1140 East Arques Avenue
Sunnyvale, CA 94085-4602

SGI values your comments and will respond to them promptly.

Introduction to CXFS



Caution: CXFS is a complex product. To ensure that it is installed and configured in an optimal manner, you **must** purchase initial setup services from SGI. You should read through the following chapters, before attempting to install and configure a CXFS cluster:

- Chapter 1, "Introduction to CXFS" on page 1
- Chapter 2, "SGI RAID for CXFS Clusters" on page 63
- Chapter 3, "Switches" on page 67
- Chapter 4, "CXFS License Keys" on page 75
- Chapter 5, "Preinstallation Steps" on page 93
- Chapter 6, "IRIX CXFS Installation" on page 101 and/or Chapter 7, "SGI ProPack CXFS Installation" on page 113
- Chapter 8, "Postinstallation Steps" on page 127
- Chapter 9, "Best Practices" on page 139
- Chapter 10, "Initial Setup of the Cluster" on page 163

If you are using a multiOS cluster, also see the *CXFS MultiOS Client-Only Guide for SGI InfiniteStorage*. If you are using coexecution with FailSafe, also see the *FailSafe Administrator's Guide for SGI InfiniteStorage*.

This chapter discusses the following:

- "What is CXFS?"
- "Comparison of XFS and CXFS" on page 3
- "Comparison of Network and CXFS Filesystems" on page 7
- "Cluster Environment" on page 9
- "Hardware and Software Support" on page 38
- "Overview of FailSafe Coexecution" on page 41
- "CXFS Tools Overview" on page 42
- "Guaranteed-Rate I/O (GRIO) Version 2 and CXFS" on page 44
- "XVM Failover and CXFS" on page 45
- "GPT Labels and CXFS" on page 54
- "Installation and Configuration Overview" on page 56

SGI ProPack refers to SGI ProPack 5 for Linux running the default kernel on SGI Altix systems or the `smp` kernel of SGI Altix XE systems.

What is CXFS?

CXFS is clustered XFS, a clustered filesystem for high-performance computing environments.

CXFS allows groups of computers to coherently share XFS filesystems among multiple hosts and storage devices while maintaining high performance. CXFS runs on storage area network (SAN) disks, such as Fibre Channel. A SAN is a high-speed, scalable network of servers and storage devices that provides storage resource consolidation, enhanced data access/availability, and centralized storage management. CXFS filesystems are mounted across the cluster by CXFS management software. All files in the filesystem are available to all nodes that mount the filesystem. All shared filesystems must be built on top of XVM volumes.

Comparison of XFS and CXFS

CXFS uses the same filesystem structure as XFS. A CXFS filesystem is initially created using the same `mkfs` command used to create standard XFS filesystems.

The primary difference between XFS and CXFS filesystems is the way in which filesystems are mounted and managed:

- In XFS:
 - Filesystems are mounted with the `mount` command directly by the system during boot via an entry in `/etc/fstab` or by the IRIX Filesystem Manager.
 - A filesystem resides on only one host.
 - The `/etc/fstab` file contains static information about filesystems. For more information, see the `fstab` man page.
- In CXFS:
 - Filesystems are mounted using the CXFS Manager graphical user interface (GUI) or the `cxfs_admin` command.¹
 - A filesystem is accessible to those hosts (nodes) in the cluster that are defined to mount it. CXFS filesystems are mounted across the cluster by CXFS management software. All files in the filesystem are visible to those hosts that are defined to mount the filesystem.
 - One node coordinates the updating of *metadata* (information that describes a file, such as the file's name, size, location, and permissions) on behalf of all nodes in a cluster; this is known as the *metadata server*.

There is one *active metadata server* per CXFS filesystem; there can be multiple active metadata servers in a cluster, one for each CXFS filesystem.

- The filesystem information is stored in the *cluster database (CDB)*, which contains persistent static configuration information about the filesystems, nodes, and cluster. The CXFS cluster daemons manage the distribution of multiple synchronized copies of the cluster database across the *CXFS administration nodes* in the pool. The administrator can view the database and modify it using the GUI or `cxfs_admin`.

¹ The CXFS GUI and `cxfs_admin` are the preferred CXFS configuration tools. For information about `cmgr`, see Appendix G, "Reference to `cmgr` Tasks" on page 557.

The GUI shows the static and dynamic state of the cluster. For example, suppose the database contains the static information that a filesystem is enabled for mount; the GUI will display the dynamic information showing one of the following:

- A blue icon indicating that the filesystem is mounted (the static and dynamic states match).
- A grey icon indicating that the filesystem is configured to be mounted but the procedure cannot complete because CXFS services have not been started (the static and dynamic states do not match, but this is expected under the current circumstances). See "CXFS Services" on page 25.
- An error (red) icon indicating that the filesystem is supposed to be mounted (CXFS services have been started), but it is not (the static and dynamic states do not match, and there is a problem).

The following commands can also be used to view the cluster state:

- `cxfs_admin` shows both the static and dynamic cluster states. This command is available on hosts that have the appropriate access and network connections.
 - `cxfs-config` shows the static cluster state. This commands are available on nodes used for cluster administration.
 - `clconf_info` shows both the static and dynamic cluster states. This command is available on nodes used for cluster administration.
 - `cxfs_info` provides status information. This command is available on nodes that are CXFS clients but are not used for administration.
- Information is **not** stored in the `/etc/fstab` file. (However, the CXFS filesystems do show up in the `/etc/mstab` file.) For CXFS, information is instead stored in the cluster database.

Supported XFS Features

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 *bandwidth* (megabytes per second), high *transaction rates* (I/O per second), and fast metadata operations.
- Dynamically allocated metadata space.
- Quotas. You can administer quotas from any IRIX or Linux node in the cluster, just as if this were a regular XFS filesystem.
- Filesystem reorganizer (defragmenter), which must be run from the CXFS metadata server for a given filesystem. See the `fsr_xfs` man page.
- Restriction of access to files using file permissions and access control lists (ACLs). You can also use logical unit (LUN) masking or physical cabling to deny access from a specific host to a specific set of disks in the SAN.
- Real-time volumes. CXFS can write to real-time files in real-time volumes on IRIX nodes. For more information about real-time volumes, see *XVM Volume Manager Administrator's Guide*.

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/O 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.

When to Use CXFS

You should use CXFS when you have multiple nodes running applications that require high-bandwidth access to common filesystems.

CXFS performs best under the following conditions:

- Data I/O operations are greater than 16 KB
- Large files are being used (a lot of activity on small files will result in slower performance)

- Read/write conditions are one of the following:
 - All processes that perform reads/writes for a given file reside on the same node.
 - The same file is read by processes on multiple nodes using buffered I/O, but there are no processes writing to the file.
 - The same file is read and written by processes on more than one node using direct-access I/O.

For most filesystem loads, the scenarios above represent the bulk of the file accesses. Thus, CXFS delivers 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.

Performance Considerations

CXFS may not give optimal performance under the following circumstances, and extra consideration should be given to using CXFS in these cases:

- When you want to access files only on the local host.
- When distributed applications write to shared files that are memory mapped.
- When exporting a CXFS filesystem via NFS, be aware that performance will be much better when the export is performed from an active CXFS metadata server than when it is performed from a CXFS client. (Exporting from a backup metadata server or client is not supported. In order to support relocation and recovery, a backup server cannot run any applications that will use the filesystem. For more information, see "Node Functions" on page 12.)
- When access would be as slow with CXFS as with network filesystems, such as with the following:
 - Small files
 - Low bandwidth
 - Lots of metadata transfer

Metadata operations can take longer to complete through CXFS 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 with CXFS 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.

Comparison of Network and CXFS Filesystems

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

Network 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 and the bandwidth. These delays are difficult to avoid in the client-server architecture of LAN-based network 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. This can be 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 fileserver can be a bottleneck for performance when multiple clients wait their turns for data, which must pass through the centralized fileserver. For example, 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 host. When many clients access the server simultaneously, the server's responsiveness degrades.

Note: You should not use multiple Samba servers to export the same CXFS filesystem. For more information, see "Samba" on page 343.

CXFS Filesystems

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

Features

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 hosts in the cluster. Each host can read and write the disks at near-local disk speeds; the data passes directly from the disks to the host requesting the I/O, without passing through a data server or over a local area network (LAN). For the data path, each host 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 systems, the bandwidth can scale to gigabytes and even tens of gigabytes per second. Compare this with a network filesystem with the data typically flowing over a 1- to 100-MB-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 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 filesystem and is not dependent on any particular host. The pathname is a normal POSIX pathname; for example, `/u/username/directory`.

Note: A Windows CXFS client uses the same pathname to the filesystem as other clients beneath a preconfigured drive letter.

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 host and processes distributed across multiple hosts have the same view of the filesystem, with performance similar on each host.

This differs from typical network filesystems, which tend to include the name of the fileserver in the pathname. 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 UNIX 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.

Restrictions

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.
- Hierarchical storage management (HSM) applications must run on the metadata server.
- The inode monitor device (`imon`) is not supported on CXFS filesystems.

The following XFS features are not supported in CXFS:

- The original XFS guaranteed-rate I/O (GRIO) implementation, GRIO version 1. (GRIO version 2 is supported, see "Guaranteed-Rate I/O (GRIO) Version 2 and CXFS" on page 44).
- Swap to a file residing on a CXFS file system.

Cluster Environment

This section discusses the following:

- "Terminology" on page 10
- "Isolating Failed Nodes: Failure Policies" on page 28
- "The Cluster Database and CXFS Clients" on page 36
- "Metadata Server Functions" on page 36

- "System View" on page 37
- "CXFS and Highly Available Services" on page 38

For details about CXFS daemons, communication paths, and the flow of metadata, see Appendix A, "CXFS Software Architecture" on page 495.

Terminology

This section defines the terminology necessary to understand CXFS. Also see the Glossary on page 683.

Cluster

A *cluster* is the set of systems (nodes) configured to work together as a single computing resource. A cluster is identified by a simple name and a cluster ID. A cluster running multiple operating systems is known as a *multiOS cluster*.

A given node may not be a member of multiple clusters.

Disks or logical units (LUNs) are assigned to a cluster by recording the name of the cluster on the disk (or LUN). Thus, if any disk is accessible (via a Fibre Channel connection) from nodes in different clusters, then those clusters must have unique names. When members of a cluster send messages to each other, they identify their cluster via the cluster ID. Cluster names and IDs must be unique.

Because of the above restrictions on cluster names and cluster IDs, and because cluster names and cluster IDs cannot be changed once the cluster is created (without deleting the cluster and recreating it), SGI advises that you choose unique names and cluster IDs for each of the clusters within your organization.

Node

A *node* is an operating system (OS) image, usually an individual computer. (This use of the term *node* does not have the same meaning as a node in an SGI Origin 3000 or SGI 2000 system.)

A given node can be a member of only one cluster.

Pool

In the GUI, the *pool* is the set of nodes from which a particular cluster may be formed. All nodes created in `cxfs_admin` are automatically part of the cluster, so the concept of the pool is obsolete when using `cxfs_admin`.

Only one cluster may be configured from a given pool, and it need not contain all of the available nodes. (Other pools may exist, but each is disjoint from the other. They share no node or cluster definitions.)

A pool is first formed when you connect to a given CXFS administration node (one that is installed with `cluster_admin`) and define that node in the cluster database using the CXFS GUI. You can then add other nodes to the pool by defining them while still connected to the first node. (If you were to connect to a different node and then define it, you would be creating a second pool).

Figure 1-1 shows the concepts of pool and cluster.

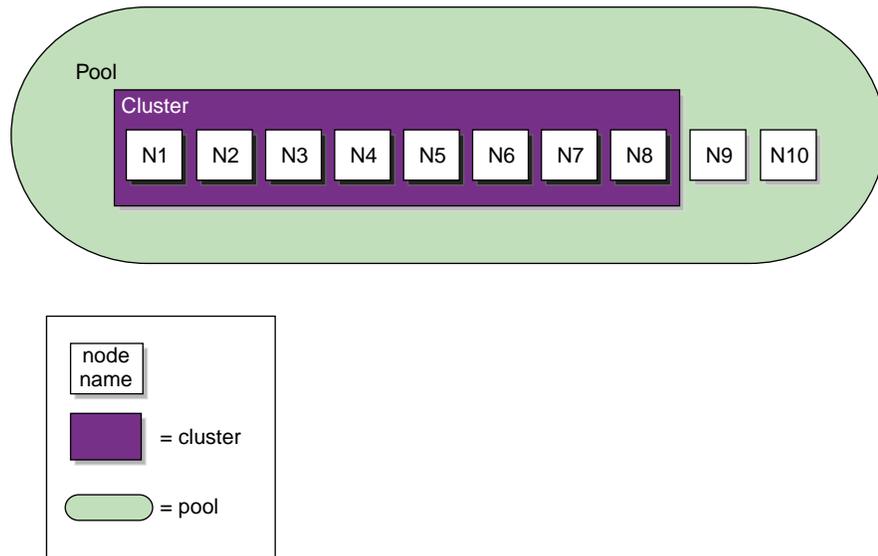


Figure 1-1 Pool and Cluster Concepts

Cluster Database

The *cluster database* contains configuration information about nodes, the cluster, logging information, and configuration parameters. The *cluster administration daemons* manage the distribution of the cluster database (CDB) across the CXFS administration nodes in the pool. See "Cluster Administration Daemons" on page 24.

The database consists of a collection of files; you can view and modify the contents of the database by using the following:

- `cxfs_admin` on a CXFS server-capable administration node or another host that has been given the proper permissions
- CXFS Manager GUI connected to a CXFS administration node
- `clconf_info` and `cxfs-config` commands on a CXFS administration node
- `cxfs_info` command on a client-only nodes

Node Functions

A node can have one of the following functions:

- *Server-capable administration node* (IRIX or SGI ProPack).

This node is installed with the `cluster_admin` software product, which contains the full set of cluster administration daemons (`fs2d`, `crsd`, `cad`, and `cmond`) and the CXFS control daemon (`clconfd`). For more details about daemons, see "Cluster Administration Daemons" on page 24, "CXFS Control Daemon" on page 26, and Appendix A, "CXFS Software Architecture" on page 495.

This node type is capable of coordinating cluster activity and metadata. *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 megabytes or gigabytes in size.

For each CXFS filesystem, one node is responsible for updating that filesystem's metadata. This node is referred to as the *metadata server*. Only nodes defined as server-capable nodes are eligible to be metadata servers.

Multiple CXFS administration nodes can be defined as *potential metadata servers* for a given CXFS filesystem, but only one node per filesystem is chosen to be the *active metadata server*. All of the potential metadata servers for a given cluster must

be either all IRIX or all SGI ProPack. There can be multiple active metadata servers in the cluster, one per CXFS filesystem.

Other nodes that mount a CXFS filesystem are referred to as *CXFS clients*. A CXFS administration node can function as either a metadata server or CXFS client, depending upon how it is configured and whether it is chosen to be the active metadata server.

Note: Do not confuse *metadata server* and *CXFS client* with the traditional data-path client/server model used by network filesystems. Only the metadata information passes through the metadata server via the private Ethernet network; the data is passed directly to and from disk on the CXFS client via the Fibre Channel connection.

You perform cluster administration tasks by using the `cxfs_admin` on any host with the appropriate access and network connections or by using the CXFS Manager GUI and connecting it to a CXFS server-capable administration node. For more details, see:

- Chapter 11, "Reference to GUI Tasks" on page 183
- Chapter 12, "Reference to `cxfs_admin` Tasks" on page 255

There should be an odd number of server-capable administration nodes with CXFS services running for quorum calculation purposes. If you have a cluster with more than two nodes, define a CXFS tiebreaker node (see "CXFS Tiebreaker" on page 23).

- *Client-only node* (all supported CXFS operating systems).

This node is one that has a minimal implementation of CXFS that runs a single daemon, the CXFS client daemon (`cxfs_client`). For more details, see Appendix A, "CXFS Software Architecture" on page 495.

This node can safely mount CXFS filesystems but it cannot become a CXFS metadata server or perform cluster administration. Client-only nodes retrieve the information necessary for their tasks by communicating with an administration node. This node does not contain a copy of the cluster database.

IRIX and SGI ProPack nodes are client-only nodes if they are installed with the `cxfs_client` software package and defined as client-only nodes. Nodes that are running supported operating systems other than IRIX or SGI ProPack are always configured as CXFS client-only nodes.

For more information, see *CXFS MultiOS Client-Only Guide for SGI InfiniteStorage*.

- *Client administration node* (IRIX only in coexecution with FailSafe).

This is a node that is installed with the `cluster_admin` software product but it cannot be a metadata server. This node type should only be used when necessary for coexecution with FailSafe. It is not supported for the SGI ProPack platform.

Figure 1-2 shows nodes in a pool that are installed with `cluster_admin` and others that are installed with `cxfs_client`. Only those nodes with `cluster_admin` have the `fs2d` daemon and therefore a copy of the cluster database.

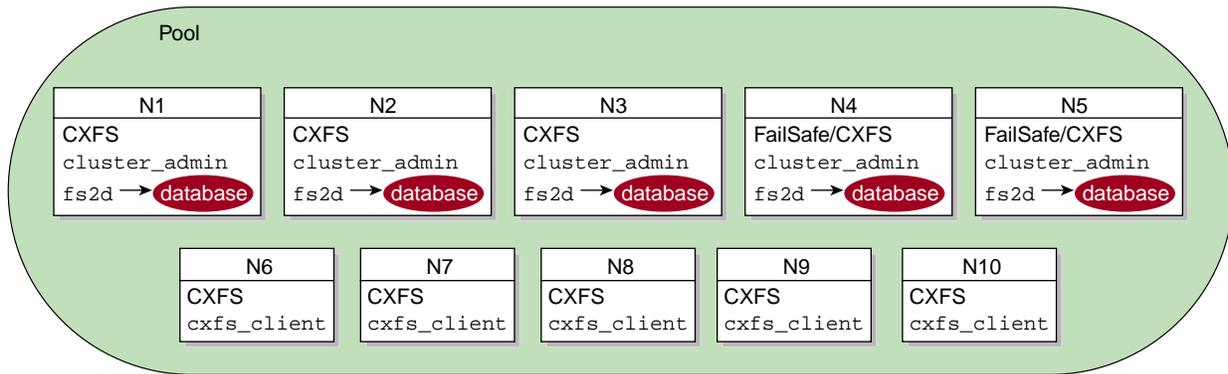


Figure 1-2 Installation Differences

A *standby node* is a server-capable administration node that is configured as a potential metadata server for a given filesystem, but does not currently run any applications that will use that filesystem. (The node can run applications that use other filesystems.)

Ideally, all administration nodes will run the same version of the operating system. However, as of IRIX 6.5.18f, SGI supports a policy for CXFS that permits a rolling annual upgrade; see "CXFS Release Versions and Rolling Upgrades" on page 307.

The following figures show different possibilities for metadata server and client configurations. The potential metadata servers are required to be CXFS administration nodes and must all run IRIX or all run SGI ProPack; the other nodes could be client-only nodes.

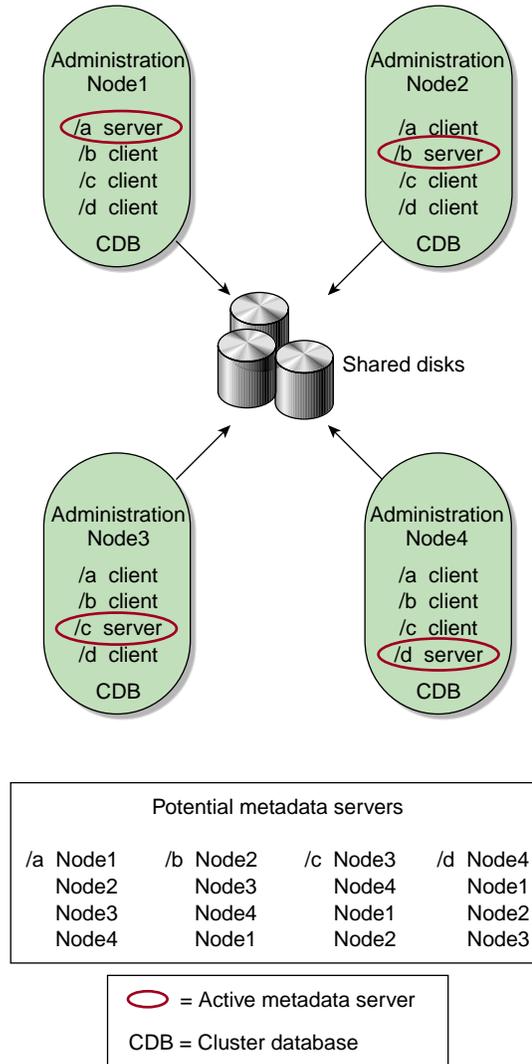


Figure 1-3 Evenly Distributed Metadata Servers

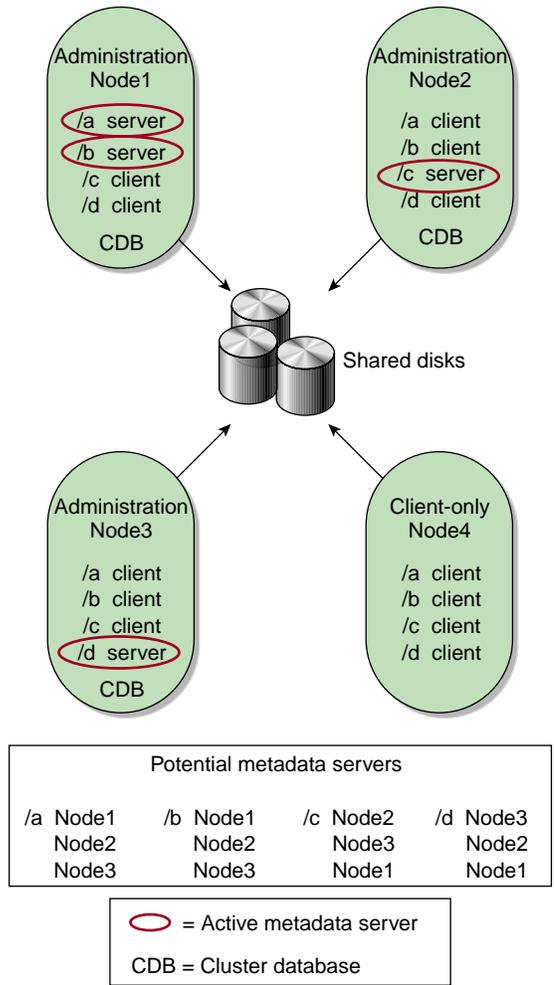


Figure 1-4 Multiple Metadata Servers

In Figure 1-4, Node4 could be running any supported OS because it is a client-only node; it is not a potential metadata server.

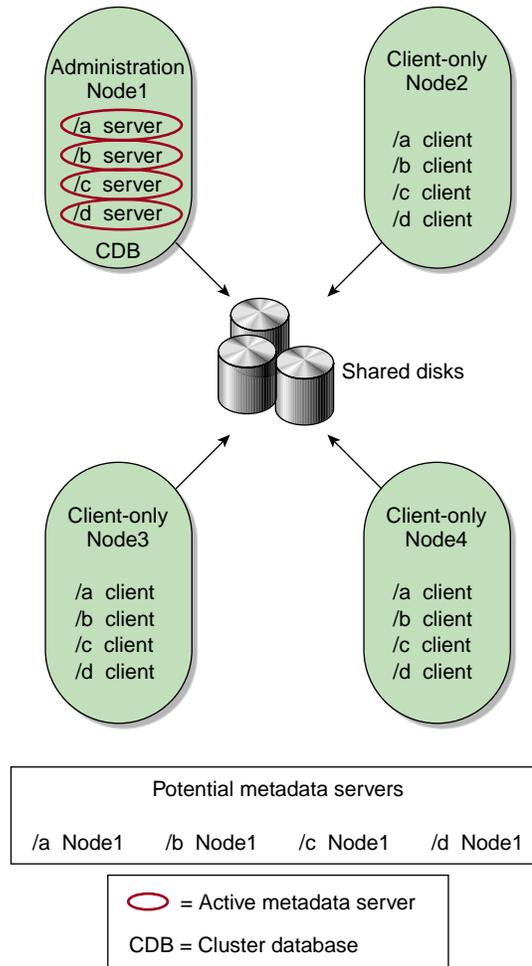


Figure 1-5 One Metadata Server

In Figure 1-5, Node2, Node3, and Node4 could be running any supported OS because they are client-only nodes; they are not potential metadata servers.

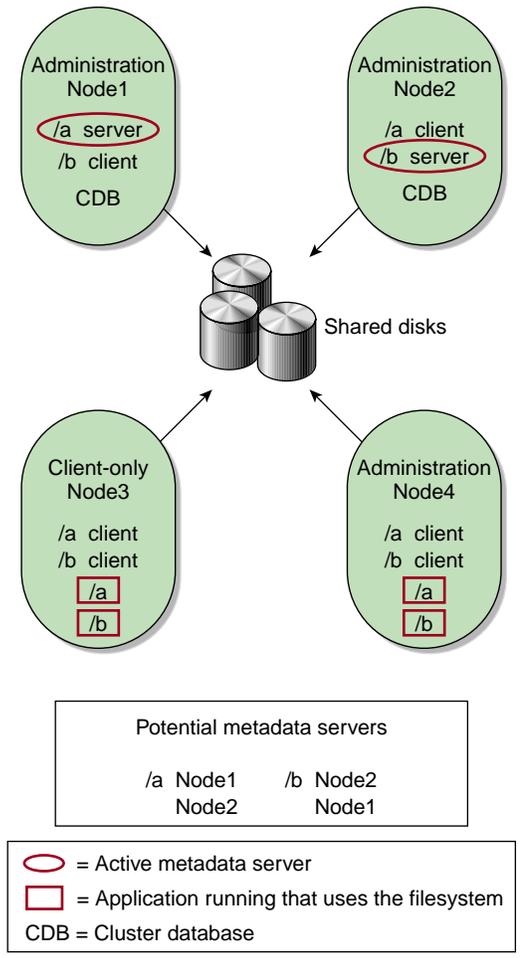


Figure 1-6 Standby Mode

Figure 1-6 shows a configuration in which Node1 and Node2 are potential metadata servers for filesystems /a and /b:

- Node1 is the active metadata server for /a
- Node2 is the active metadata server for /b

Because standby mode is used, neither Node1 nor Node2 runs applications that use /a or /b. The figure shows one client-only node, but there could be several.

Membership

The nodes in a cluster must act together to provide a service. To act in a coordinated fashion, each node must know about all the other nodes currently active and providing the service. The set of nodes that are currently working together to provide a service is called a *membership*:

- *Cluster database membership* (also known as *fs2d membership* or *user-space membership*) is the group of administration nodes that are accessible to each other. (client-only nodes are not eligible for cluster database membership.) The nodes that are part of the the cluster database membership work together to coordinate configuration changes to the cluster database.
- *CXFS kernel membership* is the group of CXFS nodes in the **cluster** that can actively share filesystems, as determined by the the CXFS kernel, which manages membership and heartbeating. The CXFS kernel membership may be a subset of the nodes defined in a cluster. All nodes in the cluster are eligible for CXFS kernel membership.

Heartbeat messages for each membership type are exchanged via a private network so that each node can verify each membership.

A cluster that is also running FailSafe has a *FailSafe membership*, which is the group of nodes that provide highly available (HA) resources for the cluster. For more information, see Appendix B, "Memberships and Quorums" on page 509, and the *FailSafe Administrator's Guide for SGI InfiniteStorage*.

Private Network

A *private network* is one that is **dedicated** to cluster communication and is accessible by administrators but not by users.

Note: A virtual local area network (VLAN) is not supported for a private network.

CXFS uses the private network for metadata traffic. The cluster software uses the private network to send the heartbeat/control messages necessary for the cluster configuration to function. Even small variations in heartbeat timing can cause problems. If there are delays in receiving heartbeat messages, the cluster software

may determine that a node is not responding and therefore revoke its CXFS kernel membership; this causes it to either be reset or disconnected, depending upon the configuration.

Rebooting network equipment can cause the nodes in a cluster to lose communication and may result in the loss of CXFS kernel membership and/or cluster database membership ; the cluster will move into a degraded state or shut down if communication between nodes is lost. Using a private network limits the traffic on the network and therefore will help avoid unnecessary resets or disconnects. Also, a network with restricted access is safer than one with user access because the messaging protocol does not prevent *snooping* (illicit viewing) or *spoofing* (in which one machine on the network masquerades as another).

Therefore, because the performance and security characteristics of a public network could cause problems in the cluster and because heartbeat is very timing-dependent, **a private network is required**. The private network should be used for metadata traffic only.

The heartbeat and control network must be connected to all nodes, and all nodes must be configured to use the same subnet for that network.



Caution: If there are any network issues on the private network, fix them before trying to use CXFS. A stable private network is important for a stable CXFS cluster network.

For more information about network segments and partitioning, see Appendix B, "Memberships and Quorums" on page 509. For information about failover from the private network to another network, see information about the `failover_net` command in `cxfs_admin` (see "Network Failover Tasks with `cxfs_admin`" on page 295). (Although the primary network must be private, the backup network may be public.) For information about using IP filtering for the private network, see Appendix C, "IP Filtering for the CXFS Private Network" on page 529.

Relocation

Relocation is the process by which the metadata server moves from one node to another due to an administrative action; other services on the first node are not interrupted.

Note: Relocation is supported only to standby nodes. Relocation is disabled by default.

A *standby node* is a metadata server-capable administration node that is configured as a potential metadata server for a given filesystem, but does not currently run any applications that will use that filesystem. To use relocation, you must not run any applications on any of the potential metadata servers for a given filesystem; after the active metadata server has been chosen by the system, you can then run applications that use the filesystem on the active metadata server and client-only nodes.

To use relocation to a standby node, you must enable relocation on the active metadata server (relocation is disabled by default.) To enable relocation, reset the `cxfs_relocation_ok` parameter as follows:

- IRIX:

- Enable:

```
irix# systune cxfs_relocation_ok 1
```

- Disable:

```
irix# systune cxfs_relocation_ok 0
```

- SGI ProPack:

- Enable at run time:

```
[root@linux64 root]# sysctl -w fs.cxfs.cxfs_relocation_ok=1
```

- Enable at reboot by adding the following line to `/etc/modprobe.conf` or `/etc/modprobe.conf.local`:

```
options sgi-cxfs cxfs_relocation_ok=1
```

- Disable:

```
[root@linux64 root]# sysctl -w fs.cxfs.cxfs_relocation_ok=0
```

- Disable at reboot by adding the following line to `/etc/modprobe.conf` or `/etc/modprobe.conf.local`:

```
options sgi-cxfs cxfs_relocation_ok=0
```

CXFS kernel membership is not affected by relocation. However, users may experience a degradation in filesystem performance while the metadata server is relocating.

The following are examples of relocation triggers:

- The system administrator uses the GUI or `cxfs_admin` to relocate the metadata server.
- The FailSafe CXFS resource relocates the IRIX CXFS metadata server. The SGI Cluster Manager for Linux CXFS plug-in relocates the SGI ProPack metadata server.
- The system administrator unmounts the CXFS filesystem on an IRIX metadata server. (Unmounting on an SGI ProPack metadata server does not trigger relocation; the SGI ProPack server will just return an `EBUSY` flag.)

Recovery

Recovery is the process by which the metadata server moves from one node to another due to an interruption in services on the first node.

Note: Recovery is supported only to standby nodes.

To use recovery to a standby node, you must not run any applications on any of the potential metadata servers for a given filesystem; after the active metadata server has been chosen by the system, you can then run applications that use the filesystem on the active metadata server and client-only nodes.

The following are examples of recovery triggers:

- A metadata server panic
- A metadata server locks up, causing heartbeat timeouts on metadata clients
- A metadata server loses connection to the heartbeat network

Figure 1-7 describes the difference between relocation and recovery for a metadata server. (Remember that there is one active metadata server per CXFS filesystem. There can be multiple active metadata servers within a cluster, one for each CXFS filesystem.)

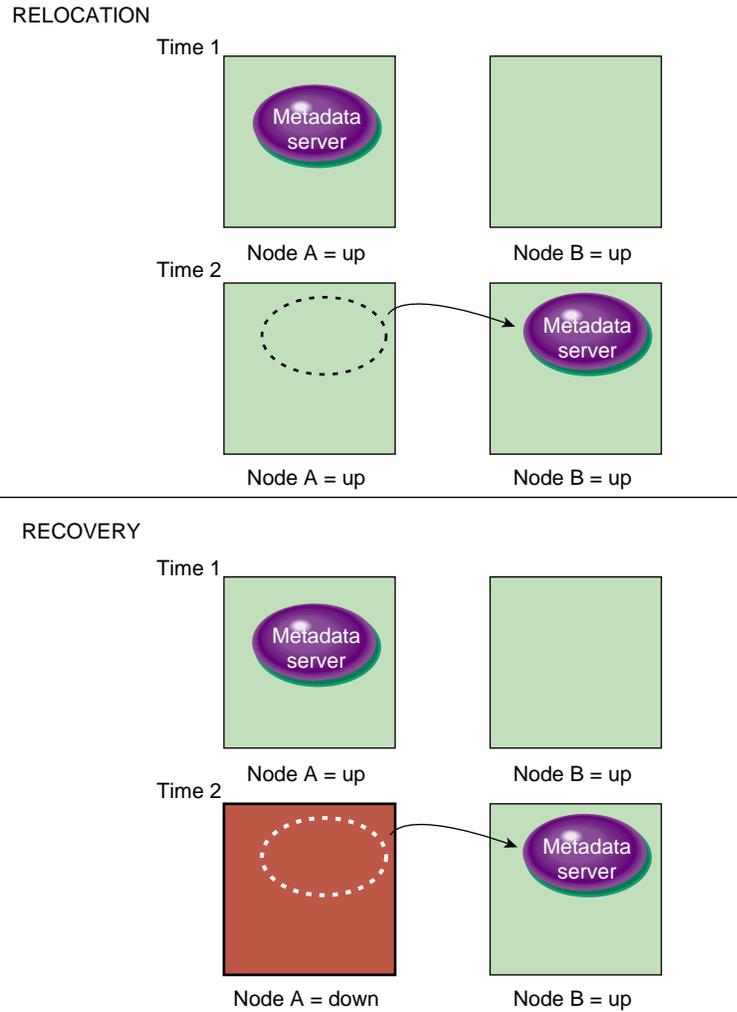


Figure 1-7 Relocation versus Recovery

CXFS Tiebreaker

The *CXFS tiebreaker node* is used in the process of computing the CXFS kernel membership for the cluster when exactly half the server-capable administration nodes in the cluster are up and can communicate with each other.

The tiebreaker is required for all clusters with more than one server-capable node and at least one client-only node. You should choose a reliable client-only node as the tiebreaker; there is no default. For a cluster that consists of **only** four or more server-capable nodes, you should choose one of them as the tiebreaker; this is the only situation in which you should choose a server-capable node as a tiebreaker.

The tiebreaker is required in addition to *I/O fencing* or *system reset*; see "Isolating Failed Nodes: Failure Policies" on page 28.

The CXFS tiebreaker differs from the FailSafe tiebreaker; see *FailSafe Administrator's Guide for SGI InfiniteStorage*.

Cluster Administration Daemons

The following set of daemons, which control various cluster infrastructure needs:

Daemon	Description
fs2d	Manages the cluster database (CDB) on the local administration node and keeps the copy on all administration nodes synchronized.
cad	Provides administration status services to the CXFS GUI.
cmnd	Manages all other cluster administration daemons and the CXFS control daemon (clconfd). The cmnd daemon starts the other daemons on the node and restarts them on failure.
crsd	Monitors the connection to other nodes. Has the ability to reset other nodes.

You can start and stop the cluster administration daemons with the following commands:

- IRIX:

```
/etc/init.d/cluster {start|stop}
```

- SGI ProPack:

```
/etc/init.d/cxfs_cluster {start|stop}
```

Note: You could also use the `restart` option to stop and start.

You can also use the following `chkconfig` commands to specify that the daemons will be restarted upon reboot:

- IRIX:

```
chkconfig cluster on
```

- SGI ProPack:

```
chkconfig cxfs_cluster on
```

For more information, see Appendix A, "CXFS Software Architecture" on page 495.

CXFS Services

The enabling/disabling of a node, which changes a flag in the cluster database. Starting or stopping CXFS services does not affect the daemons involved. The daemons that control CXFS services are as follows:

- `clconfd` on administration nodes, see "CXFS Control Daemon" on page 26.
- `cxfs_client` on client-only nodes, see "CXFS Client Daemon" on page 26.

To *start CXFS services* means to enable a node, which changes a flag in the cluster database by performing an administrative task using the CXFS GUI or `cxfs_admin`:

- "Start CXFS Services with the GUI" on page 227
- "Enable a Node with `cxfs_admin`" on page 279

To *stop CXFS services* means to disable a node, which changes a flag in the cluster database, by performing an administrative task using the GUI or `cxfs_admin`:

- "Stop CXFS Services with the GUI" on page 228
- "Disable a Node with `cxfs_admin`" on page 279

To *shutdown CXFS services* means to withdraw a node from the CXFS kernel membership, either due to the fact that the node has failed somehow or by issuing an `admin cxfs_stop` command. The node remains enabled in the cluster database. See "Forced CXFS Shutdown: Revoke Membership of Local Node" on page 337.

CXFS Control Daemon

The `clconfd` daemon, which controls CXFS services on an administration node. It does the following:

- Obtains the cluster configuration from the `fs2d` daemon and manages the local CXFS administration node's CXFS kernel membership services and filesystems accordingly
- Obtains membership and filesystem status from the kernel
- Issues reset commands to the `crsd` daemon
- Issues I/O fencing commands to configured Fibre Channel switches

You can start/stop `clconfd` with the following command on an IRIX or SGI ProPack administration node:

```
/etc/init.d/cxfs {start/stop}
```

The `clconfd` daemon may still be running when CXFS services are disabled.

You can also use the following `chkconfig` command to specify that `clconfd` will be restarted upon reboot:

- IRIX:

```
chkconfig cxfs_cluster on
```
- SGI ProPack:

```
chkconfig cxfs on
```

For more information, see Appendix A, "CXFS Software Architecture" on page 495.

CXFS Client Daemon

The `cxfs_client` daemon, which controls CXFS services on a client-only node. It does the following:

- Obtains the cluster configuration from a remote `fs2d` daemon and manages the local client-only node's CXFS kernel membership services and filesystems accordingly.
- Obtains membership and filesystem status from the kernel.

You can start/stop `cxfs_client` with the following command on a client-only IRIX or SGI ProPack node:

```
/etc/init.d/cxfs_client {start|stop}
```

Note: The path to the `cxfs_client` command varies among the other platforms supported. See the *CXFS MultiOS Client-Only Guide for SGI InfiniteStorage*.

You can also use the following `chkconfig` command to specify that `cxfs_client` will be restarted upon reboot on either IRIX or SGI ProPack:

```
chkconfig cxfs_client on
```

The `cxfs_client` daemon may still be running when CXFS services are disabled.

For more information, see Appendix A, "CXFS Software Architecture" on page 495.

Forced CXFS Shutdown

Withdraws a node from cluster membership, which disables filesystem and cluster volume access for the node. This is either due to the fact that the node has failed somehow or by issuing an `admin cxfs_stop` command. The node remains enabled in the cluster database. See "Forced CXFS Shutdown: Revoke Membership of Local Node" on page 337.

Heartbeat Monitoring

All nodes send heartbeat messages once per second. If a node does not receive a heartbeat within a defined period, that node loses membership and is denied access to the cluster's filesystems. The defined period is one of the following:

- `static`: Monitors constantly at 1-second intervals and declares a timeout after 5 consecutive missed seconds (default).
- `dynamic`: Starts monitoring only when the node is processing a message from another node (such as for token recall or XVM multicast) or when the client monitors the server because it has a message pending (for example, a token acquire or metadata operation). Once monitoring initiates, it monitors at 1-second intervals and declares a timeout after 5 consecutive missed seconds, just like static monitoring. Dynamic heartbeat monitoring is appropriate for clusters that have clients with heavy workloads; using it avoids inappropriate loss of membership.

However, it may take longer to recover a client's tokens and other state information when there is an actual problem.

You can set the heartbeat monitor for the entire cluster by using the `cxfs_admin` command. See "Create or Modify a Cluster with `cxfs_admin`" on page 282. If you use dynamic heartbeat monitoring, you must not use the `Shutdown` fail policy for client-only nodes; it can be slower to recover because failure detection may take longer if no operations are pending against a node that fails. `Shutdown` is not allowed as a fail policy because of the dynamic nature and potentially asymmetric heartbeat monitor between two nodes. For example, the server may begin monitoring heartbeat for a client, but that client may not currently be monitoring heartbeat of the server, and therefore the nodes may not discover they have lost membership in a timely manner.

Isolating Failed Nodes: Failure Policies

A failed node must be isolated from the rest of the cluster so that it cannot corrupt data in the shared CXFS filesystem. CXFS uses the following methods to isolate failed nodes. You can specify up to three methods by defining the `failpolicy` in the `cxfs_admin` command or the `failure action` in the GUI. The second method will be completed only if the first method fails; the third method will be completed only if both the first and second methods fail. The possible methods are:

- `Fence`, which disables a node's Fibre Channel ports so that it cannot access I/O devices, and therefore cannot corrupt data in the shared CXFS filesystem. When fencing is applied, the rest of the cluster can begin immediate recovery.
- `Reset`, which performs a system reset via a system controller.
- `FenceReset`, which fences the node and then, if the node is successfully fenced, performs an asynchronous system reset; recovery begins without waiting for reset acknowledgment. If used, this fail policy method should be specified first. If the fencing action fails, the reset is not performed; therefore, `reset` alone is also required for all server-capable nodes (unless there is a single server-capable node in the cluster).
- `Shutdown`, which tells the other nodes in the cluster to wait for a period of time (long enough for the node to shut itself down) before reforming the CXFS kernel membership. (However, there is no notification that the node's shutdown has actually taken place.)



Caution: Because there is no notification that a shutdown has occurred, if you have a cluster with no tiebreaker, you must not use the `shutdown` setting for any server-capable node in order to avoid multiple clusters being formed. See "Shutdown" on page 35.

You must not use the Shutdown failure policy on client nodes if you choose dynamic monitoring.

The following are valid failure policy sets:

Note: If the failure hierarchy contains `reset` or `fencerreset`, the reset might be performed before the system kernel core-dump can complete, resulting in an incomplete core-dump.

- Server-capable nodes:

FenceReset, Reset (Preferred)
FenceReset
Reset
Reset, Fence
(none)

- Client-only nodes with static heartbeat monitoring:

Fence, Shutdown (Preferred)
Fence
Fence, Reset
Fence, Reset, Shutdown
FenceReset
FenceReset, Reset
FenceReset, Reset, Shutdown
FenceReset, Shutdown
Reset
Reset, Fence
Reset, Fence, Shutdown
Reset, Shutdown
Shutdown
(none)

- Client-only nodes with dynamic heartbeat monitoring:

Fence (Most common)
Fence, Reset
FenceReset
FenceReset, Reset
Reset
Reset, Fence
(none) (Not Preferred)

For information about heartbeat monitoring, see "Heartbeat Monitoring" on page 27.

Note: If you choose no method, or if the fail policy does not include Shutdown and all of the other actions fail, CXFS will stall membership until the failed node either attempts to join the cluster again or until the administrator intervenes by using `cms_intervene`. Objects held by the failed node stall until membership finally transitions and initiates recovery. For more information, see the `cms_intervene(1M)` man page.

The rest of this section provides more details. See also "Protect Data Integrity on All Nodes" on page 146. For more information about setting the policies, see:

- "Define a Node with the GUI" on page 208
- "Create or Modify a Node with `cxfs_admin`" on page 270

Fence

I/O fencing does the following:

- Preserves data integrity by preventing I/O from nodes that have been expelled from the cluster
- Speeds the recovery of the surviving cluster, which can continue immediately rather than waiting for an expelled node to reset under some circumstances

To support I/O fencing, platforms require a Fibre Channel switch; for supported switches, see the release notes. You must put switches used for I/O fencing on a network other than the primary CXFS private network so that problems on the CXFS private network can be dealt with by the fencing process and thereby avoid data corruption issues. The network to which the switch is connected must be accessible by all administration nodes in the cluster.

Note: I/O fencing differs from zoning. *Fencing* is a generic cluster term that means to erect a barrier between a host and shared cluster resources. *Zoning* is the ability to define logical subsets of the switch (zones), with the ability to include or exclude hosts and media from a given zone. A host can access only media that are included in its zone. Zoning is one possible implementation of fencing.

Zoning implementation is complex and does not have uniform availability across switches. Therefore, SGI chose to implement a simpler form of fencing: enabling/disabling a host's Fibre Channel ports.

When a node joins the CXFS kernel membership, the worldwide port name (WWPN) of its host bus adapter (HBA) is stored in the cluster database. If there are problems with the node, the I/O fencing software sends a message via the `telnet` protocol to the appropriate switch and disables the port.



Caution: You must keep the `telnet` port free in order for I/O fencing to succeed.

Brocade switches running 4.x.x.x or later firmware by default permit multiple `telnet` sessions. However, in the case of a network partition, a server-capable administration node from each side of the partition will attempt to issue the fence commands, but only the node that is able to log in will succeed. Therefore, on a Brocade switch running 4.x.x.x or later firmware, you must modify the `admin` account to restrict it to a single `telnet` session. For details, see the release notes.

The switch then blocks the problem node from communicating with the storage area network (SAN) resources via the corresponding HBA. Figure 1-8 on page 33, describes this.

If users require access to nonclustered LUNs or devices in the SAN, these LUNs/devices must be accessed or mounted via an HBA that has been explicitly masked from fencing. For details on how to exclude HBAs from fencing for nodes, see:

- "Define a Switch with the GUI" on page 233
- "Create a Switch with `cxfs_admin`" on page 296

For nodes running other supported operating systems, see *CXFS MultiOS Client-Only Guide for SGI InfiniteStorage*.

To recover, the affected node withdraws from the CXFS kernel membership, unmounts all file systems that are using an I/O path via fenced HBA(s), and then rejoins the cluster. This process is called *fencing recovery* and is initiated automatically. Depending on the failure action hierarchy that has been configured, a node may be reset (rebooted) before initiating fencing recovery. For information about setting the failure action hierarchy, see "Create a Switch with `cxfs_admin`" on page 296 and "Define a Node with the GUI" on page 208.

In order for a fenced node to rejoin the CXFS kernel membership, the current cluster leader must lower its fence to allow it to reprobe its XVM volumes and then remount its filesystems. If a node fails to rejoin the CXFS kernel membership, it may remain fenced. This is independent of whether the node was rebooted, because fencing is an operation applied on the switch, not the affected node. In certain cases, it may therefore be necessary to manually lower a fence. For instructions, see "Lower the I/O Fence for a Node with the GUI" on page 237, and "Using `hafence` to Manipulate a Switch" on page 318.



Caution: When a fence is raised on an HBA, **no further I/O is possible to the SAN via that HBA until the fence is lowered.** This includes the following:

- I/O that is queued in the kernel driver, on which user processes and applications may be blocked waiting for completion. These processes will return the `EIO` error code under UNIX, or display a warning dialog that I/O could not be completed under Windows.
 - I/O issued via the affected HBAs to nonclustered (local) logical units (LUNs) in the SAN or to other Fibre Channel devices such as tape storage devices.
-

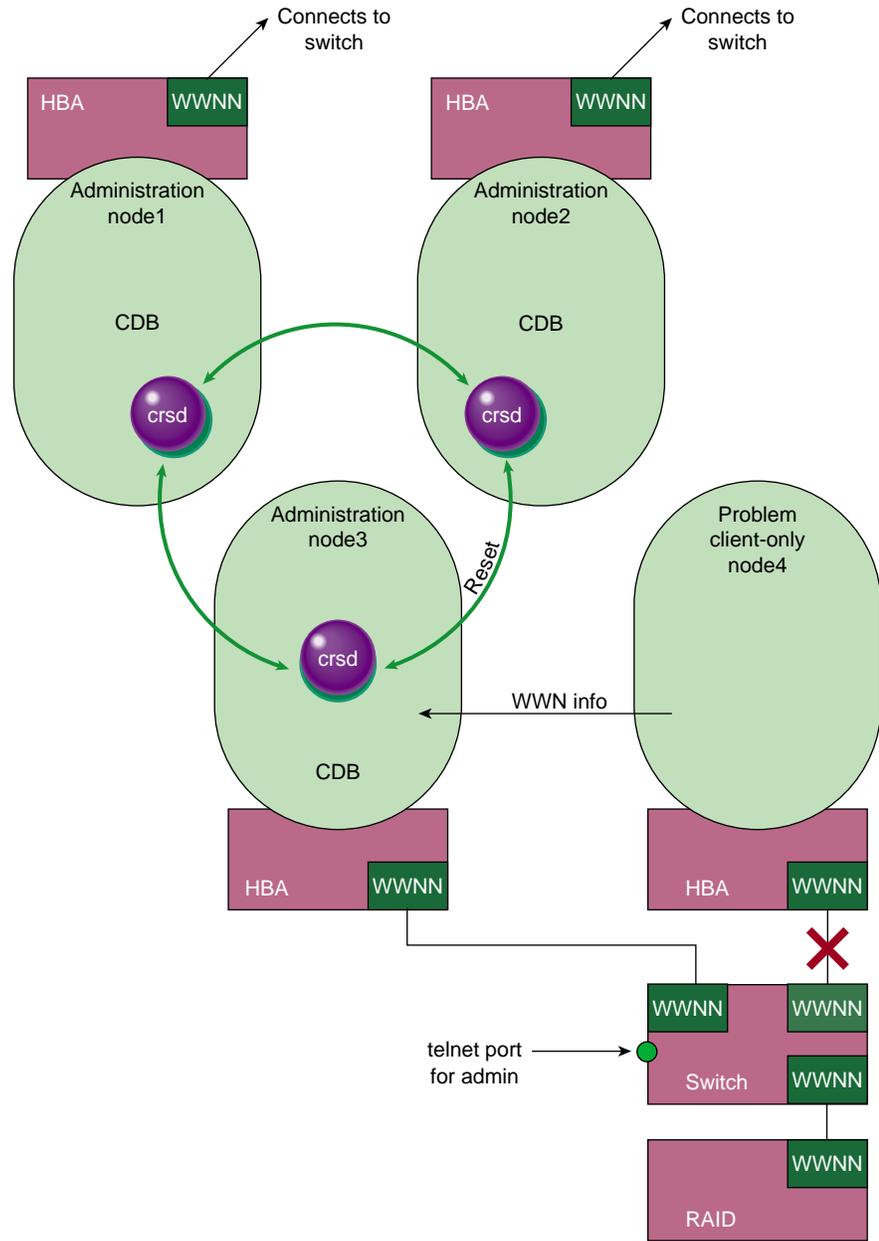


Figure 1-8 I/O Fencing

On client-only nodes with system reset capability, you would want to use Fence for data integrity protection when CXFS is just a part of what the node is doing and therefore losing access to CXFS is preferable to having the system rebooted. An example of this would be a large compute server that is also a CXFS client. However, Fence cannot return a nonresponsive node to the cluster; this problem will require intervention from the system administrator.

For more information, see "Switches and I/O Fencing Tasks with the GUI" on page 233 , "Create or Modify a Node with `cxfs_admin`" on page 270, and "Switch Tasks with `cxfs_admin`" on page 295.

Note: You cannot use I/O fencing for FailSafe nodes. FailSafe nodes require the system reset capability.

Reset

System reset is recommended because if a server hangs, it must be rebooted as quickly as possible to get it back in service, which is not available with I/O fencing. In addition, data corruption is more likely to occur with a rogue metadata server, not a rogue client. (If fencing were to be used on a metadata server and fail, the cluster would have to either shutdown or hang. A fencing failure can occur if an administrator is logged into the switch.)

System reset may be either serial reset or, for a system with an L2 system controller or a baseboard management controller (BMC), over the network.

The system reset can use the following methods:

- `powerCycle` shuts off power to the node and then restarts it
- `reset` simulates the pressing of the reset button on the front of the machine
- NMI (nonmaskable interrupt) performs a core-dump of the operating system kernel, which may be useful when debugging a faulty machine

Note: NMI is not available on systems containing a BMC.

On IRIX, the system reset connection has the same connection configuration as FailSafe; for more information, contact SGI professional or managed services. Also see Appendix F, "System Reset Configuration" on page 539.

You would want to use `Reset` for I/O protection on a client-only node that has a system controller when CXFS is a primary activity and you want to get it back online fast; for example, a CXFS fileserver.

Shutdown

You should only use the `shutdown` failpolicy for client-only nodes that use static heartbeat monitoring.

In the case of a cluster with no tiebreaker node, it is possible that using the `shutdown` setting could cause a split-brain scenario in which multiple clusters could be formed and data could therefore be corrupted.

If the CXFS private network between `NodeA` and `NodeB` fails, the following could occur:

1. Each node will try to fence the other. (That is, `NodeA` will try to fence `NodeB`, and `NodeB` will try to fence `NodeA`).
2. If the fence fails, each node will try to reset the other.
3. If the system reset fails, each assumes that the other will shut itself down. Each will wait for a few moments and will then try to maintain the cluster.
4. If the shutdown of `NodeA` is not successful, `NodeA` will try to maintain the cluster. If the shutdown of `NodeB` is not successful, `NodeB` will also try to maintain the cluster. This could result in two clusters that are unaware of each other (a split-brain situation) and data corruption will likely occur.

Suppose another configuration, in which neither node has `shutdown` set:

NodeA	NodeB
-----	-----
fence	fence
reset	reset

If the CXFS private network between `NodeA` and `NodeB` fails in this situation, each node would first try to fence the other and then try to reset the other, as before. However, if both of those actions fail, each would assume that the state of the other node is unknown. Therefore, neither node would try to maintain the cluster. The cluster will go down, but no data corruption will occur.

The split-brain problem may be avoided by using a tiebreaker node or by not using the shutdown setting on any server-capable node. You must not use shutdown if you use dynamic heartbeat monitoring.

The Cluster Database and CXFS Clients

The distributed cluster database (CDB) is central to the management of the CXFS cluster. Multiple synchronized copies of the database are maintained across the CXFS administration nodes in the pool (that is, those nodes installed with the `cluster_admin` software package). For any given CXFS Manager GUI `orcxfss_admin` task, the CXFS cluster daemons must apply the associated changes to the cluster database and distribute the changes to each CXFS administration node before another task can begin.

The client-only nodes in the pool do not maintain a local synchronized copy of the full cluster database. Instead, one of the daemons running on a CXFS administration node provides relevant database information to those nodes. If the set of CXFS administration nodes changes, another node may become responsible for updating the client-only nodes.

Metadata Server Functions

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

- Metadata logging
- File locking
- Buffer coherency
- Filesystem block allocation

All CXFS requests for metadata are routed over a TCP/IP network and through the 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 installed with the `cluster_admin` product can be defined as a server-capable administration node.
- Although you can configure multiple server-capable CXFS administration nodes to be potential metadata servers for a given filesystem, only the first of these nodes to mount the filesystem will become the *active* metadata server. The list of potential metadata servers for a given filesystem is ordered, but because of network latencies and other unpredictable delays, it is impossible to predict which node will become the active metadata server.
- A single server-capable node in the cluster can be the active metadata server for multiple filesystems at once.
- There can be multiple server-capable nodes that are active metadata servers, each with a different set of filesystems. However, a given filesystem has a single active metadata server on a single node.
- If the last potential metadata server for a filesystem goes down while there are active CXFS clients, all of the clients will be forced out of the filesystem. (If another potential metadata server exists in the list, recovery will take place. For more information, see "Metadata Server Recovery" on page 332.)
- If you are exporting the CXFS filesystem to be used with other NFS clients, the filesystem should be exported from the active metadata server for best performance. For more information on NFS exporting of CXFS filesystems, see "CXFS Mount Scripts" on page 325.

For more information, see "Flow of Metadata for Reads and Writes" on page 504.

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 pathnames to the files. CXFS lets you scale the shared-filesystem 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 fileserver. It combines the speed of near-local disk access with the flexibility, scalability, and reliability of clustering.

CXFS and Highly Available Services

You can use one of the following products or to provide highly available services (such as for NFS or Web) running on a CXFS filesystem:

- IRIX systems: IRIX FailSafe on IRIX systems
- SGI ProPack systems: SGI Cluster Manager for Linux (requires both the base and storage software plug-in packages)

The CXFS plug-in moves the CXFS metadata server along with applications that must run on the metadata server, such as DMF. This combination of CXFS and FailSafe or SGI Cluster Manager for Linux provides high-performance shared data access for highly available applications.

CXFS and IRIX FailSafe share the same infrastructure. SGI Cluster Manager for Linux has a separate infrastructure.

Hardware and Software Support

This section discusses the following:

- "Requirements"
- "Compatibility" on page 41

Requirements

CXFS requires the hardware and software specified in the release notes:

- Metadata servers running one of the following operating systems:
 - IRIX 6.5.30
 - SGI ProPack 5 SP2 (SLES 10 SP 1) on SGI Altix ia64 or SGI Altix XE x86_64

Note: For other client-only nodes, see the release notes and *CXFS MultiOS Client-Only Guide for SGI InfiniteStorage*

- At least one host bus adapter (HBA):
 - IRIX:
 - QLogic QLA2200, QLA2200F, QLA2310, QLA2310F, QLA2342, QLA2344
 - LSI Logic LSI7104XP-LC, LSI7204XP-LC
 - SGI ProPack:
 - QLogic QLA2200 (PPP and copper only), QLA2310, QLA2342, or QLA2344
 - LSI Logic LSI7104XP or LSI7204XP

Note: The LSI HBA requires the 01030600 firmware.

- Metadata servers that are dedicated to CXFS and filesystems work. See "Choose a Metadata Server that is Dedicated to CXFS Work" on page 144.
- All server-capable administration nodes must run the same type of operating system.
- A supported SAN hardware configuration.

Note: For details about supported hardware, see the Entitlement Sheet that accompanies the release materials. Using unsupported hardware constitutes a breach of the CXFS license.

- Use a network switch. (A network hub is not supported.) The switch should be at least 100baseT.
- A private 100baseT or Gigabit Ethernet TCP/IP network connected to each node.

Note: When using Gigabit Ethernet, do not use jumbo frames. For more information, see the `tgconfig` man page.

- Serial lines and/or supported Fibre Channel switches. For supported switches, see the release notes.

Either system reset or I/O fencing is required for all nodes. SGI recommends system reset for potential metadata servers. A cluster should have an odd number of server-capable nodes with CXFS services running.

- At least one host bus adapter (HBA) as specified in the release notes.
- RAID hardware as specified in the release notes.
- Adequate compute power for CXFS nodes, particularly metadata servers, which must deal with the required communication and I/O overhead. There should be at least 2 GB of RAM on the system.

A metadata server must have at least 1 processor and 1 GB of memory more than what it would need for its normal workload (non-CXFS work). In general, this means that the minimum configuration would be 2 processors and 2 GB of memory. If the metadata server is also doing NFS or Samba serving, then more memory is recommended (and the `nbuf` and `ncsize` kernel parameters should be increased from their defaults).

CXFS makes heavy use of memory for caching. If a very large number of files (tens of thousands) are expected to be open at any one time, additional memory over the minimum is also recommended. Use the following to determine the amount of memory required for your system:

$$2KB \times \text{number_of_inodes} = \text{metadata_server_memory}$$

In addition, about half of a CPU should be allocated for each Gigabit Ethernet interface on the system if it is expected to be run a close to full speed.

- To avoid problems during metadata server recovery/relocation, all potential metadata servers should have as much memory as the active metadata server.
- Licenses for CXFS and XVM. See the general release notes Chapter 4, "CXFS License Keys" on page 75.
- The XVM volume manager, which is provided as part of the IRIX release.
- If you use I/O fencing and `ipfilterd` on a node, the `ipfilterd` configuration must allow communication between the node and the `telnet` port on the switch.

A cluster is supported with as many as 64 nodes, of which as many as 16 can be server-capable administration nodes.

A cluster in which both CXFS and FailSafe are run (known as *coexecution*) is supported with a maximum of 64 nodes, as many as 8 of which can run FailSafe. The administration nodes must run IRIX; FailSafe is not supported on SGI ProPack nodes. Even when running with FailSafe, there is only one pool and one cluster. See "Overview of FailSafe Coexecution" on page 41, for further configuration details.

Compatibility

CXFS is compatible with the following:

- Data Migration Facility (DMF) and Tape Management Facility (TMF).
- Trusted IRIX. CXFS has been qualified in an SGI Trusted IRIX cluster with the Data Migration Facility (DMF) and Tape Management Facility (TMF).

If you want to run CXFS and Trusted IRIX, all server-capable administration nodes must run Trusted IRIX. Client-only nodes can be running IRIX. For more information, see Chapter 16, "Trusted IRIX and CXFS" on page 397.

- FailSafe (coexecution). See the "Overview of FailSafe Coexecution" on page 41, and the *FailSafe Administrator's Guide for SGI InfiniteStorage*.
- SGI Cluster Manager for Linux. See the *SGI InfiniteStorage Cluster Manager for Linux Administrator's Guide*.

Overview of FailSafe Coexecution

CXFS 6.5.10 or later and IRIS FailSafe 2.1 or later (plus relevant patches) may be installed and run on the same system.

A subset of nodes in a coexecution cluster can be configured to be used as FailSafe nodes; a coexecution cluster can have up to eight nodes that run FailSafe.

The cluster database contains configuration information about nodes, the cluster, logging information, and configuration parameters. If you are running CXFS, it also contains information about CXFS filesystems and CXFS metadata servers, which coordinate the information that describes a file, such as the file's name, size, location, and permissions; there is one active metadata server per CXFS filesystem. If you are running FailSafe, it also contains information about resources, resource groups, and failover policies. Figure 1-9 depicts the contents of a coexecution cluster database.

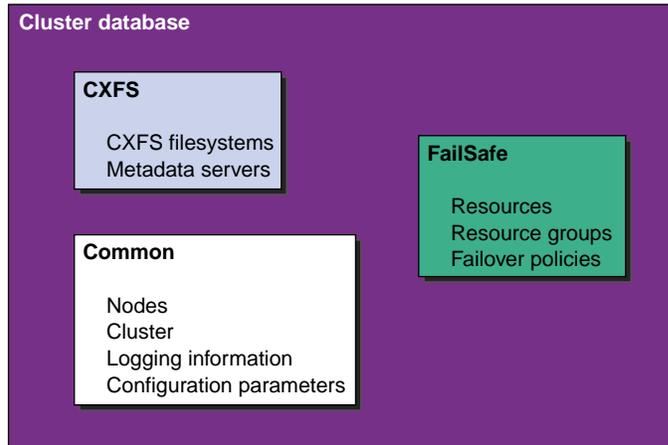


Figure 1-9 Contents of a Coexecution Cluster Database

In a coexecution cluster, a subset of the nodes can run FailSafe but all of the nodes must run CXFS. If you have both FailSafe and CXFS running, the products share a single cluster and a single database. There are separate configuration GUIs for FailSafe and CXFS. You can also view cluster information with the `clconf_info` command.

The administration nodes can perform administrative tasks for FailSafe or CXFS and they run the `fs2d` cluster database daemon, which manages the cluster database and propagates it to each administration node in the pool. All FailSafe nodes are administration nodes, but some CXFS nodes do not perform administration tasks and are known as client-only nodes.

For more information, see Chapter 15, "Coexecution with FailSafe" on page 389.

CXFS Tools Overview

CXFS provides a set of tools to manage the cluster. These tools execute only on the appropriate node types:

- Administration nodes:
 - `cxfsmgr`, which invokes the CXFS Manager graphical user interface (GUI)

Note: The GUI must be connected to a CXFS administration node, but it can be launched elsewhere; see "Starting the GUI" on page 184.

- `cxfs_admin`
- `clconf_info`
- `cxfs-config`
- Client-only nodes:
 - `cxfs_info`
 - `cxfs_admin`

Note: You can run `cxfs_admin` from any host that has the appropriate access permissions and network connections.

You can perform CXFS configuration tasks using the GUI or `cxfs_admin`. These tools update the cluster database, which persistently stores metadata and cluster configuration information.

Although these tools use the same underlying software command line interface (CLI) to configure and monitor a cluster, the GUI provides the following additional features, which are particularly important in a production system:

- You can click any blue text to get more information about that concept or input field. Online help is also provided with the **Help** button.
- 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 configuration information. In every task, the cluster configuration will not update until you click **OK**.
- Tasks take you step-by-step through configuration and management operations, making actual changes to the cluster configuration as you complete a task.

- The graphical tools can be run securely and remotely on any IRIX workstation or any computer that has a Java-enabled web browser, including Windows and Linux computers and laptops.

The `cxfs_admin` command waits for a command to be completed before continuing and provides <TAB> completion of commands. All input is validated before a command is completed. `cxfs_admin` also provides a step-by-step mode with auto-prompting and scripting capabilities. `cxfs_admin` provides better state information than the GUI, `clconf_info`, or `cxfs_info`.

The `cxfs_admin` command is convenient when performing basic configuration tasks or isolated single tasks in a production environment, or when running scripts to automate some cluster administration tasks.

You can use the `config` command in `cxfs_admin` to output the current configuration to a file and later recreate the configuration by using a command line option.

After the associated changes are applied to all online database copies in the pool, the view area in the GUI will be updated. You can use the GUI or the `cxfs_admin` and `clconf_info` commands to view the state of the database. (The database is a collection of files, which you cannot access directly.) On a client-only node, you can use the `cxfs_info` command.

For more details, see the following:

- "GUI Overview" on page 183
- "cxfs_admin Overview" on page 255
- "Saving and Recreating the Current Configuration with `cxfs_admin`" on page 300
- Chapter 17, "Monitoring Status" on page 399

Guaranteed-Rate I/O (GRIO) Version 2 and CXFS

CXFS supports guaranteed-rate I/O (GRIO) version 2 clients on all platforms, with a GRIO server on IRIX nodes or SGI ProPack nodes.

GRIO is disabled by default on CXFS Linux client-only nodes.

To enable GRIO on Linux client-only nodes, change the following line in `/etc/cluster/config/cxfs_client.options` from:

```
export GRIO2=off
```

to:

```
export GRIO2=on
```

Once installed in a cluster, you can run the following commands from any node in the cluster as a superuser:

- `grioadmin` provides stream and bandwidth management
- `griogps` is the comprehensive stream quality-of-service monitoring tool

Run the above tools with the `-h` (help) option for a full description of all available options.

The paths to the GRIO commands differ by platform. See Appendix D, "Operating System Path Differences" on page 533.

For details about GRIO, see the *Guaranteed-Rate I/O Version 2 Guide*. For other platform-specific limitations and considerations, see the *CXFS MultiOS Client-Only Guide for SGI InfiniteStorage*.

XVM Failover and CXFS

There are two versions of XVM failover. You must choose the appropriate version for your CXFS cluster. This section discusses the following:

- "XVM Failover Version 1 (V1) for Clusters with Only IRIX Nodes" on page 45
- "XVM Failover Version 2 (V2) for MultiOS Clusters" on page 46
- "RAID Units and XVM Failover V2" on page 52

XVM Failover Version 1 (V1) for Clusters with Only IRIX Nodes

You can use XVM failover V1 if your cluster has only IRIX nodes.

For failover V1 you can configure TP9300, TP9400, TP9500, and TP9700 RAID units with either SGIRDAC or SGI AVT host type. Where possible (when the cluster contains only IRIX nodes), SGIRDAC is the preferred method with failover V1.

Failover V1 uses the `scsifo(1M)` command and `/etc/failover.conf` file. For more information, see the `scsifo(1M)` and `failover(1M)` man pages.

Note: Failover V1 is independent of `xvm`, and the `xvm` failover V2 commands do not support failover V1. See "XVM Failover Version 2 (V2) for MultiOS Clusters" on page 46.

XVM Failover Version 2 (V2) for MultiOS Clusters

You should use XVM failover version 2 (V2) if you have a cluster that contains nodes running operating systems other than IRIX in order to get the maximum bandwidth and avoid LUN movement between RAID controllers (multiple nodes accessing the same LUN through different RAID controllers can degrade performance considerably). CXFS supports XVM failover V2 on all platforms.

This section discusses the following:

- "Overview of XVM Failover V2" on page 47
- "Configuring XVM Failover V2 using `failover2.conf`" on page 48
- "Generating a `failover2.conf` File" on page 49
- "Failover V2 Example Using Two Affinities" on page 50
- "Failover V2 Example Using Four Affinities" on page 51
- "Failover V2 Example for Two RAID Controllers" on page 52
- "XVM Commands Related to Failover V2" on page 52

On IRIX, the example file is installed in `/etc/failover2.conf`. SGI ProPack, the example file is installed in `/etc/failover2.conf.example`. For other platform-specific examples of `failover2.conf`, see the *CXFS MultiOS Client-Only Guide for SGI InfiniteStorage*.

For more information, see the comments in the `failover2.conf` file and the *XVM Volume Manager Administrator's Guide*.

Overview of XVM Failover V2

XVM failover V2 allows XVM to use multiple paths to LUNs in order to provide fault tolerance and static load balancing. Paths to LUNs are representations of links from the client HBA ports to the fabric and from the fabric to the RAID controller; they do not represent the path through the fabric itself.

In general, you want to evenly distribute the I/O to LUNs across all available host bus adapters and RAID controllers and attempt to avoid blocking in the SAN fabric. The ideal case, from a performance standpoint, is to use as many paths as connection endpoints between two nodes in the fabric as possible with as few blocking paths as possible in the intervening SAN fabric.

The example in Figure 1-10 shows two RAID controllers and the LUN's they own. All LUNs are visible from each controller, therefore, each LUNs can be accessed by each path (1A, 1B, 2A, and 2B). However, the controller for RAID A is preferred for LUN 0 and LUN 2, and the controller for RAID B is preferred for LUN 1 and LUN 3.

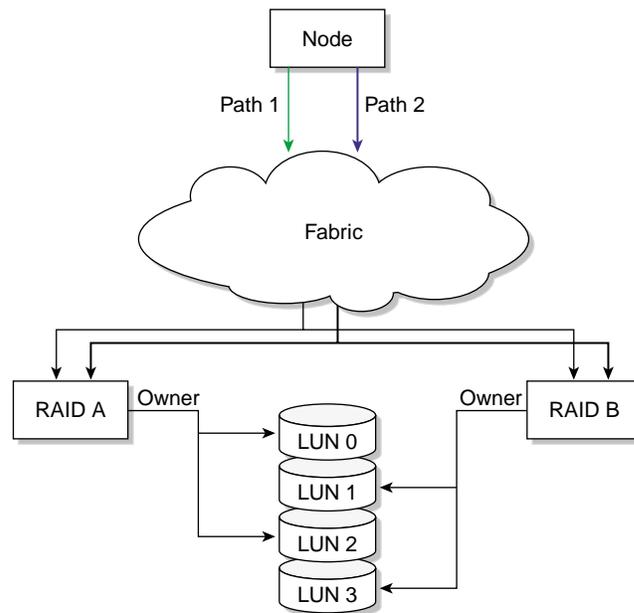


Figure 1-10 Disk Paths

Note: The Mac OS X platform provides dynamic load balancing between all paths to the same RAID controller. In this case, the system will only show one path per controller to each LUN with local HBA ports or individual paths not visible.

Configuring XVM Failover V2 using `failover2.conf`

XVM failover V2 stores path information in the `failover2.conf` file. You must configure the `failover2.conf` file for each node. The entries in this file define failover attributes associated with a path to the storage. Entries can be in any order.

The `failover2.conf` file uses the following keywords:

- `preferred` indicates the best path for accessing each XVM physvol. This is the path that will be used at startup barring failure. There is no default preferred path.
- `affinity` groups all of the paths to a particular RAID controller that can be used in harmony without causing LUN ownership changes for a LUN between RAID controllers, which would result in poor disk performance. An affinity group for a LUN should not contain paths that go to different RAID groups. The `affinity` value also determines the order in which these groups will be tried in the case of a failure, from lowest number to highest number. The valid range of `affinity` values is 0 (lowest) through 15 (highest). The path used starts with the affinity of the currently used path and increases from there. For example, if the currently used path is `affinity=2`, all `affinity=2` paths are tried, then all `affinity=3`, then all `affinity=4`, and so on; after `affinity=15`, failover V2 wraps back to `affinity=0` and starts over. Before you configure the `failover2.conf` file, the initial value for all paths defaults to `affinity=0`.

SGI recommends that the `affinity` values for a particular RAID controller be identical on every node in the CXFS cluster.

You may find it useful to specify affinities starting at 1. This makes it easy to spot paths that have not yet been configured because they are assigned a default affinity of 0. For example, if you added a new HBA but forgot to add its paths to the `failover2.conf` file, all of its paths would have an `affinity=0`, which could result in LUN ownership changes if some paths point to controller A and others point to controller B. Using this convention would not avoid this problem, but would make it easier to notice. If you use this convention, you must do so for the entire cluster.

Note: If you use the method where you do not use affinity 0 and you do not define all of the paths in the `failover2.conf` file, you will have a affinity group using an unknown controller. If in the example where you are using affinities 1 and 2, if you are using affinity 2 as your current path and there is a failover, you will failover to affinity 0, which could use the same RAID controller and thus fail again or might use the other RAID controller. If there are multiple unspecified paths in the affinity 0 group, you might be mixing different RAID controllers in the same affinity group. This is only a performance issue, but you should fix any paths using the default `affinity=0` value by adding them to the `failover2.conf` file and using an appropriate affinity value.

You can use the `affinity` value in association with the `XVM foswitch` command to switch an `XVM physvol` to a physical path of a defined `affinity` value. For more information, see "XVM Commands Related to Failover V2" on page 52 and the *XVM Volume Manager Administrator's Guide*.

Generating a `failover2.conf` File

The easiest method to generate a `failover2.conf` file is to run the following command (on platforms other than Windows²):

```
# xvm show -v phys | grep affinity > templatefile
```

The entries in the output only apply to already-labeled devices. Values within `< >` angle brackets are comments; you can delete them or ignore them.

If all nodes have correctly configured `failover2.conf` files, an affinity change in one node will signal all other nodes in the cluster that the node has changed affinity for a LUN, allowing the other nodes to change to the same affinity (the same RAID controller). You can also use the `foswitch -cluster` command to cause all nodes in a cluster to either return to their preferred paths or move to a specific affinity. See "XVM Commands Related to Failover V2" on page 52.

² For information about generating a `failover2.conf` file for Windows, see the *CXFS MultiOS Client-Only Guide for SGI InfiniteStorage*

Failover V2 Example Using Two Affinities

The following example groups the paths for lun3 and the paths for lun4:

```
/dev/xscsi/pci0004:00:01.1/node200900a0b813b982/port1/lun3/disc  affinity=1
/dev/xscsi/pci0004:00:01.1/node200900a0b813b982/port2/lun3/disc  affinity=1
/dev/xscsi/pci0004:00:01.0/node200900a0b813b982/port1/lun3/disc  affinity=1
/dev/xscsi/pci0004:00:01.0/node200900a0b813b982/port2/lun3/disc  affinity=1    preferred
/dev/xscsi/pci0004:00:01.1/node200800a0b813b982/port1/lun3/disc  affinity=3
/dev/xscsi/pci0004:00:01.0/node200800a0b813b982/port1/lun3/disc  affinity=3
/dev/xscsi/pci0004:00:01.1/node200800a0b813b982/port2/lun3/disc  affinity=3
/dev/xscsi/pci0004:00:01.0/node200800a0b813b982/port2/lun3/disc  affinity=3

/dev/xscsi/pci0004:00:01.1/node200900a0b813b982/port1/lun4/disc,  affinity=1
/dev/xscsi/pci0004:00:01.1/node200900a0b813b982/port2/lun4/disc,  affinity=1
/dev/xscsi/pci0004:00:01.0/node200900a0b813b982/port1/lun4/disc,  affinity=1
/dev/xscsi/pci0004:00:01.0/node200900a0b813b982/port2/lun4/disc,  affinity=1
/dev/xscsi/pci0004:00:01.1/node200800a0b813b982/port1/lun4/disc,  affinity=3
/dev/xscsi/pci0004:00:01.1/node200800a0b813b982/port2/lun4/disc,  affinity=3  preferred
/dev/xscsi/pci0004:00:01.0/node200800a0b813b982/port1/lun4/disc,  affinity=3
/dev/xscsi/pci0004:00:01.0/node200800a0b813b982/port2/lun4/disc,  affinity=3
```

The order of paths in the file is not significant. Paths to the same LUN are detected automatically. Without this file, all paths to each LUN would have affinity 0 and there would be no preferred path. Setting a preferred path ensures that multiple paths will be used for performance. If no path is designated as preferred, the path used to the LUN is arbitrary based on the order of device discovery. There is no interaction between the preferred path and the affinity values.

This file uses affinity to group the RAID controllers for a particular path. Each controller has been assigned an affinity value. It shows the following:

- There is one PCI card with two ports off of the HBA (pci04.01.1 and pci04.01.0)
- There are two RAID controllers, node200800a0b813b982 and node200900a0b813b982
- Each RAID controller has two ports that are identified by port1 or port2
- Each LUN has eight paths (via two ports on a PCI card, two RAID controllers, and two ports on the controllers)

- There are two affinity groups for each LUN, `affinity=1` and `affinity=3`
- There is a preferred path for each LUN

Failover will exhaust all paths to lun3 from RAID controller node200900a0b813b982 (with `affinity=1` and the preferred path) before moving to RAID controller node200800a0b813b982 paths (with `affinity=3`)

Failover V2 Example Using Four Affinities

The following example uses four affinities to associate the two HBA ports with each of the available two ports on the RAID's two controllers:

```
/dev/xscsi/pci0004:00:01.1/node200900a0b813b982/port1/lun4/disc  affinity=1
/dev/xscsi/pci0004:00:01.1/node200900a0b813b982/port2/lun4/disc  affinity=3
/dev/xscsi/pci0004:00:01.0/node200900a0b813b982/port1/lun4/disc  affinity=1
/dev/xscsi/pci0004:00:01.0/node200900a0b813b982/port2/lun4/disc  affinity=3
/dev/xscsi/pci0004:00:01.1/node200800a0b813b982/port1/lun4/disc  affinity=4
/dev/xscsi/pci0004:00:01.1/node200800a0b813b982/port2/lun4/disc  affinity=3 preferred
/dev/xscsi/pci0004:00:01.0/node200800a0b813b982/port1/lun4/disc  affinity=4
/dev/xscsi/pci0004:00:01.0/node200800a0b813b982/port2/lun4/disc  affinity=3
```

Each affinity associates the two host adapter ports with a single RAID controller port. The declaration of these eight associations completely defines all of the available paths to a single RAID LUN.

These eight associations also represent the order in which the paths are tried in a failover situation. Failover begins by trying the other paths within the current affinity and proceeds in an incremental manner through the affinities until either a working path is discovered or all possible paths have been tried. The paths will be tried in the following order:

1. `affinity=3` (the affinity of the current path), which is associated with RAID controller A port 2
2. `affinity=4`, which is associated with RAID controller A port 1
3. `affinity=1`, which is associated with raid controller B port 1
4. `affinity=2`, which is associated with raid controller B port 2

Failover V2 Example for Two RAID Controllers

The following example for IRIX shows two RAID controllers, 200800a0b818b4de and 200900a0b818b4de for lun4vol:

```
/dev/dsk/200800a0b818b4de/lun4vol/c2p2 affinity=1 preferred
/dev/dsk/200800a0b818b4de/lun4vol/c2p1 affinity=1
/dev/dsk/200900a0b818b4de/lun4vol/c2p2 affinity=3
/dev/dsk/200900a0b818b4de/lun4vol/c2p1 affinity=3
```

XVM Commands Related to Failover V2

The following are useful XVM commands related to failover V2:

```
xvm help -verbose foconfig
xvm help -verbose foswitch
xvm help -verbose show
xvm foconfig -init
xvm foswitch -cluster -preferred physvol/name (switch phys/name in all nodes in cluster to preferred path)
xvm foswitch -preferred physvol
xvm foswitch -affinity 1 physvol
xvm foswitch -dev newdev
xvm foswitch -cluster -affinity 1 phys
xvm foswitch -cluster -setaffinity X phys/name (switch phys/name in cluster to affinity "X")
xvm show -verbose physvol
xvm show -verbose physvol | fgrep affinity > templatefile
```

For details, see the *XVM Volume Manager Administrator's Guide*.

Note: The `xvm` command is provided on all CXFS platforms. However, client-only nodes support only read-only commands.

RAID Units and XVM Failover V2

This section discusses the following:

- "TP9100, RM610/660, and Failover V2" on page 53
- " TP9300, TP9500, TP9700, and S330 and Failover V1 and V2" on page 53
- "SGI InfiniteStorage 220 and Failover V2" on page 54

For more information about firmware levels, see "RAID Firmware" on page 64.

TP9100, RM610/660, and Failover V2

The TP9100 and RM610/660 RAID units do not have any host type failover configuration. Each LUN should be accessed via the same RAID controller for each node in the cluster because of performance reasons. These RAID configurations behave and have the same characteristics as the SGIAVT mode discussed below.

TP9100 1 GB and 2 GB:

- SGIRDAC mode requires that the array is set to `multiport`
- SGIAVT mode requires that the array is set to `multitid`

TP9300, TP9500, TP9700, and S330 and Failover V1 and V2

The TP9300, TP9500, and TP9700 RAID units will behave differently depending on the host type that is configured:

- SGIRDAC mode requires all I/O for a LUN to take place through the RAID controller that currently owns the LUN. Any I/O sent to a RAID controller that does not own the LUN will return an error to the host that sent the request. In order for the LUN to be accessed via the alternate controller in a RAID array, it requires the failover driver software on a host to send a command to the backup controller instructing it to take ownership of the specified LUN. At that point, the ownership of the LUN is transferred to the other LUN and I/O can take place via new owner. Other hosts in the cluster will detect this change and update their I/O for the LUN to use a path to the RAID controller that now owns the LUN. Only XVM failover V1 can successfully control RAID configurations in SGIRDAC mode.

SGIRDAC mode is supported under all revisions in RAID firmware section for these models.

- SGIAVT mode also has the concept of LUN ownership by a single RAID controller. However, LUN ownership change will take place if any I/O for a given LUN is received by the RAID controller that is not the current owner. The change of ownership is automatic based on where I/O for a LUN is received and is not done by a specific request from a host failover driver. The concern with this mode of operation is that when a host in the cluster changes I/O to a different RAID controller than that used by the rest of the cluster, it can result in severe performance degradation for the LUN because of the overhead involved in constantly changing ownership of the LUN. Either XVM failover V1 or V2 can

successfully control RAIDs in SGI_{IAVT} mode (TP9400 does not accept SGI_{IAVT} mode).

Failover V2 requires that you configure TP9300, TP9500, TP9700, and S330 RAID units with SGI_{IAVT} host type and the 06.12.18.xx code or later be installed.

TP9700 use of SGI_{IAVT} requires that 06.15.17xx. code or later be installed.

SGI InfiniteStorage 220 and Failover V2

XVM failover V2 support requires SGI ProPack 5 SP 1 or later.

GPT Labels and CXFS

CXFS supports XVM labels on LUNs with GUID partition table (GPT) labels as well LUNs with SGI disk volume header (DVH) labels. A CXFS cluster can contain LUNs that have GPT labels and LUNs that have DVH labels.

You can create these labels on SGI ProPack server-capable nodes and Linux third-party clients. The GPT label puts header data in sector 1 of a LUN, leaving sector 0 for a master boot record. Partition information is stored in a variable number of sectors, starting at sector 2. XVM requires two partitions on a GPT-labeled LUN, one for XVM metadata and the other for the user data. XVM assumes ownership of the LUN and access to a particular LUN could be fenced.

Note: CXFS supports a GPT-labeled LUN greater than 2 TB in size. However, being able to label a LUN does not mean that the system is able to recognize and use it. The operating systems in the cluster will determine whether you can actually use a LUN of a given size. If a LUN is set up as greater than 2-TB in size but if an OS of a node in a cluster cannot support a greater-than-2-TB LUN, then this node will not be able to share or even access data on this LUN.

When creating a GPT partition table for XVM to use, the first partition size should be at least 2 MB, just large enough to hold the XVM metadata (such as volume and slice information). The second partition for the volume data should be the rest of the LUN. You can place the start of the second partition anywhere after the first partition that will give good performance, such as on a boundary of the RAID's stripe width.

If you have a cluster with SGI ProPack server-capable nodes, you should use them to create the GPT label. If you have IRIX server-capable nodes, you must use another

client node to put the GPT label on the LUN. After these LUNs are GPT labeled and partitioned, you can use IRIX to write the XVM labels.

If the operating system is capable of specifying the start of a partition as a sector number, place the start of data exactly on a boundary for good performance:

- For SUSE LINUX Enterprise Server 9 (SLES 9) Service Pack 3, you can use the `mkpartsect` command to `parted`.
- For SLES 10, you can use the `mkpart` command to `parted` to specify the sector number using an `s` suffix. For example to make a partition starting at 2-MB into the LUN and ending at 961085440 (0x32490000), also a 2-MB byte boundary:

```
(parted) mkpart primary 4096s 961085440s
```

You can also use the `unit s` command to `parted` to set the input and display to default to sectors as the unit size. The partition for XVM data (partition 2) should have a start sector and length that is a common multiple of the RAID LUN's stripe width and the 16-KB page size for Altix or 4-KB page size for Altix XE. (If the partition is not made this way, the `xvm slice` command has options that you can use to place the slice on these boundaries.) If LUNs are then concatenated, I/O will be less likely to span RAID stripe-width boundaries or causing a read-modify-write inside the RAID if partial stripes are written.

For example, using a size with 2-MB boundaries:

```
(parted) unit s
(parted) print
Disk geometry for /dev/sdg: 0s - 71687371s
Disk label type: gpt
Number  Start   End     Size   File system  Name          Flags
1       34s    3906s   3873s
(parted) mkpart
Partition type? [primary]?
File system type? [ext2]? xfs
Start? 4096
End? 71687168
(parted) print
Disk geometry for /dev/sdg: 0s - 71687371s
Disk label type: gpt
Number  Start   End     Size   File system  Name          Flags
1       34s    3906s   3873s
2       4096s  71687168s  71683204s
```

- For other operating systems, see the operating system documentation.

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

Installation and Configuration Overview

This section discusses the following:

- "Client-Only Packages and Commands Installed" on page 56
- "Administration Packages and Commands Installed" on page 57
- "Installation Procedure Overview" on page 61

Client-Only Packages and Commands Installed

The following packages are installed on a client-only node:

- Application binaries, documentation, and support tools:

```
cxfs_client
cxfs_util
```

- Kernel libraries:

```
cxfs
eoe.sw.xvm
```

The following commands are shipped as part of the CXFS client-only package:

```
/usr/cluster/bin/cxfs_client (the CXFS client service)
/usr/cluster/bin/cxfs-config
/usr/cluster/bin/cxfsdump
```

These commands provide all of the services needed to include an IRIX or an SGI ProPack client-only node.

For more information, see the `cxfs_client` and `cxfsdump` man pages.

Table 1-1 Client-only Commands

Command	Description
<code>cxfs_client(1m)</code>	Controls the CXFS client daemon
<code>cxfs_info(1m)</code>	
<code>cxfsdump(1M)</code>	Gathers configuration information in a CXFS cluster for diagnostic purposes.
<code>cxfs_cp(1)</code>	Copies large files to and from a CXFS filesystem
<code>grioadmin(1m)</code>	Provides stream and bandwidth management
<code>griooos(1m)</code>	Provides comprehensive stream quality-of-service monitoring
<code>xvm(1m)</code>	Invokes the XVM command line interface

Administration Packages and Commands Installed

The following packages are installed on an administration node:

- Application binaries, documentation, and support tools:

```
cluster_admin
cluster_control
cluster_services
cxfs_cluster
cxfs_util
```

- Kernel libraries:

```
cxfs
eoe.sw.xvm
```

- GUI tools:

– IRIX:

```
sysadm_base
sysadm_cluster
sysadm_cxfs
sysadm_xvm
```

– SGI ProPack:

```
cxfs-sysadm_base
cxfs-sysadm_cluster
cxfs-sysadm_cxfs
cxfs-sysadm_xvm
```

The following commands are shipped as part of the CXFS administration package and are located in `/usr/cluster/bin`:

```
ClusterDiags
cad
cad_ctrl
cbeutil
cdb-create-resource-type
cdb-exitop
cdb-init-CX-access
cdb-init-local-CX-nodes
cdb-init-local-HA-nodes
cdb-init-std-nodes
cdbBackup
cdbRestore
cdbconfig
cdbdelete
cdbhdr
cdbhnew
cdbhreinit
cdbhutil
clconf_info
clconf_stats
clconf_status
clconfd
cluster_mgr
cmgr
cmon_ctrl
cmond
cms_failconf
cms_intervene
cms_status
crsd
cxfs-config
cxfs_mds
```

```

cxfs_shutdown
cxfscp
cxfsdump
cxfslicense
ha_cilog
ha_cmds
ha_cxfs
ha_exec2
ha_execute_lock
ha_filelock
ha_fileunlock
ha_gcd
ha_ifd
ha_ifdadmin
ha_macconfig2
ha_srmd
hafence
plugin-preop

```

Table 1-2 summarizes the CXFS commands of most use on an administration node.

Table 1-2 Administration Commands

Command	Description
<code>cxfscp(1)</code>	Copies large files to and from a CXFS filesystem
<code>cbeutil(1M)</code>	Accesses the back-end cluster database
<code>cdbBackup(1M)</code>	Backs up the cluster database
<code>cdbRestore(1M)</code>	Restores the cluster database
<code>cdbconfig(1M)</code>	Configures the cluster database
<code>cdbutil(1M)</code>	Accesses the cluster database by means of commands that correspond to functions in the <code>libcdb</code> library.
<code>clconf_info</code>	Provides information about the cluster
<code>clconf_stats</code>	Provides CXFS heartbeat statistics for cluster

Command	Description
<code>clconf_status</code>	Provides provides a <code>curses</code> interface to display status information gathered by the <code>cad</code> daemon (this information is also displayed by the <code>cxfsmgr</code> command)
<code>clconfd</code>	Implements the cluster control daemon
<code>cluster_status(1M)</code>	Obtains configuration and status information
<code>cms_failconf(1M)</code>	Configures the action taken by the surviving nodes when a CXFS node loses membership (normally, you will use the GUI or <code>cxfs_admin</code> to perform these actions.)
<code>cmond(1M)</code>	Provides a framework for starting, stopping and monitoring process groups.
<code>crsd(1M)</code>	Controls system controllers on remote nodes in a cluster
<code>cxfs_admin(1M)</code>	Configures and administers the cluster database.
<code>cxfs-config(1M)</code>	Displays and checks configuration information in a CXFS cluster.
<code>cxfsdump(1M)</code>	Gathers configuration information in a CXFS cluster for diagnostic purposes.
<code>cxfslicense(1M)</code>	Reports the status of license keys
<code>cxfsmgr(1M)</code>	Invokes the CXFS GUI, which provides access to the tasks that help you set up and administer your CXFS filesystems and provides icons representing status and structure
<code>cxfs_shutdown</code>	Shuts down CXFS in the kernel and CXFS daemons
<code>fs2d(1M)</code>	Implements the distributed cluster database
<code>grioadmin(1m)</code>	Provides stream and bandwidth management
<code>griogos(1m)</code>	Provides comprehensive stream quality-of-service monitoring
<code>haStatus(1M)</code>	Obtains configuration and status information
<code>ha_cilog(1M)</code>	Logs messages using the cluster logging subsystem (used only for FailSafe coexecution)
<code>ha_cmds(1M)</code>	Provides node membership services (used only for FailSafe coexecution)
<code>ha_exec2(1M)</code>	Executes a command on a cluster node (used only for FailSafe coexecution)
<code>ha_gcd(1M)</code>	Provides an atomic messaging service (used only for FailSafe coexecution)
<code>ha_ifd(1M)</code>	Provides support for all system resource manager (SRM) actions for IP addresses (used only for FailSafe coexecution)

Command	Description
ha_ifdadmin(1M)	Provides information about IP addresses from the <code>ha_ifd</code> daemon (used only for FailSafe coexecution)
ha_macconfig2(1M)	Displays and changes the MAC address of a network interface (used only for FailSafe coexecution)
ha_srmd(1M)	Provides resource management services (used only for FailSafe coexecution)
hafence(1M)	Administer the CXFS I/O fencing configuration stored in the cluster database (normally, you will perform this task using the GUI or <code>cxfs_admin</code>)
listclients(1M)	Lists system administration clients
sysadmd(1M)	Allows clients to perform remote system administration
xvmgr(1M)	Invokes the XVM GUI, which provides access to the tasks that help you set up and administer your logical volumes and provides icons representing status and structure

Installation Procedure Overview

Following is the order of installation and configuration steps:

1. Install the operating system (if not already done). See the CXFS release notes for supported levels.
2. Install and verify the RAID. See Chapter 2, "SGI RAID for CXFS Clusters" on page 63
3. Install and verify the switch. See Chapter 3, "Switches" on page 67.
4. Obtain and install the CXFS license keys and (if needed) XVM license keys. See Chapter 4, "CXFS License Keys" on page 75.
5. Prepare the node, including adding a private network.
6. Install the CXFS software. For details, see:
 - Chapter 6, "IRIX CXFS Installation" on page 101
 - Chapter 7, "SGI ProPack CXFS Installation" on page 113
 - Installation chapters in *CXFS MultiOS Client-Only Guide for SGI InfiniteStorage*

7. Configure the cluster to define the new node in the pool, add it to the cluster, start CXFS services, and mount filesystems. See "Guided Configuration Tasks" on page 205.

SGI RAID for CXFS Clusters

This chapter discusses SGI RAID for CXFS clusters:

- "RAID Hardware" on page 63
- "RAID Firmware" on page 64
- "Number of LUNs Supported" on page 65
- "RAID Verification" on page 66

For additional updates, see the CXFS release notes.

RAID Hardware

CXFS supports the following RAID hardware:

SGI InfiniteStorage 10000
SGI InfiniteStorage 6700
SGI InfiniteStorage 4500
SGI InfiniteStorage 4000
SGI InfiniteStorage 220 (Fibre Channel)
SGI RM610
SGI RM660
SGI InfiniteStorage 6700
SGI TP9700
SGI TP9500S (serial ATA)
SGI TP9500
SGI TP9400
SGI TP9300S (serial ATA)
SGI TP9300
SGI TP9100
SGI S330

The SGI RAID will be initially installed and configured by SGI personnel.

RAID Firmware

SGI RAID supports the following firmware:

Note: SGI InfiniteStorage 220 does not support online updates of the controller firmware.

- SGI RM610 and RM660 running version 5.12b or later.
- SGI InfiniteStorage 6700 supports controller firmware version V3.00.
- The TP9700 9.14 CD contains the required controller firmware and NVSRAM files. The 06.14.xx.xx code or later must be installed.

Note: If running with IRIX 6.5.26, the TP9700 requires IRIX patch 5822 or its successors.

- The TP9500S 8.0 CD contains the required controller firmware and NVSRAM files. The 05.41.xx.xx code or later must be installed.
- The TP9400/TP9500 6.0 CD contains the required controller firmware and NVSRAM files. The 05.30.xx.xx code or later must be installed.
- The TP9400 4.0 CD contains the required controller firmware and NVSRAM files for the 4774 or 4884 units:
 - If you have a 4774 unit, the 04.01.xx.xx , 04.02.xx.xx, or 05.30.xx.xx code or later must be installed
 - If you have a 4884 unit, the 04.02.xx.xx code is installed by default
- The TP9300S 8.0 CD contains the required controller firmware and NVSRAM files. The 05.41.xx.xx code or later must be installed if using 2882 controllers, or 05.42.xx.xx code or later if using 2822 controllers.

Note: The initial TP9300S used 2882 controllers in the controller module. This product was later replaced with a 2822 controllers (still using the TP9300S marketing code). With the release of the 2822 controller, SATA disk drives can be installed in the controller module (the 2882 did not have disk drives installed in the controller module).

- The TP9300 7.0 CD contains the required controller firmware and NVSRAM files. The 05.33.xx.xx code or later must be installed.
- The TP9100 4.0 CD contains the required version 7.75 controller firmware for the 1-Gbit TP9100. Supported via special request with optical attach (other conditions may apply).
- The TP9100 5.0 CD contains the required version 8.40 firmware for the 2-Gbit TP9100. (Direct connect is supported only for IRIX nodes using reset lines.)

Note: The TP9100 is limited to 64 host connections.

- The TP9300 8.42 CD (TPSSM 8.42) contains the required 8.42 firmware for the S330. See also "XVM Failover and CXFS" on page 45.

Number of LUNs Supported

By default, the RAID firmware supports a maximum number of logical units (LUNs). If additional LUNs are required, you must obtain a separate software-enabling key; this key will support a larger number of LUNs in separate partitions, which requires that the Fibre Channel ports be mapped to a partition. Contact your SGI sales representative for the SGI software partitioning key.

The maximum depends upon the code installed, as shown in Table 2-1.

Table 2-1 Number of LUNs Supported

Firmware Level	Default LUN Maximum	LUN Maximum with a Partitioning Key
04.01.xx.xx	32	128
04.02.xx.xx	32	128
05.30.xx.xx	32	1024
05.33.xx.xx	32	2048
05.40.xx.xx	256	2048
06.14.xx.xx	32	2048

RAID Verification

To verify that the SGI RAID is properly installed and ready for use with CXFS, you can dump the RAID's profile and verify the controller software revisions.

Switches

This chapter discusses the following:

- "Brocade Switch" on page 67
- "QLogic Fibre Channel Switch" on page 72

Brocade Switch

This section discusses the following:

- "Brocade Firmware" on page 67
- "Verifying the Brocade Switch Firmware Version" on page 69
- "Verifying the Brocade License" on page 69
- "Limiting telnet Sessions" on page 70
- "Changing the Brocade FC Cable Connections" on page 71

Brocade Firmware

All Brocade switches contained within the SAN fabric must have the appropriate Brocade firmware, shown in Table 3-1.

Note: There are issues when upgrading from firmware v4.1.1. See Technical Information Bulletin 201240 on Supportfolio for details:

<http://support.sgi.com>

Table 3-1 Brocade Firmware

Switch	Ports	Speed (Gb/s)	Minimum Firmware
200E	8, 16	4	5.2.2
2400	8	1	2.6.2d
2800	16	1	2.6.2d
3200	8	2	3.2.1c
3250	8	2	5.2.2
3252	8	2	5.2.2
3800	16	2	3.2.1c
3850	16	2	5.2.2
3852	16	2	5.2.2
3900	32	2	5.2.2
4100	32	4	5.2.2
4900	16, 32, 64	4	5.2.2
5000	16, 32, 64	4	5.2.2
12000	32, 64, dual 64	2	5.0.5d
24000	32, 64, 128	2	5.2.2
48000	32 through 256	4	5.2.2

If the current firmware level of the switches must be upgraded, please contact your local SGI service representative or customer support center.

The Brocade switch must be configured so that its Ethernet interface is accessible (using `telnet`) from all CXFS administration nodes. The fencing network connected to the Brocade switch must be physically separate from the private heartbeat network.



Caution: The admin state must be free in order for I/O fencing to succeed.

Switches using 4.x.x.x or later firmware permit multiple `telnet` sessions. However, CXFS I/O fencing requires a `telnet` lockout for global mutual exclusion when a fencing race occurs. Therefore, you must configure these switches to set the

maximum allowed simultaneous telnet sessions for the admin user to 1. (Brocade switches running 3.x.x.x firmware are shipped with the required restrictions configured by default).

Verifying the Brocade Switch Firmware Version

To verify the firmware version, log into the switch as user admin and use the version command, as shown in the following example:

```
workstation% telnet brocade1
Trying 169.238.221.224...
Connected to brocade1.example.com
Escape character is '^]'.

Fabric OS (tm) Release v2.6.0d

login: admin
Password:
brocade1:admin> version
Kernel:      5.4
Fabric OS:   v2.6.0d          <== Firmware Revision
Made on:    Fri May 17 16:33:09 PDT 2002
Flash:      Fri May 17 16:34:55 PDT 2002
BootProm:   Thu Jun 17 15:20:39 PDT 1999
brocade1:admin>
```

Verifying the Brocade License

To verify the Brocade license, log into the switch as user admin and use the licenseshow command, as shown in the following example:

```
brocade:admin> licenseshow
dcRyzyScSedSz0p:
  Web license
  Zoning license
  SES license
  Fabric license
SQQQSyddQ9TRRdUP:
  Release v2.2 license
```

Limiting telnet Sessions

You must limit the maximum allowed simultaneous telnet session

Brocade 200E/3250/3252/3850/3852/3900/4100/4900/5000 and telnet

To limit the maximum allowed simultaneous telnet sessions for the admin user to 1 on the Brocade 200E/3250/3252/3850/3852/3900/4100/4900/5000, do the following:

1. Connect to the switch via the telnet command and log in as root.
2. Issue the sync command to avoid filesystem corruption:

```
# sync
```
3. Edit the /etc/profile file to change the max_telnet_sessions from 2 to 1 and place the information in a new file. For example:

```
# cd /etc  
# sed -e 's/max_telnet_sessions=2/max_telnet_sessions=1/' profile >profile.new
```

4. Distribute the edited profile file to both partitions on both central processors. For example:

```
# cp profile.new profile  
# cp profile.new /mnt/etc/profile
```

5. Issue the sync command again to avoid filesystem corruption:

```
# sync
```

Brocade 12000/24000/48000 and telnet

To limit the maximum allowed simultaneous telnet sessions for the admin user to 1 on the Brocade 12000/24000/48000, do the following:

1. Connect to the switch via the telnet command and log in as root.
2. Use the haShow command to make sure that both central processors are up. This is indicated by the message Heartbeat Up within the output of the haShow command. If it is not up, wait a few minutes and run haShow again to check for the status.

3. Issue the `sync` command on the filesystems to avoid filesystem corruption:

```
# rsh 10.0.0.5 sync
# rsh 10.0.0.6 sync
```

4. Edit the `/etc/profile` file to change the `max_telnet_sessions` from 2 to 1 and place the information in a new file. For example:

```
# cd /etc
# sed -e 's/max_telnet_sessions=2/max_telnet_sessions=1/' profile >profile.new
```

5. Distribute the new profile to both partitions and central processors. For example:

```
# rcp /etc/profile.new 10.0.0.5:/etc/profile
# rcp /etc/profile.new 10.0.0.5:/mnt/etc/profile
# rcp /etc/profile.new 10.0.0.6:/etc/profile
# rcp /etc/profile.new 10.0.0.6:/mnt/etc/profile
```

6. Issue the `sync` command again to avoid filesystem corruption:

```
# rsh 10.0.0.5 sync
# rsh 10.0.0.6 sync
```

Changing the Brocade FC Cable Connections

To change Brocade Fibre Channel cable connections used by nodes in the CXFS cluster, do the following:

1. Cleanly shut down CXFS services on the nodes affected by the cable change. Use the CXFS GUI or `cxfs_admin`.
2. Rearrange the cables as required.
3. Restart CXFS services.
4. Reconfigure I/O fencing if required. You must perform this step if I/O fencing is enabled on the cluster and if you added/removed any Brocade switches. You must use the CXFS GUI or `cxfs_admin` to add or remove switches from the CXFS configuration as required.
5. If any CXFS client nodes are connected to a new (or different) Brocade switch, restart CXFS services on those nodes. This will ensure that the CXFS administration servers can correctly identify the Brocade ports used by all clients.

QLogic Fibre Channel Switch

All QLogic Fibre Channel (FC) switches contained within the SAN fabric must have the appropriate QLogic firmware installed, as shown in Table 3-2.

Table 3-2 QLogic FC Switch Firmware

QLogic FC Switch Model	SANbox Name	Minimum Firmware
SB2A-16A/B	2-16	4.0
SB2B-08A/B	2-8	4.0
SB2C-16BSE	2-64	4.0
SB5200-08/12/16/20A	5200	V5.0.1.10.0
SB9200-32B	9200	V6.2.0.8.0

For more information, see the QLogic *SANbox2-64 Switch Management User's Guide*.



Caution: The admin state is required for I/O fencing. To avoid interference with fencing, release admin mode as soon as possible. Do not leave admin mode sessions open.

The default port configuration on a QLogic 9200 FC switch is not compatible with the CXFS environment. To use the appropriate port configuration, change the following parameters:

LinkSpeed	Set to the appropriate value, such as 2 for 2 GB/s. (In some cases, <code>Auto</code> does not function properly.)
PortType	Enter the appropriate type, usually <code>F</code> . (You cannot use the <code>GL</code> autonegotiated mode.)
NoClose	Set to <code>True</code> to prevent the Fibre Channel circuit from shutting down during a host reboot.
IOStreamGuard	Set to <code>Enable</code> if the port is connected to a host HBA or to <code>Disable</code> if the port is connected to a storage HBA.

(You cannot use Auto mode because most HBAs cannot negotiate this.)

To modify these parameters, use the admin command. For example, for a port connected to an SGI Altix:

```
SANbox #> admin start
```

```
SANbox (admin) #> config edit
```

```
The config named default is being edited.
```

```
SANbox (admin-config) #> set config port 31
```

A list of attributes with formatting and current values will follow.
Enter a new value or simply press the ENTER key to accept the current value.
If you wish to terminate this process before reaching the end of the list
press 'q' or 'Q' and the ENTER key to do so.

```
Configuring Port Number: 31
```

```
-----
```

```
AdminState      (1=Online, 2=Offline, 3=Diagnostics, 4=Down) [Online]
LinkSpeed       (1=1Gb/s, 2=2Gb/s, 4=4Gb/s, A=Auto)   [Auto ]  2
PortType        (GL / G / F / FL / Donor)              [GL    ]  F
SymPortName     (string, max=32 chars)                          [Port31] Altix45
ALFairness      (True / False)                                    [False ]
DeviceScanEnable (True / False)                                    [True  ]
ForceOfflineRSCN (True / False)                                    [False ]
ARB_FF          (True / False)                                    [False ]
InteropCredit   (decimal value, 0-255)                          [0     ]
ExtCredit       (dec value, increments of 15, non-loop only) [0     ]
FANEnable       (True / False)                                    [True  ]
AutoPerfTuning  (True / False)                                    [True  ]
MSEnable        (True / False)                                    [True  ]
NoClose         (True / False)                                    [False ]  True
IOStreamGuard   (Enable / Disable / Auto)            [Auto  ]  Enable
PDISCPingEnable (True / False)                                    [True  ]
```

Finished configuring attributes.

This configuration must be saved (see config save command) and
activated (see config activate command) before it can take effect.

3: Switches

To discard this configuration use the config cancel command.

....

```
SANbox (admin-config) #> config save
```

The config named default has been saved.

```
SANbox (admin) #> config activate
```

The currently active configuration will be activated.

Please confirm (y/n): [n] **y**

```
SANbox (admin) #> admin end
```

```
SANbox #> show config port 31
```

Configuration Name: default

Port Number: 31

AdminState	Online
LinkSpeed	2Gb/s
PortType	F
SymbolicName	Altix45
ALFairness	False
DeviceScanEnabled	True
ForceOfflineRSCN	False
ARB_FF	False
InteropCredit	0
ExtCredit	0
FANEnabled	True
AutoPerfTuning	True
MSEnabled	True
NoClose	True
IOStreamGuard	Enabled
PDISCPingEnabled	True

CXFS License Keys

Note: On SGI Altix and SGI Altix XE hardware platforms running SGI ProPack 5.0 or later, *CPU count* is the number of processor sockets. On all other hardware platforms, or SGI Altix XE hardware running any other operating system, *CPU count* is the number of processor cores.

The licensing used for SGI ProPack server-capable nodes is based on the SGI License Key (LK) software. For IRIX server-capable nodes, the licensing is based on the FLEXlm product from Macrovision Corporation. In all clusters, only server-side licensing is supported.

For the purposes of licensing, hyperthreaded CPUs are counted as a single processor, while multicore processors are counted as multiple processors. Therefore, a dual-core processor will be counted as 2 CPUs for the purposes of licensing with CXFS. A hyperthreaded CPU or hyperthreaded core would only be counted as a single CPU for licensing purposes.

This section discusses the following:

- "Server-Side Licensing Overview" on page 76
- "Licensing on IRIX Clusters" on page 82
- "Licensing on SGI ProPack 5 Clusters" on page 83
- "Verifying the License Keys" on page 84
- "For More Information About Licensing" on page 92

Server-Side Licensing Overview

Note: As of CXFS 4.2, all server-capable nodes running 4.2 and client-only nodes running 4.2 require server-side licensing. If **all** existing client-only nodes are running a prior supported release, they may continue to use client-side license as part of the rolling upgrade policy until they are upgraded to 4.2. All client-only nodes in the cluster must use the same licensing type — if any client-only node in the cluster is upgraded to 4.2 or if a new 4.2 client-only node is added, then all nodes must use server-side licensing. For more information, see "CXFS Release Versions and Rolling Upgrades" on page 307.

CXFS server-side licensing uses license keys on the CXFS server-capable nodes; it does not require node-locked license keys on CXFS client-only nodes. The license keys are node-locked to each server-capable node and specify the number and size of client-only nodes that may join the cluster membership.

Server-side licensing provides flexibility when changing the CXFS cluster configuration, such as the following: adding nodes, changing the number of CPUs in one host, or using a license key part-time from different nodes connected to the cluster.

Licensing Requirements

Server-side licensing requires the following license keys on each server-capable administration node:

- CXFS_SS feature license key. The server license key specifies the maximum number of CPUs on the server. This license key is node-locked to the server.
- Client license keys, which specify the number and/or size of client-only nodes that may join the cluster. See "Server-Side Client License Keys" on page 77.

No license keys are required on the client-only nodes themselves.

Note: Other CXFS-aware products also require license keys:

- XVM cluster mirroring requires a license key on server-capable nodes in order for cluster nodes to access the cluster mirror. On CXFS client-only nodes, the user feature where applicable is honored after the `cxfs_client` service is started. XVM cluster mirroring on clients is also honored if it is enabled on the server. All CXFS client nodes need an appropriate mirror license key in order to access local mirrors.
 - Guaranteed rate I/O version 2 (GRIOv2) requires a license key on the server-capable nodes
-

Server-Side Client License Keys

There are two classes of server-side client license keys:

- *Workstation client license keys* specify the number of nodes with as many as 4 CPUs running one of the following platforms:

Linux on i386 architecture (as reported by the `uname -i` command)
Mac OS X
Windows

For example, an 8-node workstation client license key will allow up to eight nodes running any combination of the supported workstation platforms to join CXFS membership.

On Monday, you could have eight Mac OS X 4-CPU nodes, on Tuesday you could have four Mac OS X 4-CPU nodes and four Windows 4-CPU nodes.

- *Enterprise client license keys* specify the total number of CPUs running one of the following platforms:

AIX
IRIX
Linux on i386 architecture (more than 4 CPUs)
Linux on x86_64 or ia64 architecture
Mac OS X (more than 4 CPUs)
SGI ProPack 5 (CPU count is the number of sockets, not cores)
Solaris
Windows (more than 4 CPUs)

For example, a 32-CPU enterprise license key will allow sixteen 2-CPU nodes, eight 4-CPU nodes, or one 32-CPU node to join membership. If your cluster contained an SGI ProPack node with 4 sockets (4 dual-core CPUs), it would use 4 of the licenses.

License Key Replication on Server Nodes

The purchase of a workstation or enterprise license entitles you to generate a license key on each server-capable administration node in the cluster. Every server-capable administration node in the cluster should install a set of client license keys. A server will generate warnings in the system log if the license keys on one server-capable node are not equivalent to other server-capable nodes in the cluster.

Note: Server-side licensing does not introduce a point-of-failure in the CXFS cluster. If the metadata server fails and the cluster recovers to a backup server that has fewer/smaller client license keys, the client-only nodes that are currently in the cluster membership will remain in the membership. However, additional client-only nodes that attempt to join membership will fail until the membership count is reduced to below the license key entitlement on the active metadata server.

Cumulative Client License Keys

The number of client license keys is cumulative. To add more client-only nodes, you can purchase additional workstation or enterprise licenses as appropriate (you do not have to upgrade existing license keys).

For example, if you already have a 32-CPU enterprise license key and want to add another 32-CPU enterprise-class machine, you purchase another 32-CPU enterprise license. You must install this new license key key on every server-capable administration node in the cluster.

Examples of License Keys Required for Cluster Changes

The following figures show examples of the license keys that are required for cluster changes.

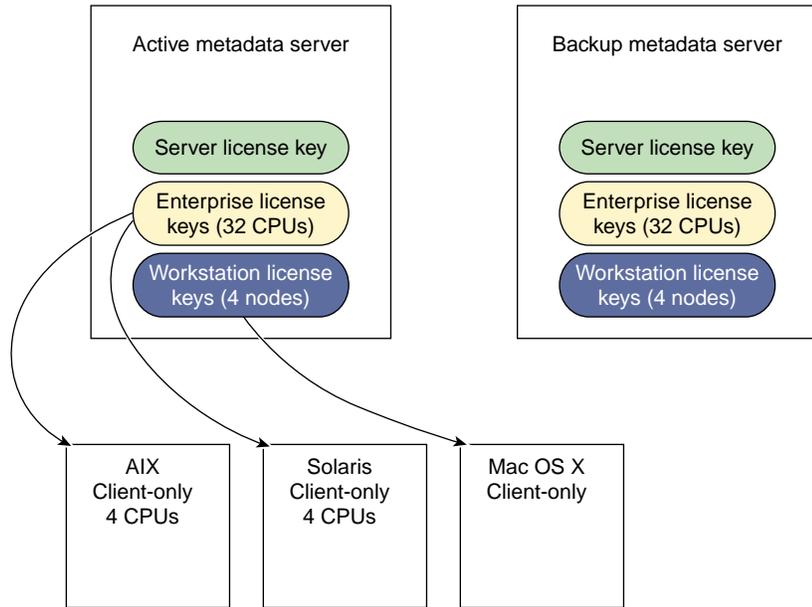


Figure 4-1 Server-Side License Keys

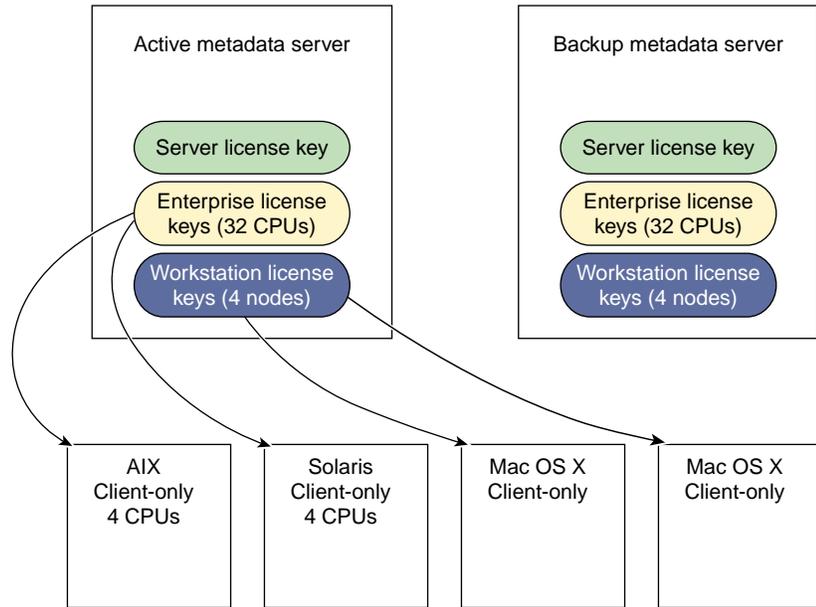


Figure 4-2 Server-Side License Keys: Adding a New Client-Only Node

The following table further illustrates the progressive changes in license keys required by a cluster as nodes are added and removed.

Table 4-1 Examples of License Keys Required for Cluster Changes

Action	Resulting Configuration	Licensing
Initial configuration	1 x 4-CPU server-capable node 4 x 2-CPU Windows clients	Purchase one 8-CPU CXFS_SS server license key and one 5-node workstation license key. Generate the 5-node workstation license key (CXFS_SS_CLIENT_WRK) for the server-capable node (the extra license is for future expansion).
Add a 2-CPU x86 (32-bit) Linux client	1 x 4-CPU server-capable node 4 x 2-CPU Windows clients 1 x 2-CPU x86 (32-bit) Linux client	No change, the 5-node workstation license key is now fully utilized.
Add an 8-CPU SGI ProPack 4 client	1 x 4-CPU sever-capable node 4 x 2-CPU Windows clients 1 x 2-CPU x86 (32-bit) Linux client 1 x 8-CPU SGI ProPack client	Purchase an 8-CPU enterprise license key. Generate the 8-CPU enterprise license key (CXFS_SS_CLIENT_ENT) for the server-capable node.
Add another 4-CPU server-capable node	2 x 4-CPU server-capable node 4 x 2-CPU Windows clients 1 x 2-CPU x86 (32-bit) Linux client 1 x 4-CPU SGI ProPack client	Purchase another 4-CPU server license. Generate both workstation and enterprise client license keys for the new server-capable node from the original license keys.
Add an 8-CPU Solaris node	2 x 4-CPU server-capable nodes 4 x 2-CPU Windows clients 1 x 2-CPU x86 (32-bit) Linux client 1 x 4-CPU SGI ProPack client 1 x 8-CPU Solaris client	Purchase an 8-CPU enterprise client license key. Generate the 8-CPU enterprise license key on each server-capable node.
Add a 4-CPU Mac OS X client for occasional use in the cluster	2 x 4-CPU server-capable nodes 4 x 2-CPU Windows clients 1 x 2-CPU x86 (32-bit) Linux client 1 x 4-CPU SGI ProPack client 1 x 4-CPU Mac OS X client	No change if one of the other workstation-class clients is dropped out of the cluster when the Mac OS X client is required.

Licensing on IRIX Clusters

This section discusses licensing on clusters with IRIX server-capable nodes:

- "Gathering the Host Information for IRIX" on page 82
- "Obtaining the License Keys for IRIX" on page 82
- "Installing the License Keys on IRIX" on page 83

Gathering the Host Information for IRIX

When you order CXFS, you will receive an entitlement ID. You must submit the system host ID, host name, and entitlement ID when requesting your permanent CXFS license key.

To obtain the host information for an IRIX server-capable node, execute the following command, execute the following FLEXlm command:

```
/usr/sbin/lmhostid
```

This command displays the FLEXlm host identifier, as shown in the following example:

```
% /usr/sbin/lmhostid
lmhostid - Copyright (c) 1989-2004 by Macrovision Corporation. All rights reserved.
The FLEXlm host ID of this machine is "690c9f5c"
```

When you are asked for the license manager host identifier, provide this FLEXlm host identifier.

For more information, see the *FLEXlm End User Manual*, which provides detailed information on using and administering software licenses. It is included in the IRIX system software documentation; the online version is in the subsystem `license_eoe.books.FlexLM_UG`.

Obtaining the License Keys for IRIX

To obtain your CXFS and XVM license keys, see information provided in your customer letter and the following web page:

<http://www.sgi.com/support/licensing>

Installing the License Keys on IRIX

For IRIX server-capable nodes, you will install the license keys in the following location:

```
/var/flexlm/license.dat
```

Do the following:

1. Create the `/var/flexlm` license key directory if necessary.

For example:

```
# mkdir -p /var/flexlm
```

2. Copy the key to the `license.dat` file.

Licensing on SGI ProPack 5 Clusters

This section discusses licensing on clusters with SGI ProPack 5 server-capable nodes:

- "Gathering the Host Information for SGI ProPack 5" on page 83
- "Obtaining the License Keys from SGI for SGI ProPack 5" on page 84
- "Installing the License Keys on SGI ProPack 5" on page 84

Gathering the Host Information for SGI ProPack 5

When you order CXFS, you will receive an entitlement ID. You must submit the system host ID, host name, and entitlement ID when requesting your permanent CXFS license key.

To obtain the host information for an SGI ProPack 5 server-capable node, execute the following command (assuming that the LK rpm from SGI ProPack has been installed):

```
/usr/sbin/lk_hostid
```

For example, the following shows that the serial number is N0000302 and the license ID is e000012e:

```
[root@linux64 root]# /usr/sbin/lk_hostid
```

```
N0000302 e000012e socket=16 core=16 processor=16
```

Obtaining the License Keys from SGI for SGI ProPack 5

To obtain your CXFS and XVM license keys, see information provided in your customer letter and the following web page:

<http://www.sgi.com/support/licensing>

Installing the License Keys on SGI ProPack 5

For SGI ProPack 5 server-capable nodes, you will install the license keys in the following location:

```
/etc/lk/keys.dat
```

Do the following:

- Create the `/etc/lk` license key directory if necessary. For example:

```
[root@linux root]# mkdir -p /etc/lk
```

- Copy the keys to the `keys.dat` file.

Note: SGI ProPack 5 client-only nodes require server-side licensing; therefore, you do not install licenses on any SGI ProPack 5 client-only nodes.

Verifying the License Keys

On IRIX server-capable nodes, use `cxfslicense` to verify keys.

On SGI ProPack 5 server-capable nodes, use either `cxfslicense` or `lk_verify` to verify licenses.

Verifying the License Keys with `cxfslicense` on Server-Capable Nodes

To verify that the license keys have been installed properly, use the `cxfslicense -d` command after installing the CXFS software. Licensing errors will be reported to the `fs2d` log.

For example, on an IRIX server-capable node, `cxfslicense -d` will report:

```
irix# /usr/cluster/bin/cxfslicense -d
No license required for XVM user feature.
Found XLV license of version 2.0 for XVM.
Found XVM_PLEX_CLUSTER license of version 4.0 for XVM.

Found 8 CPU version 4.0 license for CXFS_SS
Server-side licensing is available

Found license for 5 of CXFS_SS_CLIENT_WRK 4.0 serial 130448
Found license for 10 of CXFS_SS_CLIENT_WRK 4.0 serial 130452
Found license for 256 of CXFS_SS_CLIENT_ENT 4.0 serial 130503
Found license for 256 of CXFS_SS_CLIENT_ENT 4.0 serial 130505
```

For example, on an SGI ProPack server-capable node, `cxfslicense -d` will report:

```
[root@linux64 root]# /usr/cluster/bin/cxfslicense -d
Found 1 XVM_STD_IPF license(s) of version 1.0 for XVM.
Found 1 XVM_PLEX_IPF license(s) of version 1.0 for XVM.
Found 1 XVM_PLEX_CLUSTER license(s) of version 4.0 for XVM.

License(s) found: 1
Found 'CPU 8' serial 130441
Found 8 CPU version 4.0 license for CXFS_SS
Server-side licensing is available

License(s) found: 3
Found license for 1 CXFS_SS_CLIENT_WRK of 4.0 serial 130446
Found license for 5 CXFS_SS_CLIENT_WRK of 4.0 serial 130450
Found license for 10 CXFS_SS_CLIENT_WRK of 4.0 serial 130455
License(s) found: 4
Found license for 256 CXFS_SS_CLIENT_ENT of 4.0 serial 130459
Found license for 256 CXFS_SS_CLIENT_ENT of 4.0 serial 130507
Found license for 256 CXFS_SS_CLIENT_ENT of 4.0 serial 130509
Found license for 256 CXFS_SS_CLIENT_ENT of 4.0 serial 130511
```

If no valid license is found for SGI ProPack, `cxfslicense -d` will report:

```
[root@linux64 root]# /usr/cluster/bin/cxfslicense -d
Didn't find XVM_STD_IPF license of version 1.0 for XVM
Didn't find XVM_PLEX_IPF license(s) of version 1.0 for XVM.
Didn't find XVM_PLEX_CLUSTER license of version 4.0 for XVM.
```

Cannot find valid version 4.0 license for CXFS_SS

No CXFS server-side license, any server-side client licenses will be ignored.

No licenses available for CXFS_SS_CLIENT_WRK 4.0.
No licenses available for CXFS_SS_CLIENT_ENT 4.0.

Error: No valid CXFS licenses found for this server.

On an IRIX server-capable node with a workstation license key but no enterprise license key:

```
irix# /usr/cluster/bin/cxfslicense -d
No license required for XVM user feature.
Found XLV license of version 2.0 for XVM.
Found XVM_PLEX_CLUSTER license of version 4.0 for XVM.
```

Found 8 CPU version 4.0 license for CXFS_SS
Server-side licensing is available

Found license for 10 of CXFS_SS_CLIENT_WRK 4.0 serial 091119
No licenses available for CXFS_SS_CLIENT_ENT 4.0

On an IRIX server-capable node with no workstation or enterprise license keys:

```
irix# /usr/cluster/bin/cxfslicense -d
No license required for XVM user feature.
Found XLV license of version 2.0 for XVM.
Found XVM_PLEX_CLUSTER license of version 4.0 for XVM.

Found 8 CPU version 4.0 license for CXFS_SS
Server-side licensing is available

No licenses available for CXFS_SS_CLIENT_WRK 4.0
No licenses available for CXFS_SS_CLIENT_ENT 4.0

No client licenses for server-side licensing are available,
CXFS clients will need a client-side license.
```

On an SGI ProPack server-capable node with no workstation or enterprise license keys:

```
[root@linux64 root]# /usr/cluster/bin/cxfslicense -d
Found 1 XVM_STD_IPF license(s) of version 1.0 for XVM.
Found 1 XVM_PLEX_IPF license(s) of version 1.0 for XVM.
Found 1 XVM_PLEX_CLUSTER license(s) of version 4.0 for XVM.

Found 8 CPU version 4.0 license for CXFS_SS
Server-side licensing is available

No licenses available for CXFS_SS_CLIENT_WRK 4.0
No licenses available for CXFS_SS_CLIENT_ENT 4.0

No client licenses for server-side licensing are available,
CXFS clients will need a client-side license.
```

If you do not have the CXFS license key properly installed, you will see the following error on the console when trying to run CXFS:

```
Starting CXFS services> ...
CXFS not properly licensed for this host. Run
"/usr/cluster/bin/cxfslicense -d"
for detailed failure information. After fixing the
license, please run "/usr/cluster/bin/cxfs_cluster restart".
```

An error such as the following example will appear in the SYSLOG file (line breaks added here for readability):

```
Jan 25 10:24:03 ncc1701:Jan 25 10:24:03 cxfs_client:
cis_main FATAL: cxfs_client failed the CXFS license check.
Use the cxfslicense command to diagnose the license problem
```

The following will appear in the client-log file:

- Successful:
 - Server license key granted, regardless of local client license key:
Server-side license granted
- Unsuccessful (CXFS will not start):
 - Server denies a license key, regardless of local license key presence:
A server-side license could not be granted

On an administration node, the error will appear in the clconfd log.

The `cxfs_admin status` command displays the number of server-side license keys that have been issued to clients. See "Check Cluster/Node/Filesystem/Switch Status with `cxfs_admin`" on page 404.

Verifying the LK License Keys with `lk_verify`

On SGI ProPack 5 nodes, you can use the `lk_verify -A` command to verify LK licenses. To see more output, use the `v` option. For example:

```
# lk_verify -A -vvv
lk_check          : count=8

1 /etc/lk/keys.dat:005      product=CXFS_SS, version=4.000, count=0, begDate=1165350795, \
  expDate=1173160799, licenseID=23d5fd92, key=8jC8S9n6SmX2fhQA8KkYR6CZkaETLhLo, \
  info='CXFS SVR 8 CPU',attr='CPU 8', vendor='Silicon Graphics, Inc.', \
  ref_id='123651'
  Verdict:                SUCCESS. Nodelock.
                           Available since 8 days on 05-Dec-2006 14:33:15.
                           Will expire in 81 days on 05-Mar-2007 23:59:59

Attribute 1 of 4 : info=CXFS SVR 8 CPU
```

```
Attribute 2 of 4 : attr=CPU 8
Attribute 3 of 4 : vendor=Silicon Graphics, Inc.
Attribute 4 of 4 : ref_id=123651
```

```
2 /etc/lk/keys.dat:011      product=CXFS_SS_CLIENT_WRK, version=4.000, count=0, begDate=1165350798, \
expDate=1173160799, licenseID=23d5fd92, key=iRG6qLefp/9IdKSw/CHzuvastr4He09o, \
info='CXFS WRK 5 NODE',attr='NODE 5', vendor='Silicon Graphics, Inc.', \
ref_id='123653'
Verdict:                  SUCCESS. Nodelock.
                          Available since 8 days on 05-Dec-2006 14:33:18.
                          Will expire in 81 days on 05-Mar-2007 23:59:59
```

```
Attribute 1 of 4 : info=CXFS WRK 5 NODE
Attribute 2 of 4 : attr=NODE 5
Attribute 3 of 4 : vendor=Silicon Graphics, Inc.
Attribute 4 of 4 : ref_id=123653
```

```
3 /etc/lk/keys.dat:017      product=CXFS_SS_CLIENT_WRK, version=4.000, count=0, begDate=1165350801, \
expDate=1173160799, licenseID=23d5fd92, key=yoxsppVTgJm4QDWyDE+ZPmzRKAbXjgqF, \
info='CXFS WRK 10 NODE',attr='NODE 10', vendor='Silicon Graphics, Inc.', \
ref_id='123656'
Verdict:                  SUCCESS. Nodelock.
                          Available since 8 days on 05-Dec-2006 14:33:21.
                          Will expire in 81 days on 05-Mar-2007 23:59:59
```

```
Attribute 1 of 4 : info=CXFS WRK 10 NODE
Attribute 2 of 4 : attr=NODE 10
Attribute 3 of 4 : vendor=Silicon Graphics, Inc.
Attribute 4 of 4 : ref_id=123656
```

```
4 /etc/lk/keys.dat:023      product=CXFS_SS_CLIENT_ENT, version=4.000, count=0, begDate=1165350807, \
expDate=1173160799, licenseID=23d5fd92, key=bUgMNwrFPFUMWjfyXCXKWGl3/07lgef, \
info='CXFS ENT 2 CPU',attr='CPU 2', vendor='Silicon Graphics, Inc.', \
ref_id='123659'
Verdict:                  SUCCESS. Nodelock.
                          Available since 8 days on 05-Dec-2006 14:33:27.
                          Will expire in 81 days on 05-Mar-2007 23:59:59
```

4: CXFS License Keys

Attribute 1 of 4 : info=CXFS ENT 2 CPU
Attribute 2 of 4 : attr=CPU 2
Attribute 3 of 4 : vendor=Silicon Graphics, Inc.
Attribute 4 of 4 : ref_id=123659

5 /etc/lk/keys.dat:029 product=CXFS_SS_CLIENT_ENT, version=4.000, count=0, begDate=1165350811, \
expDate=1173160799, licenseID=23d5fd92, key=Vf3xpUzAlsYcPRnidJUCCLiS0OGhQ5wB, \
info='CXFS ENT 8 CPU',attr='CPU 8', vendor='Silicon Graphics, Inc.', \
ref_id='123705'

Verdict: SUCCESS. Nodelock.
Available since 8 days on 05-Dec-2006 14:33:31.
Will expire in 81 days on 05-Mar-2007 23:59:59

Attribute 1 of 4 : info=CXFS ENT 8 CPU
Attribute 2 of 4 : attr=CPU 8
Attribute 3 of 4 : vendor=Silicon Graphics, Inc.
Attribute 4 of 4 : ref_id=123705

6 /etc/lk/keys.dat:035 product=XVM_PLEX_CLUSTER, version=4.000, count=0, begDate=1165350536, \
expDate=1173160799, licenseID=23d5fd92, key=WqAdcLRHOBHGnajibTk6V4tf6UAibn9Iw, \
info='XVM PLEX CLUSTER 64',attr='NODE 64', vendor='Silicon Graphics, Inc.', \
ref_id='624457'

Verdict: SUCCESS. Nodelock.
Available since 8 days on 05-Dec-2006 14:28:56.
Will expire in 81 days on 05-Mar-2007 23:59:59

Attribute 1 of 4 : info=XVM PLEX CLUSTER 64
Attribute 2 of 4 : attr=NODE 64
Attribute 3 of 4 : vendor=Silicon Graphics, Inc.
Attribute 4 of 4 : ref_id=624457

7 /etc/lk/keys.dat:041 product=XVM_PLEX_IPF, version=1.000, count=0, begDate=1165350591, \
expDate=1173160799, licenseID=23d5fd92, key=jHgCqnM0ZBXgLqIkUhPcGwOYsnWTFaZN, \
info='XVM PLEX ALTIX',attr='4 CPU', vendor='Silicon Graphics, Inc.', \
ref_id='624430'

Verdict: SUCCESS. Nodelock.
Available since 8 days on 05-Dec-2006 14:29:51.
Will expire in 81 days on 05-Mar-2007 23:59:59

```

Attribute 1 of 4 : info=XVM PLEX ALTIX
Attribute 2 of 4 : attr=4 CPU
Attribute 3 of 4 : vendor=Silicon Graphics, Inc.
Attribute 4 of 4 : ref_id=624430

```

```

8 /etc/lk/keys.dat:047      product=XVM_STD_IPF, version=1.000, count=0, begDate=1165350618, \
  expDate=1173160799, licenseID=23d5fd92, key=GC4Wau6X1Etd38GxoYDHE3m4y3tX0iVV, \
  info='XVM Standalone ALTIX',attr='4 CPU', vendor='Silicon Graphics, Inc.', \
  ref_id='624429'
  Verdict:                SUCCESS. Nodelock.
                          Available since 8 days on 05-Dec-2006 14:30:18.
                          Will expire in 81 days on 05-Mar-2007 23:59:59

```

```

Attribute 1 of 4 : info=XVM Standalone ALTIX
Attribute 2 of 4 : attr=4 CPU
Attribute 3 of 4 : vendor=Silicon Graphics, Inc.
Attribute 4 of 4 : ref_id=624429

```

Verifying the XVM Mirror Licenses on Client-Only Nodes

To view the current status of XVM mirror licenses on client-only nodes, use the following command and search for the line containing the keyword mirrors:

```
xvm show -subsystem
```

For example:

```

# xvm show -subsystem
XVM Subsystem Information:
-----
apivers:                26
config gen:             33
privileged:             1
clustered:              1
cluster initialized:    1
user license enabled:   1
local mirrors enabled:  1
cluster mirrors enabled: 1
snapshot enabled:       1

```

```
snapshot max blocks:    -1
snapshot blocks used:   0
```

For More Information About Licensing

To request software keys or information about software licensing, see the following web page:

<http://www.sgi.com/support/licensing>

If you do not have access to the web, please contact your local Customer Support Center.

For more information about installing IRIX FLEXlm software license keys, see the *IRIX 6.5 Installation Instructions* booklet.

For more information on FLEXlm, you may order the *Flexible License Manager End User Manual* from Macrovision Corporation.

Preinstallation Steps

When you install the CXFS software, you must modify certain system files. The network configuration is critical. Each node in the cluster must be able to communicate with every other node in the cluster by both logical name and IP address without going through any other network routing; proper name resolution is key. SGI recommends static routing.

This section provides an overview of the steps that you should perform on your nodes prior to installing the CXFS software. It contains the following sections:

- "Hostname Resolution and Network Configuration Rules" on page 93
- "Configuring Network Interfaces" on page 94
- "Verifying the Private and Public Networks" on page 98
- "Configuring the Serial Ports for IRIX Administration Nodes" on page 99

Hostname Resolution and Network Configuration Rules



Caution: It is critical that you understand these rules before attempting to configure a CXFS cluster.

Use the following hostname resolution rules and recommendations when defining a node:

- The first node you define in the pool must be an administration node.
- Hostnames cannot begin with an underscore (_) or include any white-space characters.
- The private network IP addresses on a running node in the cluster cannot be changed while CXFS services are active.
- You must be able to communicate directly between every node in the cluster (including client-only nodes) using IP addresses and logical names, without routing.

- A private network must be dedicated to be the heartbeat and control network. No other load is supported on this network.
- The heartbeat and control network must be connected to all nodes, and all nodes must be configured to use the same subnet.

If you change hostname resolution settings in the `/etc/nsswitch.conf` file after you have defined the first administration node (which creates the cluster database), you must re-create the cluster database.

Configuring Network Interfaces

When configuring your network, remember the following:

- You must be able to communicate between every node in the cluster directly using IP address and logical name, without routing.
- You must dedicate a private network to be your heartbeat and control network. No other load is supported on this network.
- The heartbeat and control network must be connected to all nodes, and all nodes must be configured to use the same subnet for that network.
- Because CXFS heartbeat is done using IP multicast, the private network must be multicast-capable. This means that all of the interfaces must have multicast enabled (which is the default) and all of the external networking hardware (such as switches)must support IP multicast.

Adding a Private Network

The following procedure provides an overview of the steps required to add a private network.

Note: A private network is required for use with CXFS.

You may skip some steps, depending upon the starting conditions at your site.

1. Edit the `/etc/hosts` file so that it contains entries for every node in the cluster and their private interfaces as well.

The `/etc/hosts` file has the following format, where *primary_hostname* can be the simple hostname or the fully qualified domain name:

```
IP_address    primary_hostname    aliases
```

You should be consistent when using fully qualified domain names in the `/etc/hosts` file. If you use fully qualified domain names on a particular node, then all of the nodes in the cluster should use the fully qualified name of that node when defining the IP/hostname information for that node in their `/etc/hosts` file.

The decision to use fully qualified domain names is usually a matter of how the clients are going to resolve names for their client/server programs (such as NFS), how their default resolution is done, and so on.

Even if you are using the domain name service (DNS) or the network information service (NIS), you must add every IP address and hostname for the nodes to `/etc/hosts` on all nodes. For example:

```
190.0.2.1 server1-example.com server1
190.0.2.3 stocks
190.0.3.1 priv-server1
190.0.2.2 server2-example.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.

For more information, see the `hosts` and `resolve.conf` man pages.

Note: Exclusive use of NIS or DNS for IP address lookup for the nodes will reduce availability in situations where the NIS or DNS service becomes unreliable.

2. Edit the `/etc/nsswitch.conf` file so that local files are accessed before either NIS or DNS. That is, the `hosts` line in `/etc/nsswitch.conf` must list `files` first.

For example:

```
hosts:        files nis dns
```

(The order of `nis` and `dns` is not significant to CXFS, but `files` must be first.)

3. Configure your private interface according to the instructions in the Network Configuration section of your Linux distribution manual. To verify that the private interface is operational, use the `ifconfig -a` command. For example:

```
[root@linux64 root]# ifconfig -a

eth0      Link encap:Ethernet  HWaddr 00:50:81:A4:75:6A
          inet addr:192.168.1.1  Bcast:192.168.1.255  Mask:255.255.255.0
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:13782788  errors:0  dropped:0  overruns:0  frame:0
          TX packets:60846  errors:0  dropped:0  overruns:0  carrier:0
          collisions:0  txqueuelen:100
          RX bytes:826016878 (787.7 Mb)  TX bytes:5745933 (5.4 Mb)
          Interrupt:19  Base address:0xb880  Memory:fe0fe000-fe0fe038

eth1      Link encap:Ethernet  HWaddr 00:81:8A:10:5C:34
          inet addr:10.0.0.10  Bcast:10.0.0.255  Mask:255.255.255.0
          UP BROADCAST MULTICAST  MTU:1500  Metric:1
          RX packets:0  errors:0  dropped:0  overruns:0  frame:0
          TX packets:0  errors:0  dropped:0  overruns:0  carrier:0
          collisions:0  txqueuelen:100
          RX bytes:0 (0.0 b)  TX bytes:0 (0.0 b)
          Interrupt:19  Base address:0xef00  Memory:febfd000-febfd038

lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          UP LOOPBACK RUNNING  MTU:16436  Metric:1
          RX packets:162  errors:0  dropped:0  overruns:0  frame:0
          TX packets:162  errors:0  dropped:0  overruns:0  carrier:0
          collisions:0  txqueuelen:0
          RX bytes:11692 (11.4 Kb)  TX bytes:11692 (11.4 Kb)
```

This example shows that two Ethernet interfaces, `eth0` and `eth1`, are present and running (as indicated by `UP` in the third line of each interface description).

If the second network does not appear, it may be that a network interface card must be installed in order to provide a second network, or it may be that the network is not yet initialized.

4. (Optional) Make the modifications required to use CXFS connectivity diagnostics. See "IRIX Modifications for CXFS Connectivity Diagnostics" on page 111, and "SGI ProPack Modifications for CXFS Connectivity Diagnostics" on page 124.

Configuring IRIX Interfaces

To configure IRIX network interfaces, do the following:

1. Ensure that name services are available. See step 1 in "Adding a Private Network" on page 94.
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` in `/etc/config/netif.options` to the number of interfaces. For fewer than eight interfaces, the line is as follows:

```
if_num=8
```

4. Repeat steps 1 through 3 for the other nodes.
5. Edit the `/etc/config/routed.options` file on each IRIX node so that the routes are not advertised over the control network. 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 the advertising of routes
- 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)

Verifying the Private and Public Networks

For each private network on each node in the pool, verify access with the ping command. Enter the following, where *nodeIPaddress* is the IP address of the node:

```
ping nodeIPaddress
```

For example:

```
[root@linux64 root]# ping 10.0.0.1
PING 10.0.0.1 (10.0.0.1) from 128.162.240.141 : 56(84) bytes of data.
64 bytes from 10.0.0.1: icmp_seq=1 ttl=64 time=0.310 ms
64 bytes from 10.0.0.1: icmp_seq=2 ttl=64 time=0.122 ms
64 bytes from 10.0.0.1: icmp_seq=3 ttl=64 time=0.127 ms
```

Also execute a ping on the public networks. If ping fails, follow these steps:

1. Verify that the network interface was configured up using `ifconfig`. For example:

```
[root@linux64 root]# ifconfig eth1
eth1      Link encap:Ethernet  HWaddr 00:81:8A:10:5C:34
          inet addr:10.0.0.10  Bcast:10.0.0.255  Mask:255.255.255.0
          UP BROADCAST MULTICAST  MTU:1500  Metric:1
          RX packets:0 errors:0 dropped:0 overruns:0 frame:0
          TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:100
          RX bytes:0 (0.0 b)  TX bytes:0 (0.0 b)
          Interrupt:19 Base address:0xef00 Memory:febfd000-febfd038
```

In the third output line above, UP indicates that the interface was configured up.

2. Verify that the cables are correctly seated.

Repeat this procedure on each node.

Configuring the Serial Ports for IRIX Administration Nodes

If one IRIX administration node is configured to reset another IRIX administration node, you must turn off the `getty` process for the tty ports to which the reset serial cables are connected. You must do this on the IRIX administration node performing the reset (not the node receiving the reset). To do this, perform the following steps on each IRIX administration node; if you have a cluster with nodes running other operating systems, see the *CXFS MultiOS Client-Only Guide for SGI InfiniteStorage*.

1. Determine which port is used for the reset line. `ttyd2` is the most commonly used port, except on Origin 300 and Origin 350 system, where `ttyd4` is commonly used.
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, for an Origin 3000:

```
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
```


IRIX CXFS Installation



Caution: CXFS is a complex product. To ensure that it is installed and configured in an optimal manner, you **must** purchase initial setup services from SGI. This chapter is not intended to be used directly by the customer, but is provided for reference. You should read through the following chapters, before attempting to install and configure a CXFS cluster:

- Chapter 1, "Introduction to CXFS" on page 1
- Chapter 2, "SGI RAID for CXFS Clusters" on page 63
- Chapter 3, "Switches" on page 67
- Chapter 4, "CXFS License Keys" on page 75
- Chapter 5, "Preinstallation Steps" on page 93
- Chapter 6, "IRIX CXFS Installation" on page 101 (this chapter)
- Chapter 8, "Postinstallation Steps" on page 127
- Chapter 9, "Best Practices" on page 139
- Chapter 10, "Initial Setup of the Cluster" on page 163

Also see the *CXFS MultiOS Client-Only Guide for SGI InfiniteStorage*. If you are using coexecution with FailSafe, also see the *FailSafe Administrator's Guide for SGI InfiniteStorage*.

On IRIX nodes, CXFS supports either an *administration node* containing the cluster administration daemons (`fs2d`, `crsd`, `cad`, and `cmond`), the CXFS control daemon (`clconfd`), and the cluster database or a *client-only node* containing the `cxfs_client` daemon. The software you install on a node determines the node type.

Nodes that you intend to run as metadata servers must be installed as administration nodes; all other nodes should be client-only nodes.

This chapter discusses the following:

- "IRIX Administration Software Installation" on page 102
 - "IRIX Client-only Software Installation" on page 107
 - "IRIX Modifications for CXFS Connectivity Diagnostics" on page 111
-

Note: CXFS does not support a miniroot installation.

IRIX Administration Software Installation

Only those nodes that are potential CXFS metadata servers should be installed as CXFS administration nodes. All other nodes should be client-only nodes.

Note: An IRIX node can be either be a CXFS administration node (for which you install `cluster_admin`) or a client-only node (for which you install `cxfs_client`). You cannot install both `cluster_admin` and `cxfs_client` on the same node. This procedure installs an administration node; to install a client-only node, see "IRIX Client-only Software Installation" on page 107.

Installing the CXFS software for a CXFS administration node requires approximately 30.3 MB of space.

Note: You cannot combine the IRIX operating system installation and the CXFS installation. You must install the operating system first.

To install the required IRIX software for a CXFS administration node, do the following on each administration node:

1. Upgrade to IRIX 6.5.x according to the *IRIX 6.5 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 server) Install the version of the serial port server driver that is appropriate to the operating system. Use the CD that accompanies the serial port server. Reboot the system after installation.

For more information, see the documentation provided with the serial port server.

3. Insert the *CXFS MultiOS Server 4.2* CD into the CD drive.
4. Read the release notes for the CXFS IRIX platform to learn about any late-breaking changes in the installation procedure. CXFS release notes have the following chapters:

1	Introduction
2	Installation Information
3	Changes and Additions
4	Bug Fixes
5	Known Problems and Workarounds
6	Documentation Errors
7	Activating Your CXFS <i>x.x</i> and Cluster XVM for 6.5.x License With FLEXlm

You can view the release notes as follows:

- To view the release notes before they are installed, choose the following from the desktop Toolchest to bring up the **Software Manager** window:

```
System
  > Software Manager
```

Choose **Customize Installation** by typing `/CDROM/dist` into the **Available Software** box. A list of products available for installation will come up. If the product name is highlighted (similar to an HTML link), then there are release notes available. Click on the link to bring up the **Release Notes** window.

- If you do not have access to a graphics terminal, you must install the release notes and then use the `relnotes` command to view the CXFS release notes. For example:

```
# inst
...
Inst> from /CDROM/dist
Inst> keep *
```

```
Inst> install cxfs.man.relnotes
Inst> go
...
Inst> quit
```

View the release notes from the current login session or another session:

```
# /usr/sbin/relnotes cxfs ChapterNumber
```

5. Insert *IRIX CD-ROM #1* into the CD drive.
6. Start up `inst` and instruct `inst` to read the already inserted CD-ROM:

```
# inst
...
Inst> open /CDROM/dist
```



Caution: Do not install to an alternate root using the `inst -r` option. Some of the exit operations (exitops) do not use pathnames relative to the alternate root, which can result in problems on both the main and alternate root filesystem if you use the `-r` option. For more information, see the `inst` man page.

7. Install the XVM `oe.books.xvm` subsystem:

```
Inst> keep *
Inst> install oe.books.xvm
Inst> go
...
Inst> quit
```

8. (Optional) If you want to use Performance Co-Pilot to run XVM statistics, install the default `pcp_oe` subsystems. This installs the Performance Co-Pilot PMDA (the agent to export XVM statistics) as an exit operation (exitop).

```
Inst> keep *
Inst> install pcp_oe default
Inst> go
...
Inst> quit
```

9. Insert *IRIX CD-ROM #3* into the CD drive.

10. Start up `inst` and instruct `inst` to read the CD:

```
irix# inst
...
Inst> open /CDROM/dist
```

11. Insert the *CXFS MultiOS Server 4.2* CD into the CD drive.

Note: If you have a system running an earlier version of IRIX with CXFS installed and try to upgrade IRIX without also installing the required CXFS CD, you will get a conflict. You must either install the CXFS CD or remove CXFS.

12. Instruct `inst` to read the CD:

```
Inst> open /CDROM/dist
```

13. Choose the CXFS software to install:

```
Inst> keep *
Inst> install cxfs
Inst> keep cxfs.sw.grio2_cell
Inst> install cxfs_admin
Inst> install cxfs_cluster
Inst> install cxfs_util
Inst> install cluster_admin
Inst> install cluster_control
Inst> install cluster_services
Inst> install eoe.sw.xvm
Inst> install sysadm_base
Inst> install sysadm_cluster
Inst> install sysadm_cxfs
Inst> install sysadm_xvm
Inst> install applicablepatches
```

14. (Optional) To install guaranteed-rate I/O version 2 (GRIOv2) as well, choose the GRIOv2 software:

```
Inst> install eoe.sw.grio2
Inst> install cxfs.sw.grio2_cell
```

15. Install the chosen software:

```
Inst> go
...
Inst> quit
```

This installs the following packages:

```
cluster_admin.man.man
cluster_admin.sw.base
cluster_control.man.man
cluster_control.sw.base
cluster_control.sw.cli
cluster_services.man.man
cluster_services.sw.base
cluster_services.sw.cli
cxfs.books.CXFS_AG
cxfs.man.relnotes
cxfs.sw.cxfs
cxfs.sw.grio2_cell           (Optional)
cxfs.sw.xvm_cell
cxfs_admin.sw.base
cxfs_admin.sw.man
cxfs_cluster.man.man
cxfs_cluster.sw.base
cxfs_cluster.sw.cli
cxfs_util.man.man
cxfs_util.sw.base
eoe.sw.grio2                 (Optional)
eoe.sw.xvm
patch_cxfs.eoe_sw.base
patch_cxfs.eoe_sw64.lib
patch_sysadm_xvm.man.pages
patch_sysadm_xvm.man.relnotes
patch_sysadm_xvm.sw.client
patch_sysadm_xvm.sw.desktop
patch_sysadm_xvm.sw.server
patch_sysadm_xvm.sw.web
sysadm_base.man.priv
sysadm_base.man.relnotes
sysadm_base.man.server
sysadm_base.sw.client
```

```
sysadm_base.sw.dso
sysadm_base.sw.priv
sysadm_base.sw.server
sysadm_cluster.man.relnotes
sysadm_cluster.sw.client
sysadm_cluster.sw.server
sysadm_cxfs.man.pages
sysadm_cxfs.man.relnotes
sysadm_cxfs.sw.client
sysadm_cxfs.sw.desktop
sysadm_cxfs.sw.server
sysadm_cxfs.sw.web
sysadm_xvm.man.pages
sysadm_xvm.man.relnotes
sysadm_xvm.sw.client
sysadm_xvm.sw.desktop
sysadm_xvm.sw.server
sysadm_xvm.sw.web
```

The process may take a few minutes to complete.

16. Use the `cxfslicense -d` command to verify licenses. If there are errors, verify that you have obtained and installed the CXFS license keys. For more information, see "Verifying the License Keys with `cxfslicense` on Server-Capable Nodes" on page 84.
17. Reboot the system.

IRIX Client-only Software Installation

An IRIX node can be either be a CXFS administration node (for which you install `cluster_admin`) or a client-only node (for which you install `cxfs_client`). You cannot install both `cluster_admin` and `cxfs_client` on the same node. This procedure installs a client-only node; to install an administration node, see "IRIX Administration Software Installation" on page 102.

Note: You cannot combine the IRIX operating system installation and the CXFS installation. You must install the operating system first.

To install the required IRIX software, do the following on each IRIX client-only node:

1. Upgrade to IRIX 6.5.x according to the *IRIX 6.5 Installation Instructions*.

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

```
irix# uname -aR
```

2. (For sites with a serial port server) Install the version of the serial port server driver that is appropriate to the operating system. Use the CD that accompanies the serial port server. Reboot the system after installation.

For more information, see the documentation provided with the serial port server.

3. Insert the *CXFS MultiOS Client 4.2* CD into the CD drive.
4. Read the release notes for the CXFS IRIX platform to learn about any late-breaking changes in the installation procedure. CXFS release notes have the following chapters:

- 1 Introduction
- 2 Installation Information
- 3 Changes and Additions
- 4 Bug Fixes
- 5 Known Problems and Workarounds
- 6 Documentation Errors
- 7 Activating Your CXFS *x.x* and Cluster XVM for 6.5.x License With FLEXlm

You can view the release notes as follows:

- To view the release notes before they are installed, choose the following from the desktop Toolchest to bring up the **Software Manager** window:

```
System  
  > Software Manager
```

Choose **Customize Installation** by typing `/CDROM/dist` into the **Available Software** box. A list of products available for installation will come up. If the product name is highlighted (similar to an HTML link), then there are release notes available. Click on the link to bring up the **Release Notes** window.

- If you do not have access to a graphics terminal, you must install the release notes and then use the `relnotes` command to view the CXFS release notes. For example:

```

irix# inst
...
Inst> open /CDROM/dist
Inst> keep *
Inst> install cxfs.man.relnotes
Inst> go
...
Inst> quit

```

View the release notes from the current login session or another session:

```

irix# /usr/sbin/relnotes cxfs ChapterNumber

```

5. Insert *IRIX CD-ROM #1* into the CD drive.
6. Start up `inst` and instruct it to read the CD:

```

# inst
...
Inst> open /CDROM/dist

```



Caution: Do not install to an alternate root using the `inst -r` option. Some of the exit operations (exitops) do not use pathnames relative to the alternate root, which can result in problems on both the main and alternate root filesystem if you use the `-r` option. For more information, see the `inst` man page.

7. Install the XVM `coe.books.xvm` subsystem:

```

Inst> keep *
Inst> install coe.books.xvm
Inst> go
...
Inst> quit

```

8. (Optional) If you want to use Performance Co-Pilot to run XVM statistics, install the default `pcp_eoe` subsystems. This installs the Performance Co-Pilot PMDA (the agent to export XVM statistics) as an exit operation (`exitop`).

```
Inst> keep *
Inst> install pcp_eoe default
Inst> go
...
Inst> quit
```

9. Insert the *CXFS MultiOS Client 4.2* CD into the CD drive.
10. Start up `inst` and instruct it to read the CD:

```
irix# inst
...
Inst> open /CDROM/dist
```

If you have a system running an earlier version of IRIX with CXFS installed and try to upgrade IRIX without also installing the required CXFS CD, you will get a conflict. You must either install the CXFS CD or remove CXFS.



Caution: Do not install to an alternate root using the `inst -r` option. Some of the exit operations (`exitops`) do not use pathnames relative to the alternate root, which can result in problems on both the main and alternate root filesystem if you use the `-r` option. For more information, see the `inst` man page.

11. Choose the CXFS software to install:

```
Inst> keep *
Inst> install cxfs
Inst> keep cxfs.sw.grio2_cell
Inst> install cxfs_client
Inst> install cxfs_util
Inst> install eoe.sw.xvm
Inst> install applicablepatches
```



Caution: If you do not install `cxfs_client`, the `inst` utility will not detect a conflict, but the CXFS cluster will not work. You **must** install the `cxfs_client` subsystem.

12. (Optional) To install guaranteed-rate I/O version 2 (GRIOv2) as well, choose the GRIOv2 software:

```
Inst> install cxfs.sw.grio2_cell
Inst> install eoe.sw.grio2
```

13. Install the chosen software:

```
Inst> go
...
Inst> quit
```

This installs the following packages:

```
cxfs.books.CXFS_AG
cxfs.man.relnotes
cxfs.sw.cxfs
cxfs.sw.grio2_cell    (Optional)
cxfs.sw.xvm_cell
cxfs_client.man.man
cxfs_client.sw.base
cxfs_util.man.man
cxfs_util.sw.base
eoe.sw.grio2          (Optional)
eoe.sw.xvm
patch_cxfs.eoe_sw.base
patch_cxfs.eoe_sw64.lib
```

The process may take a few minutes to complete.

14. Reboot the system.

IRIX Modifications for CXFS Connectivity Diagnostics

If you want to use the connectivity diagnostics provided with CXFS, ensure that the `.rhosts` file on each administration node allows all the nodes in the cluster to have access to each other in order to run remote commands such as `rsh`. The connectivity tests execute a `ping` command from the local node to all nodes and from all nodes to the local node. To execute `ping` on a remote node, CXFS uses `rsh` (user `root`). For example, suppose you have a cluster with three nodes: `cxfs0`, `cxfs1`, and `cxfs2`.

The `.rhosts` file on each administration node will be as follows (prompt denotes node name):

```
cxfs0# cat /.rhosts
cxfs1 root
cxfs1-priv root
cxfs2 root
cxfs2-priv root
```

```
cxfs1# cat /.rhosts
cxfs0 root
cxfs0-priv root
cxfs2 root
cxfs2-priv root
```

```
cxfs2# cat /.rhosts
cxfs0 root
cxfs0-priv root
cxfs1 root
cxfs1-priv root
```

Make sure that the mode of the `.rhosts` file is set to 600 (read and write access for the owner only).

After you have completed running the connectivity tests, you may wish to disable `rsh` on all cluster nodes.

SGI ProPack CXFS Installation



Caution: CXFS is a complex product. To ensure that it is installed and configured in an optimal manner, you **must** purchase initial setup services from SGI. This chapter is not intended to be used directly by the customer, but is provided for reference. You should read through the following chapters, before attempting to install and configure a CXFS cluster:

- Chapter 1, "Introduction to CXFS" on page 1
- Chapter 2, "SGI RAID for CXFS Clusters" on page 63
- Chapter 3, "Switches" on page 67
- Chapter 4, "CXFS License Keys" on page 75
- Chapter 5, "Preinstallation Steps" on page 93
- Chapter 7, "SGI ProPack CXFS Installation" on page 113 (this chapter)
- Chapter 8, "Postinstallation Steps" on page 127
- Chapter 9, "Best Practices" on page 139
- Chapter 10, "Initial Setup of the Cluster" on page 163

If you are using a multiOS cluster, also see the *CXFS MultiOS Client-Only Guide for SGI InfiniteStorage*. If you are using coexecution with FailSafe, also see the *FailSafe Administrator's Guide for SGI InfiniteStorage*.

On SGI ProPack for Linux nodes, CXFS supports either an *administration node* containing the cluster administration daemons (`fs2d`, `crsd`, `cad`, and `cmond`), the CXFS control daemon (`clconfd`), and the cluster database or a *client-only node* containing the `cxfs_client` daemon. The software you install on a node determines the node type.

Note: SGI ProPack is an overlay product that adds or enhances features in the supported Linux base distributions.

Nodes that you intend to run as metadata servers must be installed as administration nodes; all other nodes should be client-only nodes.

This chapter discusses the following:

- "SGI ProPack Limitations and Considerations" on page 114
- "SGI ProPack Administration Software Installation" on page 116
- "SGI ProPack Client-Only Software Installation" on page 120
- "SGI ProPack Installation Verification" on page 124
- "SGI ProPack Modifications for CXFS Connectivity Diagnostics" on page 124

After completing these steps, see Chapter 10, "Initial Setup of the Cluster" on page 163. For details about specific configuration tasks, see Chapter 11, "Reference to GUI Tasks" on page 183.

SGI ProPack Limitations and Considerations

The following sections highlight limitations and considerations for SGI ProPack nodes.

See also Appendix E, "Filesystem Specifications" on page 537.

Limitations and Considerations for Any SGI ProPack Node

The following limitations and considerations apply to any SGI ProPack node (client-only or server-capable):

- By default, DMAPi is turned off on SGI ProPack 5 systems. When you install DMF on a server-capable node, it automatically enables DMAPi. However, if you want to mount filesystems on an SGI ProPack 5 client-only node with the `dmi` mount option, you must ensure that the `DMAPI_PROBE` system tunable parameter on the node is set to `yes` in the `/etc/sysconfig/sysctl` file. Changes to the file will be processed on the next reboot. After setting that system configuration file, you can immediately enable DMAPi by executing the following:

```
sysctl -w fs.xfs.probe_dmapi=1
```

If you run a DMAPi application other than DMF, you must also change parameter on the SGI ProPack 5 server-capable nodes.

- On SGI ProPack systems, the `mkfs.xfs` command does not discover log or realtime subvolumes. You must specify the log or realtime subvolumes on the command line. For more information, see the `mkfs.xfs(8)` man page.
- GPT partition tables, often created by operating system installers or the `parted` partitioning tool, store labels in two locations. If you reuse a disk that previously had a GPT label, you must be careful; using tools such as `fdisk` to repartition the drive will not eliminate the backup GPT label. When you reboot, EFI scans the disks before the operating system is started. It assumes any backup labels it finds are valid and restores them. This can corrupt or destroy filesystems. You can use the `parted` tool to detect this situation and fix it.

Note: The `parted` tool has a `mkpartsect` command that accepts start and end values for partitions being created in sectors rather than MB. For more information, see the *XVM Volume Manager Administrator's Guide* and http://support.sgi.com/content_request/838562/index.html on Supportfolio.

- CXFS filesystems with XFS version 1 directory format cannot be mounted on SGI ProPack nodes.
- Whenever you install a new kernel patch, you must also install the corresponding CXFS package. This is required because the kernel patch causes the kernel version number to be increased. Failure to install the corresponding CXFS package will result in the inability to run CXFS. To obtain the required CXFS package, see your SGI support contact.
- After upgrading CXFS, you should reboot the system in order to make the new updates to take effect. A reboot is not required if you are performing a fresh installation.
- The implementation of file creates using `O_EXCL` is not complete. Multiple applications running on the same node using `O_EXCL` creates as a synchronization mechanism will see the expected behavior (only one of the creates will succeed). However, applications running between nodes may not get the `O_EXCL` behavior they requested (creates of the same file from two or more separate nodes may all succeed).

Limitations and Considerations for SGI ProPack Client-Only Nodes

On systems running SUSE Linux Enterprise Server 10 (SLES 10) that are greater than 64 CPUs, there are issues with using the md driver and CXFS. The md driver holds the BKL (Big Kernel Lock), which is a single, system-wide spin lock. Attempting to acquire this lock can add substantial latency to a driver's operation, which in turn holds off other processes such as CXFS. The delay causes CXFS to lose membership. This problem has been observed specifically when an md pair RAID split is done, such as the following:

```
raidsetfaulty /dev/md1 /dev/path/to/partition
```

SGI ProPack Administration Software Installation

The CXFS software will be initially installed and configured by SGI personnel. This section provides an overview of those procedures.

Note: Version numbers shown here are examples; your installed system may differ.

A node that may be a CXFS metadata server must be installed as a CXFS administration node. All other nodes should be client-only nodes.

Installing the CXFS software for a CXFS administration node requires approximately 65 MB of space.

Do the following to install the software required for an SGI ProPack administration node:

1. Read the CXFS `README` file for the SGI ProPack platform to learn about any late-breaking changes in the installation procedure.
 2. Install the SGI ProPack release, according to the directions in the SGI ProPack documentation. Ensure that you select the `SGI Licensed` package group.
-

Note: When installing the Linux OS, disconnect the system from the fabric or ensure that the drive you are installing on is not a SAN-attached drive.

3. Install any required patches. See the SGI ProPack `releasenotes/README` file for more information.



Caution: You **must** update the operating system with all security fixes, bug fixes, and enhancements available from the operating system vendor.

4. Verify that the node is running the supported Linux distribution and SGI ProPack overlay, according to the CXFS for SGI ProPack release notes. See the `/etc/SuSE-release` and `/etc/sgi-release` files.
5. If you have previously installed XVM in standalone mode, remove any remaining `sgi-xvm-standalone` package. To find and remove the package:

```
[root@linux CXFS_CDROM]# rpm -e --allmatches `rpm -qa | grep xvm-standalone`
```

You may also need to remove weak-updates links from the `sgi-xvm-standalone` RPM. If you are running the 2.6.16.21-0.25 kernel, you would do the following:

```
[root@linux CXFS_CDROM]# rm -rf /lib/modules/2.6.16.21-0.25-default/weak-updates/os_lib
[root@linux CXFS_CDROM]# rm -rf /lib/modules/2.6.16.21-0.25-default/weak-updates/xvm
```

6. Insert and mount the *CXFS MultiOS Server 4.2* CD.
7. Install the enhanced XFS kernel module:

```
[root@linux CXFS_CDROM]# rpm -Uvh sgi-enhancedxfs-kmp-default-version-kernelrelease.architecture.rpm
Preparing...          ##### [100%]
 1:-sgi-enhancedxfs-kmp-de##### [100%]
```

Where:

- *kernelvariant* and *kernelrelease* are the kernel variant and release output by the `uname -r` command
- *version* is the version number
- *architecture* is the processor architecture type output by the `uname -i` command

Note: For SGI ProPack 5 running SLES 10, one version of CXFS may support one or more *kernelrelease* values. See the CXFS SGI ProPack release notes for the supported versions.

8. Install the CXFS server kernel module:

```
[root@linux cdrom]# rpm -Uvh sgi-cxfs-server-kmp-kernelvariant-kernelrelease-version.architecture.rpm
Preparing...                               ##### [100%]
 1:-sgi-cxfs-server-kmp-ker##### [100%]
```

9. Install the SGI XFS user-space `sgi-xfsdump` and `sgi-xfsprogs` RPMs (which replace the `xfsdump` and `xfsprogs` RPMs, respectively):

```
[root@linux cdrom]# rpm --nodeps -Uvh sgi-xfsdump* sgi-xfsprogs*
```

10. Install the CXFS application binaries, documentation, and support tools:

```
[root@linux cdrom]# rpm -Uvh cxfs-sysadm_base-lib* cxfs-xvm-cmds* cluster_admin* \
cluster_control* cluster_services* cxfs_util* cxfs_cluster* cxfs-doc*
Preparing...                               ##### [100%]
 1:cluster_admin                            ##### [ 13%]
cxfs_cluster                               0:off 1:off 2:off 3:on  4:off 5:on  6:off
cdb-exitop: initializing CDB
fs2d
cdbnew: Created standard CDB database in /var/cluster/cdb/cdb.db
cdb-exitop: success
 2:cxfs-sysadm_base-lib                    ##### [ 25%]
 3:cluster_control                         ##### [ 38%]
 4:cluster_services                        ##### [ 50%]
Started cluster control processes
cluster_control-exitop: success
cluster_ha-exitop: Added HA keys to /var/cluster/cdb/cdb.db
cluster_ha-exitop: success
 5:cxfs_util                               ##### [ 63%]
 6:cxfs-xvm-cmds                           ##### [ 75%]
boot.xvm                                   0:off 1:off 2:off 3:off 4:off 5:off 6:off
 7:cxfs_cluster                            ##### [ 88%]
cxfs                                        0:off 1:off 2:off 3:on  4:off 5:on  6:off
cluster_cx-exitop: Added CXFS keys to /var/cluster/cdb/cdb.db
cluster_cx-exitop: Added CXFS administration access keys to
/var/cluster/cdb/cdb.db

cluster_cx-exitop: success
 8:cxfs-doc                                ##### [100%]
boot.xvm                                   0:off 1:off 2:off 3:off 4:off 5:off 6:off
```

Note: If you have not yet installed the license key file, you may get a warning at this point.

11. (Optional) If you are using GRIO, install the `grio2-cmds` and `grio2-server` packages:

```
[root@linux cdrom]# rpm -Uvh grio2-cmds* grio2-server*
Preparing...                               ##### [100%]
 1:grio2-cmds                               ##### [ 50%]
grio2                                       0:off 1:off 2:off 3:on 4:off 5:on 6:off
 2:grio2-server                             ##### [100%]
```

12. Install the CXFS graphical user interface (GUI) and XVM GUI server and client packages:

```
[root@linux CXFS_CDROM]# rpm -Uvh cxfs-sysadm_base-client* cxfs-sysadm_base-server* \
cxfs-sysadm_xvm* cxfs-sysadm_c*
Preparing...                               ##### [100%]
 1:cxfs-sysadm_base-server##### [ 10%]
 2:cxfs-sysadm_base-client##### [ 20%]
 3:cxfs-sysadm_xvm-server ##### [ 30%]
 4:cxfs-sysadm_cxfs-server##### [ 40%]
 5:cxfs-sysadm_cluster_bas##### [ 50%]
 6:cxfs-sysadm_xvm-client ##### [ 60%]
 7:cxfs-sysadm_cxfs-client##### [ 70%]
 8:cxfs-sysadm_xvm-web ##### [ 80%]
 9:cxfs-sysadm_cluster_bas##### [ 90%]
10:cxfs-sysadm_cxfs-web ##### [100%]
```

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

13. Install the `cxfs_admin` tool:

```
[root@linux cdrom]# rpm -Uvh cxfs_admin*
Preparing...                               ##### [100%]
 1:cxfs_admin                               ##### [100%]
```

14. Start the file alteration monitoring (fam) service, which is required for the GUI's use of task privileges for users:

```
[root@linux CXFS_CDROM]# /etc/init.d/fam start
Starting File Access Monitoring Daemon                               done
```

You could also enable the fam service with chkconfig so that fam service automatically starts on a reboot:

```
[root@linux CXFS_CDROM]# chkconfig fam on
```

15. Use the `/usr/cluster/bin/cxfslicense -d` command to verify license keys. For more information, see "Verifying the License Keys with `cxfslicense` on Server-Capable Nodes" on page 84.
16. If your root filesystem is not of type XFS, run the commands `depmod` and `mkinitrd`.
17. (*Upgrades*) Reboot the system in order to make the new updates to take effect.



Caution: If you installed GRIOV2 in step 11, the reboot is required.

If XVM standalone was in use prior to CXFS installation, you must reboot the system before starting CXFS services to ensure that the new `xvm` modules are loaded.

SGI ProPack Client-Only Software Installation

The CXFS software will be initially installed and configured by SGI personnel. This section provides an overview of those procedures.

Note: Package version numbers shown here are examples; your installed system may differ.

SGI ProPack Client-Only Installation Overview

Installing the CXFS client software for SGI ProPack requires approximately 50–200 MB of space, depending upon the packages installed at your site.

To install the required software on an SGI ProPack node, SGI personnel will do the following:

1. Read the release notes to learn about any late-breaking changes in the installation procedure.
2. Install the SGI ProPack release, according to the directions in the SGI ProPack documentation. Ensure that you select the SGI Licensed package group. You must install the `pcp-open` package from the SGI ProPack release.

Note: When installing the Linux OS, disconnect the system from the fabric or ensure that the drive you are installing on is not a SAN-attached drive.

3. Install any required patches. See the SGI ProPack `releasenotes/README` file for more information.



Caution: You **must** update the operating system with all security fixes, bug fixes, and enhancements available from the operating system vendor.

4. Verify that the node is running the supported Linux distribution and SGI ProPack overlay, according to the CXFS for SGI ProPack release notes. See the `/etc/SuSE-release` and `/etc/sgi-release` files.
5. If you have previously installed XVM in standalone mode, remove any remaining `sgi-xvm-standalone` package. To find and remove the package:

```
[root@linux CXFS_CDROM]# rpm -e --allmatches `rpm -qa | grep xvm-standalone`
```

If installing on an SGI ProPack 5 client, you may also need to remove weak-updates links from the `sgi-xvm-standalone` RPM. If you are running the 2.6.16.21-0.25 kernel, you would do the following:

```
[root@linux CXFS_CDROM]# rm -rf /lib/modules/2.6.16.21-0.25-default/weak-updates/os_lib
```

```
[root@linux CXFS_CDROM]# rm -rf /lib/modules/2.6.16.21-0.25-default/weak-updates/xvm
```

6. Insert and mount the *CXFS MultiOS Client 4.2* CD.
7. Install the CXFS kernel modules:

Note: This procedure uses the `rpm -U` option to update RPMs, which works for an initial installation as well as updates. For an initial installation, you could also use `-i`.

```
[root@linux cdrom]# rpm -Uvh sgi-cxfs-kmp-kernelvariant-kernelrelease-version.architecture.rpm
Preparing... ##### [100%]
 1:sgl-cxfs-kmp-kernelvariant-##### [100%]
```

Where:

- *kernelvariant* and *kernelrelease* are the kernel variant and release output by the `uname -r` command
- *version* is the version number
- *architecture* is the processor architecture type output by the `uname -i` command

Note: For SGI ProPack 5 running SLES 10, one version of CXFS may support one or more *kernelrelease* values. See the CXFS SGI ProPack release notes for the supported versions.

8. Install the user-space packages:

```
[root@linux cdrom]# rpm -Uvh cxfs_client* cxfs_util* cxfs-xvm-cmds* cxfs-doc*
Preparing... ##### [100%]
 1:cxfs-xvm-cmds ##### [ 25%]
boot.xvm 0:off 1:off 2:off 3:off 4:off 5:off 6:off
 2:cxfs_util ##### [ 50%]
 3:cxfs_client ##### [ 75%]
cxfs_client 0:off 1:off 2:off 3:on 4:off 5:on 6:off
 4:cxfs-doc ##### [100%]
boot.xvm 0:off 1:off 2:off 3:off 4:off 5:off 6:off
```

Note: The order of RPMs listed on the command line is not necessarily the same order in which they will be displayed in the `rpm` command output.

9. If you are using GRIO, install the `grio2-cmds` package:

```
[root@linux cdrom]# rpm -Uvh grio2-cmds*
Preparing...          #####[100%]
   1:grio2-cmds      #####[100%]
```

10. Edit the `/etc/cluster/config/cxfs_client.options` file as necessary. See "Client-only Node System Files" on page 134 and the `cxfs_client(1M)` man page.
11. (*Upgrades*) Reboot the system to make the updates take effect:

```
[root@linux cdrom]# reboot
```



Caution: If XVM standalone was in use prior to CXFS installation, you must reboot the system before starting CXFS services to ensure that the new `xvm` modules are loaded.

12. Modify `updatedb` behavior so that it avoids CXFS filesystems. See "Modify `updatedb` to Avoid Unnecessary Load" on page 154.

Installing the Performance Co-Pilot Agent

The `cxfs_utils` package includes a Performance Co-Pilot (PCP) agent for monitoring CXFS heartbeat, CMS status and other statistics. If you want to use this feature, you must also install the following PCP packages:

- `pcp-open`
- `pcp-sgi`

These packages are included on the first and second SGI ProPack CDs (respectively). You can obtain the open source PCP package from [ftp://oss.sgi.com/projects/pcp/download](http://oss.sgi.com/projects/pcp/download)

SGI ProPack Installation Verification

To verify that the CXFS software has been installed properly, use the `rpm -qa` command to display all of the installed packages. You can filter the output by searching for particular package name.

For example, to verify that the `cxfs-sysadm_base-lib` package has installed:

```
[root@linux root]# rpm -qa | grep cxfs-sysadm_base-lib
cxfs-sysadm_base-lib-3.0-sgi06092521
```

Note: The output above is an example. The version level may not match the installed software.

To verify the SGI ProPack release, display the `/etc/sgi-release` file.

SGI ProPack Modifications for CXFS Connectivity Diagnostics

If you want to use the cluster diagnostics to test node connectivity, the `root` user on the node running the CXFS diagnostics must be able to access a remote shell using the `rsh` command (as `root`) on all other nodes in the cluster. There are several ways of accomplishing this, depending on the existing settings in the pluggable authentication modules (PAM) and other security configuration files.

Following is one possible method. Do the following on all administration nodes in the cluster:

1. Install the `rsh-server` RPM using YaST.
2. Enable `rsh` by changing `disable yes` to `disable no` in the `/etc/xinetd.d/rsh` file.
3. Restart `xinetd`:

```
[root@linux root]# /etc/init.d/xinetd restart
```

4. Add the hostname of the node from which you will be running the diagnostics into the `/root/.rhosts` file. Make sure that the mode of the `.rhosts` file is set to 600 (read and write access for the owner only).

After you have completed running the connectivity tests, you may wish to disable `rsh` on all cluster nodes.

For more information, see the operating system documentation and the `hosts.equiv` man page.

Postinstallation Steps

This chapter discusses the following:

- "Configuring System Files" on page 127
- "IRIX: Configuring for Automatic Restart" on page 135
- "IRIX: Converting Filesystem Definitions for Upgrades" on page 135
- "SGI ProPack: Using `cxfs-reprobe` on Client-Only Nodes" on page 136

After completing these step discussed in this chapter, see Chapter 10, "Initial Setup of the Cluster" on page 163. For details about specific configuration tasks, see Chapter 11, "Reference to GUI Tasks" on page 183. For information about installing CXFS and Trusted IRIX, see Chapter 16, "Trusted IRIX and CXFS" on page 397. For information about upgrades, see "CXFS Release Versions and Rolling Upgrades" on page 307.

Configuring System Files

When you install the CXFS software, there are some system file considerations you must take into account. **The network configuration is critical.** Each node in the cluster must be able to communicate with every other node in the cluster by both logical name and IP address without going through any other network routing; proper name resolution is key. SGI recommends static routing.

This section discusses the following:

- `/etc/exports` on All Nodes" on page 127
- "Administration Node System Files" on page 128
- "Client-only Node System Files" on page 134

`/etc/exports` on All Nodes

The optional `/etc/exports` file on each node describes the filesystems that are being exported to NFS clients.

If the `/etc/exports` file contains a CXFS mount point, then when the system is booted NFS will export the empty mount point because the exports are done before

CXFS is running. When CXFS on the node joins membership and starts mounting filesystems, the `clconfd-pre-mount` script searches the `/etc/exports` file looking for the mountpoint that is being mounted. If found, the script unexports the mountpoint directory because if it did not the CXFS mount would fail. After successfully mounting the filesystem, the `clconfd-post-mount` script will search the `/etc/exports` file and export the mount point if it is found in the `/etc/exports` file.

For more information, see "CXFS Mount Scripts" on page 325.

Administration Node System Files

This section discusses system files on administration nodes:

- `/etc/services` on CXFS Administration Nodes" on page 128
- `cad.options` on CXFS Administration Nodes" on page 129
- `fs2d.options` on CXFS Administration Nodes" on page 130
- `clconfd.options` on CXFS Administration Nodes" on page 133

`/etc/services` on CXFS Administration Nodes

The `/etc/services` file on each CXFS administration contains entries for `sgi-cad` and `sgi-crsd`. The port numbers assigned for these processes must be the same in all nodes in the pool.

Note: You will see an `inst` message that says `sgi-cmsd` and `sgi-gcd` must be added to `/etc/services`. This is true only for coexecution with FailSafe, or when running only FailSafe; if you are running just CXFS, you do not need `sgi-cmsd`. CXFS does not require `sgi-cmsd`.

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
```

cad.options on CXFS Administration Nodes

The `cad.options` file on each CXFS administration node contains the list of parameters that the cluster administration daemon reads when the `cad` process is started. The files are located as follows:

- IRIX: `/etc/config/cad.options`
- SGI ProPack: `/etc/cluster/config/cad.options`

`cad` provides cluster information.

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

<code>--append_log</code>	Append <code>cad</code> logging information to the <code>cad</code> log file instead of overwriting it.
<code>--log_file filename</code>	<code>cad</code> log filename. Alternately, this can be specified as <code>-lf filename</code> .
<code>-vvvv</code>	Verbosity level. The number of <code>v</code> characters indicates the level of logging. Setting <code>-v</code> logs the fewest messages; setting <code>-vvvv</code> 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 the `/etc/config/cad.options` file cannot be modified using `cxfs_admin` or the GUI.

If you make a change to the `cad.options` file at any time other than initial configuration, 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:

- IRIX:


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

- SGI ProPack:

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

If you execute this command on a running cluster, it will remain up and running. However, the GUI will lose connection with the cad daemon; the GUI will prompt you to reconnect.

fs2d.options on CXFS Administration Nodes

The `fs2d.options` file on each CXFS administration node contains the list of parameters that the `fs2d` daemon reads when the process is started. (The `fs2d` daemon manages the distribution of the cluster database (CDB) across the CXFS administration nodes in the pool.) The files are located as follows:

- IRIX: `/etc/config/fs2d.options`
- SGI ProPack: `/etc/cluster/config/fs2d.options`

Table 8-1 shows the options can that can be set in the `fs2d.options` file.

Table 8-1 `fs2d.options` File Options

Option	Description
<code>-logevents <i>event name</i></code>	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 <i>log destination</i></code>	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>logfile</code> .
<code>-logfile <i>filename</i></code>	Set log filename. The default is <code>/var/cluster/ha/log/fs2d_log</code> .
<code>-logfilemax <i>maximum size</i></code>	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.

Option	Description
<code>-loglevel <i>loglevel</i></code>	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 <i>trace_class</i></code>	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 <i>filename</i></code>	Set trace filename. There is no default.
<code>-tracefilemax <i>maximum_size</i></code>	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.
<code>-[no]tracelog</code>	[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 <code>-tracelog</code> .
<code>-[no]parent_timer</code>	[Do not] exit when the parent exits. The default is <code>-noparent_timer</code> .
<code>-[no]daemonize</code>	[Do not] run as a daemon. The default is <code>-daemonize</code> .
<code>-l</code>	Do not run as a daemon.
<code>-h</code>	Print usage message.
<code>-o help</code>	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, are sent to the `/var/cluster/ha/log/fs2d_log` file. When the file size reaches 10 MB, 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.

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 CXFS administration nodes or by entering the following command:

- IRIX:

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

- SGI ProPack:

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

If you execute this command on a running cluster, it should remain up and running. However, the GUI will lose connection with the cad daemon; the GUI will prompt you to reconnect.

Example 1

The following example shows an `/etc/config/fs2d.options` file that directs logging and tracing information as follows:

- All log events are sent to:
 - IRIX: `/var/adm/SYSLOG`
 - SGI ProPack: `/var/log/messages`
- Tracing information for RPCs, updates, and transactions are sent to `/var/cluster/ha/log/fs2d_ops1`.

When the size of this file exceeds 100,000,000 bytes, 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 added for readability.)

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

Example 2

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 100,000,000 bytes is specified. `-tracelog` directs the tracing to the log file.

(Line breaks 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.
```

clconfd.options on CXFS Administration Nodes

The `clconfd.options` file on each CXFS administration node contains the list of parameters that the `clconfd` daemon reads when the process is started. The files are located as follows:

- IRIX: `/etc/config/clconfd.options`
- SGI ProPack: `/etc/cluster/config/clconfd.options`

Table 8-2 shows the options that can be set in the `fs2d.options` file.

Table 8-2 `clconfd.options` File Options

Option	Description
<code>-c CDBfile</code>	Reads the cluster database configuration from the specified <i>CDBfile</i> file. The default file is <code>/var/cluster/cdb/cdb.db</code> .
<code>-d debugfile</code>	Enables printing of ha-fence debug information to the specified file <i>debugfile</i> . The default is to print no information.
<code>-h</code>	Prints a help message for <code>clconfd.options</code> .
<code>-l</code>	Runs <code>clconfd</code> in the foreground. (For SGI development debugging purposes only. Do not use this option unless directed to do so by SGI support.) The default is to run <code>clconfd</code> in the background.
<code>-s loglevel</code>	Specifies the log level to use for logging to standard error. The default is 0 (no logging). For information about log levels, see "Configure Log Groups with the GUI" on page 231.
<code>-R</code>	Disables real-time scheduling. By default, real-time scheduling is enabled.

For example, to print hafence debug information to the file `/tmp/hafence.log`, add the following line to the `clconfd.options` file:

```
-d /tmp/hafence.log
```

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

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

Client-only Node System Files

This section discusses the `cxfs_client.options` file for IRIX and SGI ProPack client-only nodes. For client-only nodes running other operating systems, see the *CXFS MultiOS Client-Only Guide for SGI InfiniteStorage*.

On client-only nodes, you can modify the CXFS client daemon (`/usr/cluster/bin/cxfs_client`) by placing options in the `cxfs_client.options` file:

- IRIX: `/etc/config/cxfs_client.options`
- SGI ProPack: `/etc/cluster/config/cxfs_client.options`

The available options are documented in the `cxfs_client` man page.



Caution: Some of the options are intended to be used internally by SGI only for testing purposes and do not represent supported configurations. Consult your SGI service representative before making any changes.

For example, to see if `cxfs_client` is using the options in `cxfs_client.options`, enter the following:

```
irix# ps -ax | grep cxfs_client
3612 ?      S        0:00 /usr/cluster/bin/cxfs_client -i cxfs3-5
3841 pts/0  S        0:00 grep cxfs_client
```

IRIX: Configuring for Automatic Restart

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

```
# 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
```

IRIX: Converting Filesystem Definitions for Upgrades

The structure of the CXFS filesystem configuration was changed with the release of IRIX 6.5.13f. Upgrading to the 6.5.13f release provided an automatic conversion from the old structure to the new structure. However, if you are upgrading directly from 6.5.12f or earlier, (without first installing and running 6.5.13f), you must convert your CXFS filesystem definitions manually.

Upgrading from 6.5.12f or Earlier

Note: If you are upgrading from 6.5.13f or later, you do not need to follow the instructions in this section. Filesystems definitions are automatically and transparently converted when running 6.5.13f.

After upgrading from 6.5.12f or earlier, you will notice that the CXFS filesystems are no longer mounted, and that they do not appear in the GUI or `cxfs_admin` queries. To convert all of the old CXFS filesystem definitions to the new format, run the following command from one of the 6.5.14f or later nodes in the CXFS cluster:

```
# /usr/sysadm/privbin/cxfsfilesystemUpgrade
```

After running this command, the CXFS filesystems should appear in the GUI and `cxfs_admin` output, and they should be mounted if their status was enabled and CXFS services are active.



Caution: This conversion is a one-time operation and **should not** be run a second time. If you make changes to the filesystem and then run `cxfsfilesystemUpgrade` for a second time, all of your changes will be lost.

Running with All IRIX Nodes Upgraded to 6.5.14f or Later

After all of the IRIX nodes in the cluster have been upgraded to 6.5.14f or later, it is recommended that you destroy the old CXFS filesystem definitions, in order to prevent these stale definitions from overwriting the new definitions if the `cxfsfilesystemUpgrade` command were to be run again accidentally. To destroy the old CXFS filesystem definitions, enter the following, where *clustername* is the name of the cluster in the CXFS cluster database:

```
# /usr/cluster/bin/cdbutil -c "delete #cluster#clustername#Cellular#FileSystems"
```

SGI ProPack: Using `cxfs-reprobe` on Client-Only Nodes

When `cxfs_client` needs to rescan disk buses, it executes the `/var/cluster/cxfs_client-scripts/cxfs-reprobe` script. This requires the use of parameters in SGI ProPack due to limitations in the Linux SCSI layer. You can export these parameters from the `/etc/cluster/config/cxfs_client.options` file.

The `cxfs_reprobe` script detects the presence of the SCSI layer on the system and probes all SCSI layer devices by default. You can override this decision by setting `CXFS_PROBE_SCSI` to 0 to disable the probe or 1 to force the probe (default).

When a SCSI scan is performed, all buses/channels/IDs and LUNs are scanned by default to ensure that all devices are found. You can override this decision by setting one or more of the environment variables listed below. This may be desired to reduce lengthy probe times.

The following summarizes the environment variables (separate multiple values by white space and enclose within single quotation marks):

```
CXFS_PROBE_SCSI=0|1
```

Stops (0) or forces (1) a SCSI probe. Default: 1

`CXFS_PROBE_SCSI_BUSES=BusList`

Scans the buses listed. Default: All buses (-)

`CXFS_PROBE_SCSI_CHANNELS=ChannelList`

Scans the channels listed. Default: All channels (-)

`CXFS_PROBE_SCSI_IDS=IDList`

Scans the IDs listed. Default: All IDs (-)

`CXFS_PROBE_SCSI_LUNS=LunList`

Scans the LUNs listed. Default: All LUNs (-)

For example, the following would only scan the first two SCSI buses:

```
export CXFS_PROBE_SCSI_BUSES='0 1'
```

The following would scan 16 LUNs on each bus, channel, and ID combination (all on one line):

```
export CXFS_PROBE_SCSI_LUNS='0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15'
```

Other options within the `/etc/cluster/config/cxfs_client.options` file begin with a `-` character. Following is an example `cxfs_client.options` file:

```
# Example cxfs_client.options file
#
-Dnormal -serror
export CXFS_PROBE_SCSI_BUSES=1
export CXFS_PROBE_SCSI_LUNS='0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20'
```

Note: The `-` character or the term `export` must start in the first position of each line in the `cxfs_client.options` file; otherwise, they are ignored by the `/etc/init.d/cxfs_client` script.

Best Practices

This chapter summarizes configuration and administration best-practices information for CXFS:

- "Configuration Best Practices" on page 139
- "Administration Best Practices" on page 152

For the latest information and a matrix of supported CXFS and operating system software, see http://support.sgi.com/content_request/838562/index.html on Supportfolio.

Configuration Best Practices

This section discusses the following configuration topics:

- "Fix Network Issues First" on page 140
- "Use a Private Network" on page 140
- "Provide Enough Memory" on page 141
- "Use CXFS Configuration Tools Appropriately" on page 141
- "Use Netscape and the Brocade Switch GUI on IRIX Appropriately" on page 142
- "Ensure Cluster Database Membership Quorum Stability" on page 142
- "Be Consistent in Configuration" on page 143
- "Do Not Mix Metadata Operating System Flavors" on page 143
- "Use the Correct Mix of Software Releases" on page 143
- "Form a Small Functional Cluster First" on page 144
- "Choose a Metadata Server that is Dedicated to CXFS Work" on page 144
- "Use an Odd Number of Server-Capable Nodes" on page 145
- "Make Most Nodes Client-Only" on page 145
- "Use a Client-Only Tiebreaker" on page 145

- "Protect Data Integrity on All Nodes" on page 146
- "Minimize the Number of Switches" on page 149
- "Configure Filesystems Properly" on page 149
- "Verify the Configuration" on page 151
- "Use the Recovery Timeout Mechanism" on page 151
- "Use Proper Storage Management Procedures" on page 151

Fix Network Issues First

If there are any network issues on the private network, fix them before trying to use CXFS. Ensure that you understand the information in "Hostname Resolution and Network Configuration Rules" on page 93.

Use a Private Network

You must use a private network for CXFS metadata traffic:

- A private network is a requirement.
- The private network is used for metadata traffic and should not be used for other kinds of traffic.
- A stable private network is important for a stable CXFS cluster environment.
- Two or more clusters should not share the same private network. A separate private network switch is required for each cluster.
- The private network should contain at least a 100-Mbit network switch. A network hub is not supported and should not be used.
- All cluster nodes should be on the same physical network segment (that is, no routers between hosts and the switch).
- The private network must be configured as the highest priority network for the cluster. The public network may be configured as a lower priority network to be used by CXFS network failover in case of a failure in the private network.
- A virtual local area network (VLAN) is not supported for a private network.

- Use private (10.x.x.x, 176.16.x.x, or 192.168.x.x) network addresses (RFC 1918).
- When administering more than one CXFS cluster, use unique private network addresses for each cluster. If you have multiple clusters using the same public network as the backup CXFS metadata network, use unique cluster names and cluster IDs.

Provide Enough Memory

There should be at least 2 GB of RAM on the system. A metadata server must have at least 1 processor and 1 GB of memory more than what it would need for its normal workload (work other than CXFS). In general, this means that the minimum configuration would be 2 processors and 2 GB of memory. If the metadata server is also doing NFS or Samba serving, then more memory is recommended (and the `nbuf` and `ncsize` kernel parameters should be increased from their defaults). CXFS makes heavy use of memory for caching.

If a very large number of files (tens of thousands) are expected to be accessed at any one time, additional memory over the minimum is recommended to avoid throttling memory. Estimate the maximum number of inodes that will be accessed during a 2-minute window and size the metadata server memory for that number. (The inode references are not persistent in memory and are removed after about 2 minutes of non-use.) Use the following general rule to determine the amount of memory required when the number of open files at any one time may be this large:

$$2 \text{ KB} \times \#inodes = \text{metadata_server_memory}$$

To avoid problems during metadata server recovery/relocation, all potential metadata servers should have as much memory as the active metadata server

Use CXFS Configuration Tools Appropriately

The `cxfs_admin` command-line tool waits for a command to be completed before continuing and provides <TAB> completion of commands. It also provides scripting capabilities.

The GUI provides a convenient display of a cluster and its components through the view area. You should use it 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 view area 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. For more information, see "Starting the GUI" on page 184.

Note: When running the GUI on IRIX, do not move to another IRIX desktop while GUI action is taking place; this can cause the GUI to crash.

Do not attempt to make simultaneous changes using `cxfs_admin`, the GUI, and/or `cmgr`. Use one tool at a time.

Use Netscape and the Brocade Switch GUI on IRIX Appropriately

When accessing the Brocade Web Tools V2.0 through Netscape on an IRIX node, you must first enter one of the following before starting Netscape:

- For `sh` or `ksh` shells:

```
$ NOJIT=1; export NOJIT
```

- For `csh` shell:

```
% setenv NOJIT 1
```

If this is not done, Netscape will crash with a core dump.

Ensure Cluster Database Membership Quorum Stability

The cluster database membership quorum must remain stable during the configuration process. If possible, use multiple windows to display the `fs2d_log` file for each CXFS administration node while performing configuration tasks. Enter the following:

```
# tail -f /var/cluster/ha/log/fs2d_log
```

Check the member count when it prints new quorums. Under normal circumstances, it should print a few messages when adding or deleting nodes, but it should stop within a few seconds after a new quorum is adopted.

If not enough machines respond, there will not be a quorum. In this case, the database will not be propagated.

If you detect cluster database membership quorum problems, fix them before making other changes to the database. Try restarting the cluster administration daemons on

the node that does not have the correct cluster database membership quorum, or on all nodes at the same time.

Enter the following on administration nodes:

- IRIX:

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

- SGI ProPack:

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

Note: You could also use the `restart` option to stop and start.

Please provide the `fs2d` log files when reporting a cluster database membership quorum problem.

Be Consistent in Configuration

Be consistent in configuration files for nodes across the pool, and when configuring networks. Use the same names in the same order. See "Configuring System Files" on page 127.

Do Not Mix Metadata Operating System Flavors

Mixing SGI ProPack and IRIX metadata servers in one cluster is not supported. All server-capable administration nodes in a cluster must be either all SGI ProPack or all IRIX.

Use the Correct Mix of Software Releases

Create a new cluster using server-capable nodes that have the same version of the OS release and the same version of CXFS installed.

All nodes should run the same level of CXFS and the same level of operating system software, according to platform type. To support upgrading without having to take

the whole cluster down, nodes can run different CXFS releases during the upgrade process.



Caution: You must upgrade all server-capable nodes before upgrading any client-only nodes (servers must run the same or later release as client-only nodes.) Operating a cluster with clients running a mixture of older and newer CXFS versions will result in a performance loss.

For details, see the platform-specific release notes and "CXFS Release Versions and Rolling Upgrades" on page 307.

Form a Small Functional Cluster First

Ensure that you follow the instructions in "Preliminary Cluster Configuration Steps" on page 164

For large clusters, SGI recommends that you first form a functional cluster with just server-capable nodes and then build up the large cluster in small groups of client-only nodes. This method make it easier to locate and fix problems, should any occur. See "Configuring a Large Cluster" on page 176.

Choose a Metadata Server that is Dedicated to CXFS Work

The nodes that you use as potential metadata servers must be dedicated to CXFS and filesystems work (such as Samba or NFS). Standard services (such as `ftp`, DNS, and NIS) are permitted, but any other applications (such as analysis, simulation, and graphics) must be avoided.

Note: Octane, Octane2, and Tezro systems are not supported as metadata servers.

As of CXFS 3.4, non-dedicated nodes are not supported as CXFS metadata servers. Running a metadata server in a non-dedicated manner will void the support contract. If the use of an application is desired on a metadata server, SGI will provide a quotation to perform the following work:

- Audit the solution
- Design a supportable configuration

- Implement the changes

A statement of work will be created and implementation will begin after mutual agreement with the customer.

If additional products are required from SGI, the customer will be responsible for obtaining a quote and providing a purchase order before any corrective action begins. SGI will not correct unsupported configurations without compensation and reserves the right to terminate or suspend the support agreement.

SGI recommends that all potential metadata servers be configured with system reset in order to protect data integrity. See "Protect Data Integrity on All Nodes" on page 146.

Use an Odd Number of Server-Capable Nodes

Use an odd number of server-capable nodes with CXFS services running and a client-only CXFS tiebreaker node if you have more than two nodes total in the cluster. See "Use a Client-Only Tiebreaker" on page 145.

Make Most Nodes Client-Only

You should define most nodes as client-only nodes and define just the nodes that may be used for CXFS metadata as server-capable administration nodes. Use client administration nodes only for a Failsafe co-execution node that cannot be a metadata server (Failsafe requires that a node be either a server-capable administration node or a client administration node).

The advantage to using client-only nodes is that they do not keep a copy of the cluster database; they contact an administration node to get configuration information. It is easier and faster to keep the database synchronized on a small set of nodes, rather than on every node in the cluster. In addition, if there are issues, there will be a smaller set of nodes on which you must look for problem.

Use a Client-Only Tiebreaker

SGI recommends that you always define a client-only node CXFS tiebreaker for all clusters with more than one server-capable node and at least one client-only node. (Using server-capable nodes as a tiebreaker is not recommended because these nodes always affect CXFS kernel membership.)

Having a tiebreaker is critical when there are an even number of server-capable administration nodes. A tiebreaker avoids the problem of multiple-clusters being formed (also known as *split-brain syndrome*) while still allowing the cluster to continue if one of the metadata servers fails.

As long as there is a reliable client-only node in the cluster, the client-only node should be used as tiebreaker. Server-capable nodes are not recommended as tiebreaker nodes. Only if there are more than three server-capable nodes in the cluster should one of them be used as the tie-breaker.

The tiebreaker is of benefit in a cluster even with an odd number of server-capable administration nodes because when one of the server-capable administration nodes is removed from the cluster, it effectively becomes a cluster with an even-number of server-capable nodes.

Note the following:

- If exactly two server-capable nodes are configured and there are no client-only nodes, **neither** server-capable node should be set as the tiebreaker. (If one node was set as the tiebreaker and it failed, the other node would also shut down.)
- If exactly two server-capable nodes are configured and there is at least one client-only node, you should specify the client-only node as a tiebreaker.

If one of the server-capable nodes is the CXFS tiebreaker in a two-server-capable-node cluster, failure of that node or stopping the CXFS services on that node will result in a cluster-wide forced shutdown. If you use a client-only node as the tiebreaker, either server could fail but the cluster would remain operational via the other server.

- If there are an even number of servers and there is no tiebreaker set, the fail policy must not contain the `shutdown` option because there is no notification that a shutdown has occurred. See "Isolating Failed Nodes: Failure Policies" on page 28.

SGI recommends that you start CXFS services on the tiebreaker client after the metadata servers are all up and running, and before CXFS services are started on any other clients.

Protect Data Integrity on All Nodes

All nodes must be configured to protect data integrity in case of failure. System reset and/or I/O fencing is required to ensure data integrity for all nodes.

Note: No matter what the cluster components are, SGI recommends that you use a system reset configuration on server-capable nodes in order to protect data integrity and improve server reliability. I/O fencing (or system reset when available) must be used on client-only nodes.

See also "Isolating Failed Nodes: Failure Policies" on page 28.

System Reset

You should configure system reset for any potential metadata servers in order to protect data integrity. (I/O fencing is appropriate for client-only nodes.) This means that nodes without system reset capability, such as Fuel systems, should not be potential metadata servers.

Note: If the failure hierarchy contains `reset` or `fencereset`, the reset might be performed before the system kernel core-dump can complete, resulting in an incomplete core-dump.

I/O Fencing

Nodes without system reset capability (such as AIX, Linux third-party, Mac OS X, Solaris, and Windows nodes) require I/O fencing. I/O fencing is also appropriate for nodes with system controllers if they are client-only nodes.

You should use the `admin` account when configuring I/O fencing. On a Brocade switch running 4.x.x.x or later firmware, modify the `admin` account to restrict it to a single `telnet` session. For details, see the release notes.

If you use I/O fencing, you must keep the `telnet` port on the switch free at all times; **do not** perform a `telnet` to the switch and leave the session connected.

If you use I/O fencing, SGI recommends that you use a switched network of at least 100baseT.

You should isolate the power supply for the switch from the power supply for a node and its system controller. You should avoid any possible situation in which a node can continue running while both the switch and the system controller lose power. Avoiding this situation will prevent the possibility a split-brain scenario.

You must put switches used for I/O fencing on a network other than the primary CXFS private network so that problems on the CXFS private network can be dealt with by the fencing process and thereby avoid data corruption issues. The network to which the switch is connected must be accessible by all administration nodes in the cluster.

If you manually change the port status, the CXFS database will not be informed and the status output by the `cxfs_admin` command will not be accurate. To update the CXFS database, run the following command:

```
# hafence -U
```

For other details, see the release notes.

Avoid Network Partition

The worst scenario is one in which the node does not detect the loss of communication but still allows access to the shared disks, leading to data corruption. 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 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.

In a case of a true network partition, where an existing CXFS kernel membership splits into two halves (each with half the total number of server-capable nodes), the following will happen:

- If the CXFS tiebreaker and system reset or I/O fencing are configured, the half with the tiebreaker node will reset or fence the other half. The side without the tiebreaker will attempt to forcibly shut down CXFS services.
- If there is no CXFS tiebreaker node but system reset or I/O fencing is configured, each half will attempt to reset or fence the other half using a delay heuristic. One half will succeed and continue. The other will lose the reset/fence race and be rebooted/fenced.
- If there is no CXFS tiebreaker node and system reset or I/O fencing is not configured, then both halves will delay, each assuming that one will win the race

and reset the other. Both halves will then continue running, because neither will have been reset or fenced, leading to likely data corruption.

To avoid this situation, you should configure a tiebreaker node, and you must use system reset or I/O fencing. However, if the tiebreaker node (in a cluster with only two server-capable nodes) fails, or if the administrator stops CXFS services, the other node will do a forced shutdown, which unmounts all CXFS filesystems.

If the network partition persists when the losing half attempts to form a CXFS kernel membership, it will have only half the number of server-capable nodes and be unable to form an initial CXFS kernel membership, preventing two CXFS kernel memberships in a single cluster.

For more information, contact SGI professional or managed services.

Minimize the Number of Switches

CXFS is more efficient with fencing operations with a smaller number of large switches rather than a large number of smaller switches.

Configure Filesystems Properly

Configure filesystems properly:

- Use a filesystem block size that is common to all CXFS OS platforms. Each CXFS OS platform supports a unique range of filesystem block sizes, but all of them support a filesystem block size of 4096 bytes. For this reason, SGI recommends 4-KB filesystems for compatibility with all CXFS platforms. For details on the filesystem block sizes supported by each CXFS OS platform, see Appendix E, "Filesystem Specifications" on page 537 and the "Filesystem and Logical Unit Specifications" appendix in the *CXFS MultiOS Client-Only Guide for SGI InfiniteStorage*.
- Determine whether or not to have all filesystems served off of one metadata server or to use multiple metadata servers to balance the load, depending upon how filesystems will be accessed. The more often a file is accessed, the greater the stress; a filesystem containing many small files that are accessed often causes greater stress than a filesystem with a few large files that are not accessed often. CXFS performs best when data I/O operations are greater than 16 KB and large files are being accessed. (A lot of activity on small files will result in slower performance.)

- Enable the forced unmount feature for CXFS filesystems, which is off by default. Many sites have found that enabling this feature improves the stability of their CXFS clusters, particularly in situations where the filesystem must be unmounted.

On IRIX nodes, this feature uses 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. That is, it attempts to terminate any I/O going to the filesystem, so that it can unmount it promptly, rather than having to wait for the I/O to finish on its own, causing the unmount to possibly fail.

On SGI ProPack nodes, a similar function is performed with the `fuser -m -k` command and the `umount` command

This feature is available through the following CXFS GUI menu:

Tasks

> **Filesystems**

> **Unmount a CXFS Filesystem**

You can also specify this feature using `cxfs_admin` to define the filesystem. See "Unmount a CXFS Filesystem with `cxfs_admin`" on page 293, "Create or Modify a CXFS Filesystem with `cxfs_admin`" on page 287, and "Unmount CXFS Filesystems with the GUI" on page 246.

- If you are using NFS or Samba, you should have the NFS or Samba server run on the active metadata server.
- Do not use nested mount points. IRIX nodes do not permit nested mount points on CXFS filesystems; that is, you cannot mount an IRIX XFS or CXFS filesystem on top of an existing CXFS filesystem. Although it is possible to mount other filesystems on top of an SGI ProPack CXFS filesystem, this is not recommended.
- Perform reconfiguration (including but not limited to adding and deleting filesystems or nodes) during a scheduled cluster maintenance shift and not during production hours.

Verify the Configuration

You should always run the following command after any significant configuration change, or whenever problems, warnings or errors occur:

```
/usr/cluster/bin/cxfs-config -xfs -xvm
```

The CXFS GUI and `cxfs_admin` do not always prevent poor configurations. The `status` command in `cxfs_admin` will indicate some potential problems and the `cxfs-config` tool can detect a large number of potential problems.

Use the Recovery Timeout Mechanism

The recovery timeout mechanism prevents the cluster from hanging and keeps filesystems available in the event that a node becomes unresponsive.

When recovery timeout is enabled, nodes are polled for progress after a recovery has begun. If recovery for a node is not making progress according to the specified polls, the recovery is considered stalled and the node will shut down or panic. For example, to enable the recovery timeout to begin monitoring after 5 minutes, monitor every 2 minutes, declare a node's recovery stalled after 15 minutes of without progress, and panic the node with stalled recovery, you would set the following:

```
cxfs_recovery_timeout_start 300
cxfs_recovery_timeout_period 120
cxfs_recovery_timeout_stalled 900
cxfs_recovery_timeout_panic 1
```

For details about the parameters, see "Site-Changeable System Tunable Parameters" on page 348.

Use Proper Storage Management Procedures

You should configure storage management hardware and software according to its documentation and use proper storage management procedures, including the following:

- Assign IP addresses to all storage controllers and have them network-connected (but not on the private CXFS metadata network) and manageable via out-of-band management

Note: Do not use in-band management (which can cause problems if there is a loss of Fibre Channel connectivity)

- Keep a copy of the array configuration
- Monitor for read errors that do not result in drive strikes
- Keep a copy of the XVM volume configuration

Administration Best Practices

This section discusses the following administration topics:

- "Do Not Run User Jobs on Metadata Servers" on page 153
- "Do Not Run Backups on a Client Node" on page 153
- "Use cron Jobs Properly" on page 153
- "Modify updatedb to Avoid Unnecessary Load" on page 154
- "Repair Filesystems with Care" on page 154
- "Defragment Filesystems with Care" on page 155
- "Use Relocation and Recovery Properly" on page 155
- "Shut Down Nodes Unobtrusively" on page 155
- "Remove Unused Cluster Components" on page 156
- "Use fam Properly" on page 156
- "Use Trusted IRIX Consistently" on page 157
- "Upgrade the Software Properly" on page 157
- "Use Fast Copying for Large CXFS Files" on page 157
- "Log File Names and Sizes" on page 157
- "Use System Capacity Wisely" on page 158
- "Reboot Before Changing Node ID or Cluster ID" on page 158

- "Restart CXFS on a Node after an Administrative CXFS Stop" on page 158
- "Restart the Cluster In an Orderly Fashion" on page 159
- "Disable Reset Capability If You Remove Reset Lines" on page 159
- "Avoid Performance Problems with Unwritten Extent Tracking and Exclusive Write Tokens" on page 159
- "Use the Appropriate Version of `lcrash` for SGI ProPack" on page 161
- "Disk Layout Optimization for Approved Media Customers" on page 161

Do Not Run User Jobs on Metadata Servers

Do not run user jobs on the CXFS metadata server node.

Do Not Run Backups on a Client Node

SGI recommends that backups are done on the metadata server.

Do not run backups on a client node, because it causes heavy use of non-swappable kernel memory on the metadata server. During a backup, every inode on the filesystem is visited, and if done from a client, it imposes a huge load on the metadata server. The metadata server may experience typical out-of-memory symptoms, and in the worst case can even become unresponsive or crash.

Use cron Jobs Properly

Because CXFS filesystems are considered as local on all nodes in the cluster, the nodes may generate excessive filesystem activity if they try to access the same filesystems simultaneously while running commands such as `find`, `ls`, or SGI ProPack `slocate`. You should build databases for `rfind` and GNU `locate` only on the metadata server.

On IRIX systems, the default `root` `crontab` on some platforms has the following `find` job that should be removed or disabled on all nodes (line breaks added here for readability):

```
0 5 * * * /sbin/suattr -m -C CAP_MAC_READ,
CAP_MAC_WRITE,CAP_DAC_WRITE,CAP_DAC_READ_SEARCH,CAP_DAC_EXECUTE=eip
-c "find / -local -type f '(' -name core -o -name dead.letter ')' -atime +7
-mtime +7 -exec rm -f '{} ' ;'"
```

Edit the nodes' `crontab` file to only execute this `find` command on one metadata server of the cluster.

On SGI ProPack systems, there is often a `cron` job to execute `updatedb`, which can be problematic. You must remove this `cron` job or modify it to exclude CXFS directories. (On SGI ProPack for SGI ProPack systems on which you are using local XFS, you cannot add `xfs` to the `PRUNEFSS` configuration variable to exclude all CXFS filesystems because this would also exclude local XFS filesystems.)

Modify `updatedb` to Avoid Unnecessary Load

CXFS filesystems are mounted on each node in the cluster. Therefore, running the default `updatedb` or `slocate` on each SGI ProPack and Linux third-party client will cause extra unnecessary load in the cluster, and may result in racing and an incorrect filename database.

To avoid this situation, add CXFS mount points to the parameter in the following files:

- Red Hat Enterprise Linux: `PRUNEPATHS` in `/etc/updatedb.config`
- SUSE Linux Enterprise Server 10 (SLES 10): `UPDATEDB_PRUNEPATHS` in `/etc/sysconfig/locate`
- SGI ProPack 5 on SLES 10: `UPDATEDB_PRUNEPATHS` in `/etc/sysconfig/locate`

Repair Filesystems with Care

Always contact SGI technical support before using `xfs_repair` on CXFS filesystems. You must first ensure that you have an actual case of data corruption and retain valuable metadata information by replaying the XFS logs before running `xfs_repair`.



Caution: If you run `xfs_repair` without first replaying the XFS logs, you may introduce data corruption.

You should run `xfs_ncheck` and capture the output to a file before running `xfs_repair`. If running `xfs_repair` results in files being placed in the `lost+found` directory, the saved output from `xfs_ncheck` may help you to identify the original names of the files.

Only use `xfstool repair` on metadata servers and only when you have verified that all other cluster nodes have unmounted the filesystem.

When using `xfstool repair`, make sure it is run only on a cleanly unmounted filesystem. If your filesystem has not been cleanly unmounted, there will be un-committed metadata transactions in the log, which `xfstool repair` will erase. This usually causes loss of some data and messages from `xfstool repair` that make the filesystem appear to be corrupted.

If you are running `xfstool repair` right after a system crash or a filesystem shutdown, your filesystem is likely to have a dirty log. To avoid data loss, you **MUST** mount and unmount the filesystem before running `xfstool repair`. It does not hurt anything to mount and unmount the filesystem locally, after CXFS has unmounted it, before `xfstool repair` is run.

For more information, see the *IRIX Admin: Disks and Filesystems*.

Defragment Filesystems with Care

Using `xfstool fsr` to defragment CXFS filesystems is not recommended except on read-mostly filesystems because `xfstool fsr` badly fragments the free space. XFS actually does best at maintaining contiguous free space and keeping files from being fragmented if `xfstool fsr` is not run as long as there is a moderate (10% or more) free space available on the filesystem.

The `xfstool fsr` tool is useful when defragmenting specific files but not filesystems in general.

Use Relocation and Recovery Properly

Use relocation and recovery only on standby nodes. A *standby node* is a server-capable administration node that is configured as a potential metadata server for a given filesystem, but does not currently run any applications (including NFS and Samba) that will use that filesystem. The node can run applications that use other filesystems.

Shut Down Nodes Unobtrusively

Use the proper procedures for shutting down nodes. See "Removing and Restoring Cluster Members" on page 369.

When shutting down, resetting, or restarting a CXFS client-only node, do not stop CXFS services on the node. (Stopping CXFS services is more intrusive on other nodes in the cluster because it updates the cluster database. Stopping CXFS services is appropriate only for a CXFS administration node.) Rather, let the CXFS shutdown scripts on the node stop CXFS when the client-only node is shut down or restarted.

If you are going to perform maintenance on a potential metadata server, you should first shut down CXFS services on it. Disabled nodes are not used in CXFS kernel membership calculations, so this action may prevent a loss of quorum.

Remove Unused Cluster Components

As long as a server-capable node remains configured in the cluster database, it counts against cluster database quorum. However, the way it impacts the cluster depends upon the actual node count.

If a server-capable administration node is expected to be down for longer than the remaining mean-time to failure (MTTF) of another server-capable node in the cluster, you should remove it from the cluster and the pool to avoid cluster database membership and CXFS membership quorum problems. See the following sections:

- "Modify a Cluster Definition with the GUI" on page 225
- "Delete a Node with `cxfs_admin`" on page 279

You should leave a client-only node in the cluster database unless you are permanently removing it.

You should also remove the definitions for unused objects such as filesystems and switches from the cluster database. This will improve the cluster database performance and reduce the likelihood of cluster database problems.

Use `fam` Properly

If you want to use the file alteration monitor (`fam`), you must remove the `/dev/imon` file from CXFS nodes. Removing this file forces `fam` to poll the filesystem. For more information about the monitor, see the `fam` man page.

Use Trusted IRIX Consistently

If you want to run CXFS and Trusted IRIX, all server-capable nodes in the cluster must run Trusted IRIX. The client-only nodes can run IRIX. SGI ProPack and the multiOS platforms are not supported in a cluster with Trusted IRIX. You should configure your system such that all nodes in the cluster have the same user IDs, access control lists (ACLs), and capabilities.

Upgrade the Software Properly

Do the following when upgrading the software:

- Read the release notes when installing and/or upgrading CXFS. These notes contain useful information and caveats needed for a stable install/upgrade.
- Do not make any other configuration changes to the cluster (such as adding new nodes or filesystems) until the upgrade of all nodes is complete and the cluster is running normally.

Use Fast Copying for Large CXFS Files

You can use the `cxfscp(1)` command to quickly copy large files (64 KB or larger) to and from a CXFS filesystem. It can be significantly faster than `cp(1)` on CXFS filesystems because it uses multiple threads and large direct I/Os to fully use the bandwidth to the storage hardware.

Files smaller than 64 KB do not benefit from large direct I/Os. For these files, `cxfscp` uses a separate thread using buffered I/O, similar to `cp(1)`.

The `cxfscp` command is available on on IRIX, SGI ProPack, Linux, and Windows platforms. However, some options are platform-specific, and other limitations apply. For more information and a complete list of options, see the `cxfscp(1)` man page.

Log File Names and Sizes

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

Periodically, you should rotate log files to avoid filling your disk space; see "Log File Management" on page 339. If you are having problems with disk space, you may

want to choose a less verbose log level; see "Configure Log Groups with the GUI" on page 231.

Use System Capacity Wisely

To avoid a loss of connectivity between the metadata server and the CXFS clients, do not oversubscribe the metadata server or the private network connecting the nodes in the cluster. Avoid unnecessary metadata traffic.

If the amount of free memory is insufficient, a node may experience delays in heartbeating and as a result will be kicked out of the CXFS membership. To observe the amount of free memory in your system, use the `osview` tool.

See also "Out of Logical Swap Space" on page 462.

Reboot Before Changing Node ID or Cluster ID

If you want redefine a node ID or the cluster ID, you must first remove the current cluster definition for the node, then reboot. The problem is that the kernel still has the old values, which prohibits a CXFS membership from forming. However, if you perform a reboot first, it will clear the original values and you can then redefine the node or cluster ID.

Therefore, if you use `cdbreinit` on a node to recreate the cluster database, you must reboot it before changing the node IDs or the cluster ID. See "Recreating the Cluster Database" on page 489.

Restart CXFS on a Node after an Administrative CXFS Stop

If you perform an administrative CXFS stop (forced CXFS shutdown) on a node, you must perform an administrative CXFS start on that node before it can return to the cluster. If you do this while the database still shows that the node is in a cluster and is activated, the node will restart the CXFS membership daemon. Following a forced CXFS shutdown, the node can be prevented from restarting the CXFS membership daemon when CXFS is restarted by stopping CXFS services. (A forced CXFS shutdown alone does not stop CXFS services. A forced CXFS shutdown stops only the kernel membership daemon. Stopping CXFS services disables the node in the cluster database.)

For example, enter the following on the local node you wish to start:

```
cxfs_admin:clustername> disable node:nodename  
cxfs_admin:clustername> enable node:nodename
```

See also "Forced CXFS Shutdown: Revoke Membership of Local Node" on page 337.

Restart the Cluster In an Orderly Fashion

SGI recommends that you do the following to start the cluster in an orderly fashion if you have taken the entire cluster down for maintenance or because of server instability. This procedure assumes all nodes have been disabled.

1. Start CXFS services (using the CXFS GUI or `cxfs_admin`) for the potential metadata servers. Do the following for each potential metadata server if you are using the `cxfs_admin` command:

```
cxfs_admin:clustername> enable node:nodename
```

2. Start CXFS services on the client-only tiebreaker node.
3. Start CXFS services on the remaining client-only nodes.

Disable Reset Capability If You Remove Reset Lines

When reset is enabled, CXFS requires a `reset successful` message before it moves the metadata server. Therefore, if you have the reset capability enabled and you must remove the reset lines for some reason, you must also disable the reset capability. See "Modify a Node Definition with the GUI" on page 217, or "Create or Modify a Node with `cxfs_admin`" on page 270.

Avoid Performance Problems with Unwritten Extent Tracking and Exclusive Write Tokens

This section discusses performance problems with unwritten extent tracking and exclusive write tokens.

Unwritten Extent Tracking

When you define a filesystem, you can specify whether unwritten extent tracking is on (`unwritten=1`) or off (`unwritten=0`); it is on by default.

In most cases, the use of unwritten extent tracking does not affect performance and you should use the default to provide better security.

However, unwritten extent tracking can affect performance when **both** of the following are true:

- A file has been preallocated
- These preallocated extents are written for the first time with records smaller than 4 MB

For optimal performance with CXFS when **both** of these conditions are true, it may be necessary to build filesystems with `unwritten=0` (off).

Note: There are security issues with using `unwritten=0`. For more information, see the *IRIX Admin: Disks and Filesystems*.

Exclusive Write Tokens

For proper performance, CXFS should not obtain exclusive write tokens. Therefore, use the following guidelines:

- Preallocate the file.
- Set the size of the file to the maximum size and do not allow it to be changed, such as through truncation.
- Do not append to the file. (That is, `O_APPEND` is not true on the open.)
- Do not mark an extent as `written`.
- Do not allow the application to do continual `preallocation` calls.

If the guidelines are followed and there are still performance problems, you may find useful information by running the `icrash stat` command before, halfway through, and after running the MPI job. For more information, see the `icrash` man page.

Note: You must run the `sial` scripts version of `icrash` commands. See "Kernel Status Tools" on page 434.

Use the Appropriate Version of `lcrash` for SGI ProPack

If you want to use `lcrash` for troubleshooting on an SGI ProPack node, you must use the version of `lcrash` that is available from Supportfolio. Use the `-x` option to load the CXFS kerntypes:

```
# lcrash -x /boot/sgi-cxfs-kerntypes-kernelversion-architecture
```

Note: Do not use the version of `lcrash` that is shipped with SLES 9.

Disk Layout Optimization for Approved Media Customers

Approved media customers can use the XFS `filestreams` mount option with CXFS to maximize the ability of storage to support multiple real-time streams of video data. It is appropriate for workloads that generate many files that are created and accessed in a sequential order in one directory.



Caution: SGI must validate that your RAID model and RAID configuration can support the use of the `filestreams` mount option to achieve real-time data transfer and that your application is appropriate for its use. Use of this feature is complex and is reserved for designs that have been approved by SGI

The `filestreams` mount option changes the behavior of the XFS allocator in order to optimize disk layout. It selects an XFS disk block allocation strategy that does the following:

- Identifies streams writing into the same directory and locks down a region of the filesystem for that stream, which prevents multiple streams from using the same allocation groups
- Allocates the file data sequentially on disk in the order that the files are created, space permitting
- Uses different regions of the filesystem for files in different directories

Using the `filestreams` mount option can improve both bandwidth and latency when accessing the files because the RAID will be able to access the data in each directory sequentially. Therefore, multiple writers may be able to write into the same filesystem without interleaving file data on disk. Filesystem can be filled up to

approximately 94% before performance degrades. Deletion of projects does not fragment filesystem, therefore there is no need to rebuild filesystem after each project.

You can safely enable the `filestreams` mount option on an existing filesystem and later disable it without affecting compatibility. (The mount option affects where data is located in the filesystem; it does not change the format of the filesystem.) However, you may not get the full benefit of `filestreams` due to preexisting filesystem fragmentation.

For more information, contact your SGI Support person.

Initial Setup of the Cluster



Caution: CXFS is a complex product. To ensure that it is installed and configured in an optimal manner, you **must** purchase initial setup services from SGI. You should read through the following chapters, before attempting to install and configure a CXFS cluster:

- Chapter 1, "Introduction to CXFS" on page 1
- Chapter 2, "SGI RAID for CXFS Clusters" on page 63
- Chapter 3, "Switches" on page 67
- Chapter 4, "CXFS License Keys" on page 75
- Chapter 5, "Preinstallation Steps" on page 93
- Chapter 6, "IRIX CXFS Installation" on page 101 and/or Chapter 7, "SGI ProPack CXFS Installation" on page 113
- Chapter 8, "Postinstallation Steps" on page 127
- Chapter 9, "Best Practices" on page 139
- Chapter 10, "Initial Setup of the Cluster" on page 163

This chapter provides recommendations and a summary of the basic steps required to initially configure a cluster. It contains the following:

- "Preliminary Cluster Configuration Steps" on page 164
- "Initial Setup Using One of the Configuration Tools" on page 167
- "Configuring a Large Cluster" on page 176
- "Testing the System" on page 178

You should also refer to the information in "Configuration Best Practices" on page 139 and you may wish to use the worksheet provided in Appendix I, "Initial

Configuration Checklist" on page 661. If you are converting from an existing FailSafe cluster, see "Set Up an Existing FailSafe Cluster for CXFS with the GUI" on page 205.

This chapter points to detailed descriptions in the task reference chapters and in the *XVM Volume Manager Administrator's Guide*.

For information about licenses, see Chapter 4, "CXFS License Keys" on page 75.

Preliminary Cluster Configuration Steps

Complete the following steps to ensure that you are ready to configure the initial cluster:

- "Verify the License" on page 164
- "Verify that the Cluster Daemons are Running" on page 164
- "Gather the Required Information" on page 165
- "Configure for nsd Use (Optional)" on page 166
- "Verify that the chkconfig Arguments are On" on page 166

During the course of configuration, you will see various information-only messages in the log files. See "Normal Messages" on page 458.

Verify the License

Verify that you have the appropriate CXFS licenses by using the `-d` option to the `cxfslicense` command on server-capable nodes. See "Verifying the License Keys with `cxfslicense` on Server-Capable Nodes" on page 84.

Verify that the Cluster Daemons are Running

When you **first install** the software, the following daemons should be running on all administration nodes:

- `fs2d`
- `cmond`
- `cad`

- `crsd`

To determine which daemons are running on an administration node, enter the following:

- IRIX:

```
irix# /etc/init.d/cluster status
fs2d is running
cmond is running
cad is running
crsd is running
```

- SGI ProPack:

```
[root@linux64 root]# /etc/init.d/cxfs_cluster status
fs2d is running.
cmond is running.
cad is running.
crsd is running.
```

If you do not see these processes on an administration node, go to the logs to see what the problem might be. Then restart the daemons by entering the following:

- IRIX:

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

- SGI ProPack:

```
[root@linux64 root]# /etc/init.d/cxfs_cluster start
```

The `cxfs_client` daemon should be running on a client-only node. If it is not, enter the following:

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

For more information, see "Stopping and Restarting Cluster Administration Daemons" on page 489 and "Daemons" on page 495.

Gather the Required Information

You should know the fully qualified hostname of the machine from which you will do CXFS administration, which should be the first node you define in the cluster database. If you use `cxfs_admin` (see "Initial Setup Using One of the Configuration

Tools" on page 167), you should use the hostname when defining the first node in the pool. (This information is automatically supplied for you in the CXFS GUI.)

You should also know the IP addresses and hostnames of the other machines that will form the cluster and the name by which want to refer to the cluster.

Configure for `nsd` Use (Optional)

If your system uses `nsd` for hostname resolution, you must configure your system so that local files are accessed before the network information service (NIS) or the domain name service (DNS).

Verify that the `chkconfig` Arguments are On

Ensure that the appropriate `chkconfig` arguments are on. For more information, see "CXFS `chkconfig` Arguments" on page 320.

IRIX `chkconfig` Verification

For an IRIX node, ensure that `chkconfig` displays the following

```
irix# chkconfig | grep cluster
      cluster                on
      cxfs_cluster           on
```

If it does not, set the flags to on and reboot. For example:

```
irix# /etc/chkconfig cluster on
irix# /etc/chkconfig cxfs_cluster on
irix# init 6
```

Or:

```
irix# init 1
irix# /etc/chkconfig cluster on
irix# /etc/chkconfig cxfs_cluster on
irix# init 2
```

SGI ProPack chkconfig Verification

For an SGI ProPack node, use `--list` option to `chkconfig` to verify that the `chkconfig` names are set to `on` for the site's normal run levels. For example, if the normal run levels were 3 and 5:

```
[root@linux root]# chkconfig --list | grep cxfs
cxfs_cluster    0:off  1:off  2:off  3:on   4:off  5:on   6:off
cxfs            0:off  1:off  2:off  3:on   4:off  5:on   6:off
```

Note: Your site's normal run levels may differ.

If the normal run levels are set to `off`, set them to `on` and reboot. For example:

```
[root@linux root]# chkconfig cxfs_cluster on
[root@linux root]# chkconfig cxfs on
[root@linux root]# reboot
```

Initial Setup Using One of the Configuration Tools

You can create the cluster and its components using any one of the following tools, which provide similar functionality:

- "Initial Setup with the CXFS GUI" on page 167
 - "Initial Setup with the `cxfs_admin` Command" on page 172
-



Caution: You should only use one configuration tool at a time to make changes.

The following procedures provide an overview of the basic steps to set up a cluster. You will first define a server-capable node from which you perform administrative tasks, and then the other components of the cluster.

Initial Setup with the CXFS GUI

Note: For complete details about using the GUI, see "CXFS Tools Overview" on page 42 and Chapter 11, "Reference to GUI Tasks" on page 183.

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

- "Start the GUI" on page 168
- "Set Up a New Cluster with the GUI" on page 169
- "Set Up a New CXFS Filesystem with the GUI" on page 171

The CXFS administration node to which you connect the GUI affects your view of the cluster. You should wait for a change to appear in the view area before making another change; the change is not guaranteed to be propagated across the cluster until it appears in the view area. You should only make changes from one instance of the GUI at any given time; changes made by a second GUI instance may overwrite changes made by the first instance.

Start the GUI

Start the CXFS Manager by entering the following:

```
# /usr/sbin/cxfsmgr
```

You can also start the GUI from your web browser on a Microsoft Windows, Linux, or other platform. To do this, enter `http://server/CXFSManager/` (where *server* is the name of a CXFS administration node in the pool) and press **Enter**. At the resulting webpage, click the CXFS Manager icon. This method of launching CXFS Manager requires you to have enabled Java in your browser's preferences and have installed the appropriate Java plug-in. (After installing the plug-in, you must close any existing Java windows and restart your browser.) The CXFS administration node must be running a web server, such as Apache, and have the following software installed:

- IRIX: `sysadm_cxfs.sw.web`
- SGI ProPack: `cxfs-sysadm_cxfs-web`

Note: If you load the GUI using Netscape on IRIX and then switch to another page in Netscape, CXFS Manager GUI will not operate correctly. To avoid this problem, leave the CXFS Manager GUI web page up and open a new Netscape window if you want to view another page.

There are other methods of starting the GUI. For more information, see "Starting the GUI" on page 184.

Supply the name of the CXFS administration node you wish to connect to and the root password.

Figure 10-1 shows an example of the CXFS Manager window.

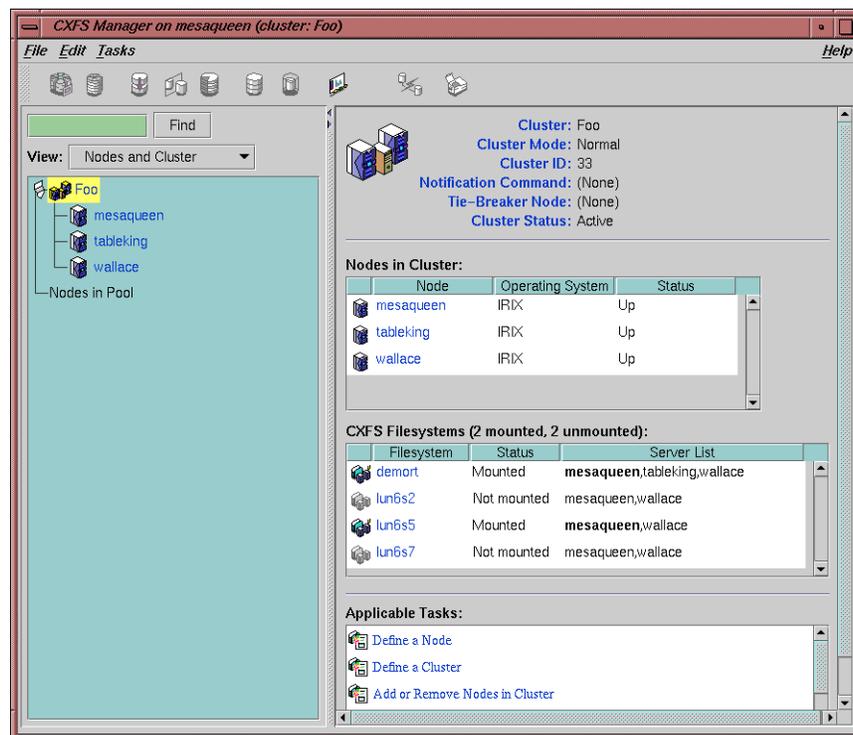


Figure 10-1 CXFS Manager

Set Up a New Cluster with the GUI

Within the CXFS tasks, you can click 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**.

Note: To specify reset method that uses Intelligent Platform Management Interface (IPMI) and baseboard management controller (BMC), you must use the `cxfs_admin` configuration tool. See "Create or Modify a Node with `cxfs_admin`" on page 270.

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

Do the following:

1. Click **Define a Node** to define the CXFS administration node to which you are connected. See "Define a Node with the GUI" on page 208.

Note: If you attempt to define a cluster or other object before the local node has been defined, you will get an error message that says:

No nodes are registered on *servername*. You cannot define a cluster until you define the node to which the GUI is connected. To do so, click "Continue" to launch the "Set Up a New Cluster" task.

2. (Optional) **After** the first node icon appears in the view area on the left, click step 2, **Define a Node**, to define the other nodes in the cluster. To use private network failover, you must use the `cxfs_admin` command's `create failover_net` command to specify the network and mask; see "Network Failover Tasks with `cxfs_admin`" on page 295. See "Define a Node with the GUI" on page 208.

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

Repeat this step for each node. For large clusters, define only the administration nodes first; see "Configuring a Large Cluster" on page 176.

3. Click **Define a Cluster** to create the cluster definition. See "Define a Cluster with the GUI" on page 224. Verify that the cluster appears in the view area. Choose **View: Nodes and Cluster**.

4. After the cluster icon appears in the view area, click **Add/Remove Nodes in Cluster** to add the nodes to the new cluster. See "Add or Remove Nodes in the Cluster with the GUI" on page 216.

Click **Next** to move to the second screen of tasks.

5. (*Optional*) Click on **Test Connectivity** to verify that the nodes are physically connected. See "Test Node Connectivity with the GUI" on page 223. (This test requires the proper configuration; see "IRIX Modifications for CXFS Connectivity Diagnostics" on page 111, "SGI ProPack Modifications for CXFS Connectivity Diagnostics" on page 124.)
6. If you are using I/O fencing, define the switch in the cluster; see the release notes for supported switches. I/O fencing is required for nodes without system controllers; see "Requirements" on page 38.
7. Click **Start CXFS Services**. See "Start CXFS Services with the GUI" on page 227.
8. Click **Close**. Clicking on **Close** exits the task; it does not undo the task.

Set Up a New CXFS Filesystem with the GUI

Note: Within the CXFS tasks, you can click 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 **OK**.

The **Set Up a New CXFS Filesystem** task 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.

Do the following:

1. Click **Start CXFS Services** if the services have not been started already. (The current status is displayed beneath the task link.) See "Start CXFS Services with the GUI" on page 227.
2. Click **Label Disks**.

Note: The disk must be initialized before being labeled. If your disk has not been initialized during factory set-up, use the IRIX `fx` command or SGI ProPack `fdisk` command to initialize the disk.

For information about XVM tasks, see the *XVM Volume Manager Administrator's Guide*.

3. Create slices, which define the physical storage, on the labeled disk. Click **Slice Disks**.
4. Create the type of filesystem you want: stripe, mirror, or concat.
5. Click **Make the Filesystem**. If you do not want to use the default options, click **Specify Sizes** and go to the next page. For more information, see the `mkfs` man page, the *IRIX Admin: Disks and Filesystems* guide, and the *XVM Volume Manager Administrator's Guide*.
6. Click **Define a CXFS Filesystem**. This task lets you define a new filesystem, set the ordered list of potential metadata servers, and set the list of client nodes for the filesystem. See "Define CXFS Filesystems with the GUI" on page 241.
7. Click **Mount a CXFS Filesystem**. This task lets you mount the filesystem on all nodes in the cluster. See "Mount CXFS Filesystems with the GUI" on page 245.

Repeat these steps for each filesystem.

Initial Setup with the `cxfs_admin` Command

Note: For the initial installation, SGI highly recommends that you use the GUI guided configuration tasks. See "Initial Setup with the CXFS GUI" on page 167. For complete details about using `cxfs_admin`, see "CXFS Tools Overview" on page 42 and Chapter 12, "Reference to `cxfs_admin` Tasks" on page 255.

You can perform configuration with `cxfs_admin` using normal mode (in which you specify each command and attribute) or in prompting mode, in which `cxfs_admin` asks you for the information it requires.

To initially configure the cluster with `cxfs_admin`, do the following (line breaks shown here for readability). A simple example of prompting mode follows the steps.

1. "Preliminary Cluster Configuration Steps" on page 164.
2. Initialize the cluster database and start `cxfs_admin`:

```
# cxfs_admin -s
```

3. Define the cluster name, where *clustername* is the logical name of the cluster:

```
cxfs_admin> create cluster name=clustername
```

For example:

```
cxfs_admin> create cluster name=mycluster
```

4. Create the first server-capable node (normally the node on which you are currently running *cxfs_admin*). (You do not need to specify the node type because it must be *server_admin*.) If you use prompting mode, the name of the local node is used as a default for *name*.



Caution: It is critical that you enter the primary hostname for the first node defined in the pool.

```
cxfs_admin> create node name=server_capable_hostname private_net=private_IPaddress
```

For example:

```
cxfs_admin> create node name=server1 private_net=10.11.20.114
```

5. Exit *cxfs_admin* and restart the CXFS cluster services:

- IRIX:

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

- SGI ProPack:

```
# /etc/init.d/cxfs stop
# /etc/init.d/cxfs_cluster stop
# /etc/init.d/cxfs_cluster start
# /etc/init.d/cxfs start
```

6. Restart *cxfs_admin*:

```
# cxfs_admin
```

Note: If you have multiple clusters using the same public network as the backup CXFS metadata network, use the `-i` option to identify the cluster name:

```
# cxfs_admin -i mycluster
```

7. (Optional) Create the failover networks:

```
cxfs_admin:cluster> create failover_net network=IPaddress1 mask=netmask
cxfs_admin:cluster> create failover_net network=IPaddress2 mask=netmask
```

For example:

```
cxfs_admin:mycluster> create failover_net network=192.168.0.0 mask=255.255.255.0
cxfs_admin:mycluster > create failover_net network=10.0.0.0 mask=255.255.255.0
```

8. Create the switches:

```
cxfs_admin:cluster> create switch name=switch_hostname [vendor=brocade|qlogic]
[user=username password=password]
```

For example:

```
cxfs_admin:mycluster> create switch name=myswitch vendor=qlogic
```

9. Create other CXFS nodes as required:

```
cxfs_admin:mycluster> create node name=nodename os=OSType private_net=IPaddress
[type=server_admin|client_only]
```

For example, for a server-capable node:

```
cxfs_admin:mycluster> create node name=server2 os=Linux private_net=10.11.20.115 \
type=server_admin
```

For example, for a client-only node, in this case running Windows:

```
cxfs_admin:mycluster> create node name=client1 os=Windows private_net=10.11.20.116 \
```

10. (Optional) Define one of the client-only nodes as the CXFS tiebreaker if using multiple server-capable nodes:

```
cxfs_admin:cluster> modify clustername tiebreaker=client_only_nodename
```

For example:

```
cxfs_admin:mycluster> modify mycluster tiebreaker=client1
```

11. Obtain a shell window for one of the CXFS administration nodes in the cluster and use the IRIX `fx(1M)` command or the Linux `parted(8)` command to create a volume header on the disk drive. For information, see the man pages, *IRIX Admin: Disks and Filesystems*, and *Linux Configuration and Operations Guide*.
12. Create the XVM logical volumes. In the shell window, use the `xvm` command line interface. For information, see the *XVM Volume Manager Administrator's Guide*.
13. Make the XFS filesystems. In the shell window, use the `mkfs` command. For information, see the *XVM Volume Manager Administrator's Guide* and *IRIX Admin: Disks and Filesystems*.
14. Create the CXFS filesystems:

```
cxfs_admin:cluster> create filesystem name=XVMvolume [mountpoint=path]
[options=mount_options]
```

For example:

```
cxfs_admin:cluster> create filesystem name=cxfsvoll
```

15. (Optional) Create private network failover:

```
cxfs_admin:cluster> network=IPaddress mask=NetMask
```

For example, to create two private networks, one on the 192.168.0.x and the other on the 10.0.0.x subnets:

```
cxfs_admin:mycluster> create failover_net network=192.168.0.0 mask=255.255.255.0
cxfs_admin:mycluster > create failover_net network=10.0.0.0 mask=255.255.255.0
```

16. View the cluster status:

```
cxfs_admin:cluster> status
```

Following is a simple example using prompting mode:

```
cxfsopus14:~ # /usr/cluster/bin/cxfs_admin -s
Connecting to the local CXFS server...
cxfs_admin:(no cluster defined)> create cluster
Specify the attributes for create cluster:
name? mycluster
```

```
cxfs_admin:mycluster> create node
Specify the attributes for create node:
  name? cxfsopus14
  private_net? 10.11.20.114
Node "cxfsopus14" has been created, waiting for it to join the cluster...
Please restart all cxfs and cluster services on the server "cxfsopus14"
to make
it join the cluster.
cxfs_admin:mycluster> create filesystem
Specify the attributes for create filesystem:
  name? thump
  options?
  forced_unmount? false
  mountpoint? /mnt/thump
  mounted? true
Filesystem "thump" has been created, waiting for it to be mounted on all
assigned nodes...
```

For more information, see Chapter 12, "Reference to cxfs_admin Tasks" on page 255 and the help command within cxfs_admin.

Configuring a Large Cluster

When configuring a large cluster, you should ensure that a small cluster containing just the server-capable administration nodes is fully functional before adding client-only nodes. By building up the cluster with client-only nodes in small groups, you will minimize concurrent operational issues and use the database most efficiently.

Do the following:

1. Create the initial cluster with just the server-capable nodes and test it:
 - a. Define all of the server-capable administration nodes.
 - b. Define the cluster.
 - c. Add all of the server-capable administration nodes to the cluster.
 - d. Create the filesystems as described in "Set Up a New CXFS Filesystem with the GUI" on page 171.
 - e. Verify that the nodes are all part of the cluster membership and that the filesystems are mounted and fully functional.
2. Add the client-only nodes to the database:
 - a. Define **all** client-only nodes.
 - b. Add **all** client-only nodes to the cluster.
3. Gradually build up the functional cluster with subsets of client-only nodes:
 - a. Start CXFS services on a **subset** of four client-only nodes.
 - b. Ensure that the nodes are part of the cluster membership and that the filesystems are fully functional.
4. Repeat step 3 as needed to complete the cluster membership.

Following is an example `cxfs_admin` script to configure a cluster. The first node line creates the first server-capable node; you can copy and repeat the second node line for each remaining server-capable or client-only node in the cluster:

```
create cluster name=clustername
create node name=nodename private_net=IPaddress
create node name=nodename os=OS private_net=IPaddress [copy and repeat]
create filesystem name=filesystemname forced_unmount=false mountpoint=/mnt/nodename mounted=true [copy and repeat]
```

Following is an example for configuring a one-node cluster that can be copied and repeated for the number of nodes required:

```
create cluster name=clustername
create node name=nodename private_net=IPaddress
create filesystem name=filesystemname forced_unmount=false mountpoint=/mnt/nodename
mounted=true
```

Testing the System

This section discusses the following:

- "Private Network Interface" on page 178
- "System Reset Connection for CXFS Administration Nodes" on page 178
- "Testing Serial Connectivity for the L2 on Altix 350 Systems" on page 180

Private Network Interface

For each private network on each node in the pool, enter the following, where *nodeIPAddress* is the IP address of the node:

```
# ping -c 3 nodeIPAddress
```

Typical ping output should appear, such as the following:

```
PING IPaddress (190.x.x.x: 56 data bytes
64 bytes from 190.x.x.x: icmp_seq=0 ttl=254 time=3 ms
64 bytes from 190.x.x.x: icmp_seq=1 ttl=254 time=2 ms
64 bytes from 190.x.x.x: icmp_seq=2 ttl=254 time=2 ms
```

If ping fails, follow these steps:

1. Verify that the network interface was configured up by using `ifconfig`. For example:

```
# ifconfig ec3
ec3: flags=c63<UP,BROADCAST,NOTRAILERS,RUNNING,FILTMULTI,MULTICAST>
inet 190.x.x.x netmask 0xffffffff broadcast 190.x.x.x
```

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

2. Verify that the cables are correctly seated.

Repeat this procedure on each node.

System Reset Connection for CXFS Administration Nodes

To test the system reset connections, do the following:

1. Ensure that the nodes and the serial port multiplexer are powered on.

2. Ensure that the `MODULES_LOADED_ON_BOOT` variable in `/etc/sysconfig/kernel` contains the `ioc4_serial` module. This module must be loaded in order for the devices to be present.

3. Start the `cmgr` command on one of the CXFS administration nodes in the pool:

```
# cmgr
```

4. Stop CXFS services on the entire cluster:

```
stop cx_services for cluster clustername
```

For example:

```
cmgr> stop cx_services for cluster cxfs6-8
```

Wait until the node has successfully transitioned to inactive state and the CXFS processes have exited. This process can take a few minutes.

5. Test the serial connections by entering one of the following:

- To test the whole cluster, enter the following:

```
test serial in cluster clustername
```

For example:

```
cmgr> test serial in cluster cxfs6-8
```

```
Status: Testing serial lines ...
```

```
Status: Checking serial lines using crsd (cluster reset services) from node cxfs8
```

```
Success: Serial ping command OK.
```

```
Status: Checking serial lines using crsd (cluster reset services) from node cxfs6
```

```
Success: Serial ping command OK.
```

```
Status: Checking serial lines using crsd (cluster reset services) from node cxfs7
```

```
Success: Serial ping command OK.
```

```
Notice: overall exit status:success, tests failed:0, total tests executed:1
```

- To test an individual node, enter the following:

```
test serial in cluster clustername node machinename
```

For example:

```
cmgr> test serial in cluster cxfs6-8 node cxfs7
Status: Testing serial lines ...
Status: Checking serial lines using crsd (cluster reset services) from node cxfs6
Success: Serial ping command OK.
```

Notice: overall exit status:success, tests failed:0, total tests executed:1

- To test an individual node using just a ping, enter the following:

```
admin ping node nodename
```

For example:

```
cmgr> admin ping node cxfs7
```

```
ping operation successful
```

6. If a command fails, make sure all the cables are seated properly and rerun the command.
7. Repeat the process on other nodes in the cluster.

Testing Serial Connectivity for the L2 on Altix 350 Systems

You can use the `cu(1)` command to test the serial reset lines if you have installed the `uucp` RPM.

The `cu` command requires that the device files be readable and writable by the user `uucp`. The command also requires the `/var/lock` directory be writable by group `uucp`.

Perform the following steps:

1. Change ownership of the serial devices so that they are in group `uucp` and owned by user `uucp`.

Note: The ownership change may not be persistent across reboots.

For example, suppose you have the following TTY devices on the IO10:

```
# ls -l /dev/ttyIOC*
crw-rw---- 1 root uucp 204, 50 Sep 15 16:20 /dev/ttyIOC0
crw-rw---- 1 root uucp 204, 51 Sep 15 16:20 /dev/ttyIOC1
crw-rw---- 1 root uucp 204, 52 Sep 15 16:20 /dev/ttyIOC2
crw-rw---- 1 root uucp 204, 53 Sep 15 16:20 /dev/ttyIOC3
```

To change ownership of them to uucp, you would enter the following:

```
# chown uucp.uucp /dev/ttyIOC*
```

2. Determine if group uucp can write to the /var/lock directory and change permissions if necessary.

For example, the following shows that group uucp cannot write to the directory:

```
# ls -ld /var/lock
drwxr-xr-t 5 root uucp 88 Sep 19 08:21 /var/lock
```

The following adds write permission for group uucp:

```
# chmod g+w /var/lock
```

3. Join the uucp group temporarily, if necessary, and use cu to test the line.

For example:

```
# newgrp uucp
# cu -l /dev/ttyIOC0 -s 38400
Connected
nodeA-001-L2>cfg
L2 192.168.0.1: - 001 (LOCAL)
L1 192.0.1.133:0:0 - 001c04.1
L1 192.0.1.133:0:1 - 001i13.1
L1 192.0.1.133:0:5 - 001c07.2
L1 192.0.1.133:0:6 - 001i02.2
```

For more information, see the cu(1) man page and the documentation that comes with the uucp RPM.

Reference to GUI Tasks

This chapter discusses the CXFS Manager graphical user interface (GUI). It contains detailed information about CXFS tasks and an overview of XVM tasks. (For details about XVM tasks, see the *XVM Volume Manager Administrator's Guide*.)

This chapter contains the following sections:

- "GUI Overview" on page 183
- "Guided Configuration Tasks" on page 205
- "Node Tasks with the GUI" on page 207
- "Cluster Tasks with the GUI" on page 223
- "Cluster Services Tasks with the GUI" on page 227
- "Switches and I/O Fencing Tasks with the GUI" on page 233
- "Filesystem Tasks with the GUI" on page 237
- "Privileges Tasks with the GUI" on page 249

Note: CXFS requires a license key to be installed on each server-capable node. If you install the software without properly installing the license key, you will get an error and will not be able to use the CXFS Manager GUI. For more information about licensing, see Chapter 4, "CXFS License Keys" on page 75.

GUI Overview

The GUI lets you set up and administer CXFS filesystems and XVM logical volumes. It also provides icons representing status and structure.

This section provides an overview of the GUI:

- "Starting the GUI"
- "GUI Windows" on page 189

- "GUI Features" on page 191
- "Key to Icons and States" on page 201

Note: CXFS is incompatible with the Red Hat cluster manager available in the Red Hat Advanced Server product.

Starting the GUI

There are several methods to start the GUI and connect to a node.

Starting the GUI on IRIX

To start the GUI, use one of the following methods:

- On an IRIX system where the CXFS GUI-client software (`sysadm_cxfs.sw.client`) and desktop support software (`sysadm_cxfs.sw.desktop`) are installed, do one of the following:

Note: SGI does not recommend this method across a wide-area network (WAN) or virtual private network (VPN), or if the IRIX system has an R5000 or earlier CPU and less than 128-MB memory.

- Enter the following command line:

```
# /usr/sbin/cxfsmgr
```

(The `cxdetail` and `cxtask` commands on IRIX administration nodes perform the identical function as `cxfsmgr`; these command names are kept for historical purposes.)

- Choose the following from the Toolchest:

System
 > **CXFS Manager**

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 &
```

If you are using WAN or VPN, see "Running the Web-based Version" on page 185.

Starting the GUI on SGI ProPack

To start the GUI on an SGI ProPack system where the CXFS GUI-client software (`cxfs-sysadm_cxfs-client`) is installed, do the following:

1. Obtain and install the J2SE 1.4.2 (latest patch) software available from <http://java.sun.com>
2. Enter the following command line:

```
# /usr/sbin/cxfsmgr
```

Running the Web-based Version

If you want to use a web-based version of the GUI, do the following:

1. Ensure that the following subsystems are installed on the CXFS administration nodes that you will connect to (by means of a Java-enabled web browser running on any platform) for performing administrative operations:

- IRIX:

```
sysadm_xvm.sw.web
sysadm_cxfs.sw.web
```

- SGI ProPack:

```
cxfs-sysadm_xvm-web
cxfs-sysadm_cxfs-web
```

These subsystems are part of the software normally installed with CXFS.

2. Ensure that a Web server is installed and running:

- IRIX:
 - `sgi_apache.sw.server` (preferred)
 - `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.

- SGI ProPack requires an apache server

3. On a PC, install the Java2 v1.4.2 or v1.5 plug-in.

On an IRIX machine that launches the GUI client from a web browser that supports Java, install the `java_plugin` subsystem from the IRIX 6.5.x CD. This is the Runtime Plug-in for IRIX, Java Edition 1.4.1, which supports JRE 1.4.1. (However, launching the GUI from a web browser is not the recommended method on IRIX. On IRIX, running the GUI client from the desktop is preferred.)

4. Add the following to your `httpd.conf` file:

```
<Location "/CXFSManager">
    Options Includes ExecCGI FollowSymLinks
    DirectoryIndex index.html index.shtml
</Location>
```

5. Close all browser windows and restart the browser.

6. Enter the URL `http://server/CXFSManager/` where *server* is the name of a CXFS administration node in the pool

7. At the resulting webpage, click the **CXFS Manager** icon.

Note: This method can be used on IRIX systems, but it is not the preferred method unless you are using WAN or VPN. If you load the GUI using Netscape on IRIX and then switch to another page in Netscape, CXFS Manager GUI will not operate correctly. To avoid this problem, leave the CXFS Manager GUI web page up and open a new Netscape window if you want to view another web page.

Running as a Non-Root User on IRIX

Running the CXFS Manager graphical user interface (GUI) from a login other than `root` requires the `sysadmdesktop` package, which is installed by default when you install IRIX. This package provides commands that allow you to give users privileges, including the privileges required to run the CXFS commands. `sysadmdesktop` (located on the *Applications CD 1 of 2 for 6.5.x*) installs the following subsystems:

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

Running the GUI from an IRIX Desktop Outside the Cluster

If you want to run the GUI client from an IRIX desktop outside of the cluster, install the following subsystems on that machine:

```
java2_eoe.sw
java2_eoe.sw32
sysadm_base.man
sysadm_base.sw.client
sysadm_cluster.sw.client
sysadm_cxfs.man
sysadm_cxfs.sw.client
sysadm_cxfs.sw.desktop
sysadm_xvm.sw.client
sysadm_xvm.sw.desktop
```



Caution: The GUI on IRIX only operates with Java2 v1.4.1 Execution Environment (Sun JRE v1.4.1). This is the version of Java that is provided with the supported IRIX 6.5.x release.

The SGI website also contains Java1. However, you cannot use this version of Java with the GUI. Using a Java version other than 1.4.1 will cause the GUI to fail.

Summary of GUI Platforms

Table 11-1 describes the platforms where the GUI may be started, connected to, and displayed.

Table 11-1 GUI Platforms

GUI Mode	Where You Start the GUI	Where You Connect the GUI	Where the GUI Displays
cxfsmgr	Any IRIX system (such as an SGI 2000 series or SGI O2 workstation) with <code>sysadm_cxfs.sw.client</code> and <code>sysadm_cxfs.sw.desktop</code> software installed An SGI ProPack system with <code>cxfs-sysadm_cxfs-client</code> installed	The CXFS administration node in the pool that you want to use for cluster administration	The system where the GUI was invoked
Toolchest	Any IRIX system (such as an SGI 2000 series or SGI O2 workstation) with <code>sysadm_cxfs.sw.client</code> and <code>sysadm_cxfs.sw.desktop</code> software installed	The CXFS administration node in the pool that you want to use for cluster administration	The system where the GUI was invoked
Web	Any system with a web browser and Java2 1.4.1 or 1.4.2 plug-in installed and enabled	The CXFS administration node in the pool that you want to use for cluster administration	The same system with the web browser

Logging In

To ensure that the required GUI 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 other user using the GUI privilege tasks; see "Privileges Tasks with the GUI" on page 249. (Under IRIX, this functionality is also available with the Privilege Manager, part of the IRIX Interactive Desktop System Administration `sysadmdesktop` product. For more information, see the *Personal System Administration Guide*.)

A dialog box will appear prompting you to log in to a CXFS host. You can choose one of the following connection types:

- **Local** runs the server-side process on the local host instead of going over the network
- **Direct** creates a direct socket connection using the `tcpmux` TCP protocol (`tcpmux` must be enabled)

- **Remote Shell** connects to the server via a user-specified command shell, such as `rsh` or `ssh`. For example:

```
ssh -l root servername
```

Note: For secure connection, choose **Remote Shell** and type a secure connection command using a utility such as `ssh`. Otherwise, the GUI will not encrypt communication and transferred passwords will be visible to users of the network.

- **Proxy** connects to the server through a firewall via a proxy server

Making Changes Safely

Do not make configuration changes on two different administration nodes in the pool simultaneously, or use the CXFS GUI, `cxfs_admin`, `cmgr`, and `xvm` commands simultaneously to make changes. You should run one instance of the `cxfs_admin` command, `cmgr` command, or the CXFS GUI on a single administration node in the pool when making changes at any given time. However, you can use any node in the pool when requesting status or configuration information. Multiple CXFS Manager windows accessed via the **File** menu are all part of the same application process; you can make changes from any of these windows.

The CXFS administration node to which you connect the GUI affects your view of the cluster. You should wait for a change to appear in the *view area* before making another change; the change is not guaranteed to be propagated across the cluster until it appears in the view area. (To see the location of the view area, see Figure 11-1 on page 190.) The entire cluster status information is sent to every CXFS administration node each time a change is made to the cluster database.

GUI Windows

Figure 11-1 shows the **CXFS Manager** window displaying information for a specific component in the *details area*. For information about using the *view area* to monitor status and an explanation of the icons and colors, see "Cluster Status" on page 402.

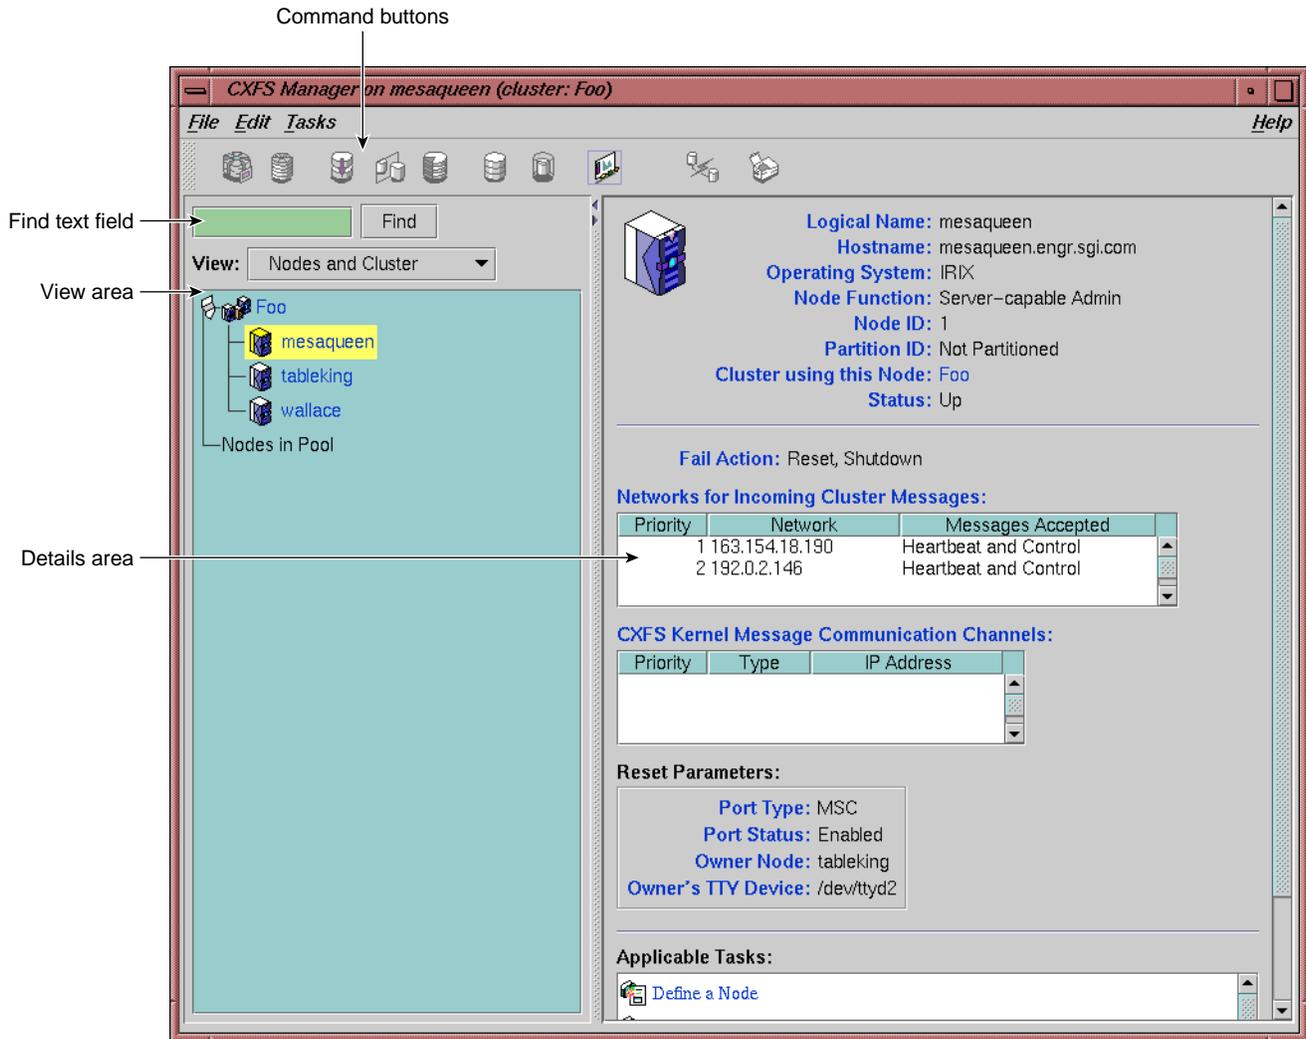


Figure 11-1 CXFS Manager GUI Showing Details for a Node

Figure 11-2 shows an example of the pop-up menu of applicable tasks that appears when you click the right mouse button on a selected item; in this example, clicking on the node name *trinity* displays a list of applicable tasks.

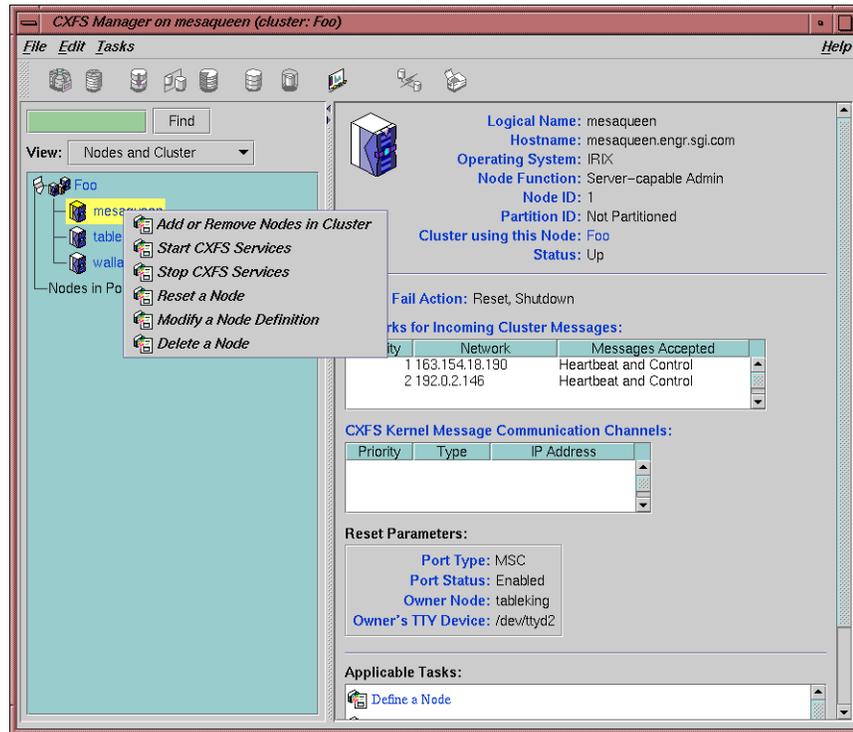


Figure 11-2 Pop-up Menu that Appears After Clicking the Right Mouse Button

GUI Features

The **CXFS Manager** GUI allows you to administer the entire CXFS cluster from a single point. It provides access to the tools that help you set up and administer your CXFS cluster:

- *Tasks* let you set up and monitor individual components of a CXFS cluster, including XVM volumes. For details about XVM tasks, see *XVM Volume Manager Administrator's Guide*.
- *Guided configuration tasks* 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 clicking their titles.

This section discusses the following:

- "GUI Window Layout" on page 192
- "File Menu" on page 193
- "Edit Menu" on page 193
- "Tasks Menu" on page 193
- "Help Menu" on page 194
- "Shortcuts Using Command Buttons" on page 194
- "View Menu" on page 196
- "Performing Tasks" on page 197
- "Using Drag-and-Drop" on page 198
- "Analyzing I/O Performance with Performance Co-Pilot on an IRIX Node" on page 198
- "Structuring Volume Topologies" on page 199
- "Configuring Disks" on page 200
- "Getting More Information" on page 200
- "Important GUI and xvm Command Differences" on page 200

GUI Window Layout

By default, the window is divided into two sections: the *view area* and the *details area* (see Figure 11-1 on page 190). The details area shows generic overview text if no item is selected in the view area. You can use the arrows in the middle of the window to shift the display.

File Menu

The **File** menu lets you display the following:

- Multiple windows for this instance of the GUI
- System log file:
 - IRIX: `/var/adm/SYSLOG`
 - SGI ProPack: `/var/log/messages`
- System administration log file:
 - IRIX: `/var/sysadm/salog`
 - SGI ProPack: `/var/lib/sysadm/salog`

The `salog` file shows the commands run directly by this instance of the GUI or some other instance of the GUI running commands on the system. (Changes should not be made simultaneously by multiple instances of the GUI or the GUI and `cxfs_admin` or `cmgr`.)

The **File** menu also lets you close the current window and exit the GUI completely.

Edit Menu

The **Edit** menu lets you expand and collapse the contents of the view area. You can choose to automatically expand the display to reflect new nodes added to the pool or cluster. You can also use this menu to select all items in the view menu or clear the current selections.

Tasks Menu

The **Tasks** menu contains the following:

- **Guided Configuration**, which contains the tasks to set up your cluster, define filesystems, create volumes, check status, and modify an existing cluster
- **Nodes**, which contains tasks to define and manage the nodes
- **Cluster**, which contains tasks to define and manage the cluster
- **Cluster Services**, which allows you to start and stop CXFS services, set the CXFS tiebreaker node, set the log configuration, and revoke or allow CXFS kernel membership of the local node

- **Switches and I/O Fencing**, which contains tasks to configure switch definitions and manage I/O fencing
- **Disks**, which contains XVM disk administration tasks
- **Volume Elements**, which contains tasks to create, delete, modify, and administer XVM volume elements
- **Filesystems**, which contains tasks to define and manage filesystems and relocate a metadata server
- **Privileges**, which lets you grant or revoke access to a specific task for one or more users
- **Find Tasks**, which lets you use keywords to search for a specific task

Help Menu

The **Help** menu provides an overview of the GUI and a key to the icons. You can also get help for certain items in blue text by clicking on them.

Shortcuts Using Command Buttons

The command buttons along the top of the GUI window provide a method of performing tasks quickly. When you click a button, the corresponding task executes using default values, usually without displaying a task window. To override the defaults, launch the task from the **Tasks** menu. Table 11-2 summarizes the shortcuts available; for details about these tasks, see the *XVM Volume Manager Administrator's Guide*.

Table 11-2 Command Buttons

Button	Task
	Labels selected unlabeled disks. If the selected disks include foreign and/or labeled disks, the Label Disks task will be run.
	Brings up the Slice Disk task with the selected disks as default inputs
	Creates a concat with a temporary name
	Creates a mirror with a temporary name
	Creates a stripe with a temporary name
	Creates a volume with a temporary name
	Creates a subvolume with a temporary name

Button	Task
	Starts the Performance Co-Pilot XVM I/O monitor <code>pmgsvm</code> on the IRIX server, displaying via X Windows to your local administration station
	Detaches the selected volume elements from their current parents
	Deletes the selected non-slice volume elements or unlabels the selected disks directly, or brings up the appropriate delete task for the selected component

View Menu

Choose what you want to view from the **View** menu:

- Nodes and cluster
- Filesystems
- Cluster volume elements
- Local volume elements
- Disks
- Switches
- Users
- Task privileges

Selecting Items to View or Modify

You can use the following methods to select items:

- Click to select one item at a time
- `Shift+click` to select a block of items

- Ctrl+click to toggle the selection of any one item

Another way to select one or more items is to type a name into the **Find** text field and then press **Enter** or click the **Find** button.

Viewing Component Details

To view the details on any component, click its name in the view area; see "Selecting Items to View or Modify" on page 196.

The configuration and status details for the component will appear in the details area to the right. At the bottom of the details area will be the **Applicable Tasks** list, which displays tasks you may wish to launch after evaluating the component's configuration details. To launch a task, click the task name; based on the component selected, default values will appear in the task window.

To see more information about an item in the details area, select its name (which will appear in blue); details will appear in a new window. Terms with glossary definitions also appear in blue.

Performing Tasks

To perform an individual task, do the following:

1. Select the task name from the **Task** menu or click the right mouse button within the view area. For example:

Task
 > **Guided Configuration**
 > **Set Up a New Cluster**

The task window appears.

As a shortcut, you can right-click an item in the view area to bring up a list of tasks applicable to that item; information will also be displayed in the details area.

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

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

3. Continue launching tasks as needed.

Using Drag-and-Drop

The GUI lets you use drag-and-drop to do the following:

- Move nodes between the pool and the cluster
- Structure volume topologies
- Administer XVM disks



Caution: Always exercise care when restructuring volume elements with drag-and-drop because data that resides on the volume element can be lost. The GUI attempts to warn the user when it can predict that there is a high likelihood of data loss. However, when a volume is not associated with a mounted filesystem, neither the `xvm` command nor the GUI can determine whether that volume holds important data.

To select multiple GUI icons, select the first icon by clicking the left mouse button, then press the `Ctrl` button while clicking on the additional icons. To select consecutive icons, select the first icon and press shift while selecting the last icon.

You cannot drag and drop between two GUI windows. You cannot drag and drop between the CXFS Manager and the IRIX Interactive Desktop Personal System Administration windows. You cannot drag and drop items onto shortcut command buttons.

See the *XVM Volume Manager Administrator's Guide* for more information about using drag-and-drop to structure volume topologies and configure disks.

Analyzing I/O Performance with Performance Co-Pilot on an IRIX Node

To analyze performance on an IRIX node, click the button to launch Performance Co-Pilot; see "Shortcuts Using Command Buttons" on page 194. The resulting Performance Co-Pilot window shows all volumes, with colored LEDs indicating read

and write I/O activity. Position the cursor over any LED and press the spacebar to view a window showing the value-color legend for the LED and the current value of the read or write rate for the corresponding XVM volume or volume element. Middle-mouse-click any LED to get a menu from which you can launch additional tools to show XVM read and write I/O activity charts and a 3D graphical view of disk activity.

Structuring Volume Topologies

To reconfigure a logical volume, do the following:

- Select the view you want:

View

> **Cluster Volume Elements**

or

View

> **Local Volume Elements**

- Select a volume element icon
- Drag the icon and drop it on another volume element icon

Icons turn blue as you drag to indicate when it is valid to drop upon them. When you drag, if the mouse cursor reaches the top or the bottom of the view area, the display will scroll automatically.

You can use drag-and-drop to operate on multiple volume elements of different types. For example, you can detach several types of volume elements by selecting items and dragging them to any **Unattached** heading, even if no selected item belongs to that category. You can select multiple items of different types and attach them to a parent. For example, you can select two concats and a stripe and use drag-and-drop to attach them to a parent concat.

You can rename volume elements by clicking a selected (highlighted) volume element and typing a new name into the text field.

Configuring Disks

To label or unlabel disks using drag-and-drop, select the following:

View
> Disks

Select an unlabeled disk then drag and drop it on the **Labeled Disks** heading, or select a labeled disk then drag and drop it on the **Unlabeled Disks** heading.

You can give away a disk using the task menu or drag-and-drop. In the **Disks** view, select a disk and then drag and drop it on the **Cluster Disks** heading.

Note: Giving away a disk presents less risk of data loss than stealing a disk.

You can label a disk by clicking a selected (highlighted) disk and typing a name into the resulting name text field.

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

Getting More Information

Click blue text to launch tasks or display one of the following:

- Term definitions
- Input instructions
- Item details
- The selected task window

Important GUI and `xvm` Command Differences

When volume elements other than volumes are created or detached, the system automatically creates a volume and a subvolume that are associated with the volume element. You can explicitly name this generated volume, in which case the volume name is stored in label space and persists across machine reboots.

The GUI does not display volumes and subvolumes that were not named explicitly. The GUI displays the children of these volumes and subvolumes as available for use or as unattached. In contrast, the `xvm` command shows all volumes and subvolumes.

The GUI displays filesystems that are on volumes that were not named explicitly, but lists the volumes as **None**. Volumes and subvolumes that the system generated automatically with temporary names are mentioned in the full paths of unattached volume elements (for example, `/vol196/datav`), but the GUI ignores them otherwise.

To reduce the risk of data loss, SGI recommends that you name volumes explicitly when using the GUI. If you have created volumes using the `xvm` command that you did not name explicitly, you can use the `xvm` tool to assign these volumes permanent names before proceeding. This can reduce the risk of data loss.

Key to Icons and States

The following tables show keys to the icons and states used in the CXFS Manager GUI.

Table 11-3 Key to Icons

Icon	Entity
	IRIX node (server-capable or client-only)
	SGI ProPack node (server-capable or client-only)
	AIX, Linux third-party, Mac OS X, Solaris, or Windows node (client-only)
	Cluster
	Expanded tree in view area

Icon	Entity
	Collapsed tree in view area
	Switch
	XVM disk
	Unlabeled disk
	Foreign disk
	Slice
	Volume
	Subvolume
	Concat
	

Icon	Entity
	Mirror
	
	Stripe
	
	Slot
	
	Local filesystem
	
	CXFS filesystem
	
	Copy on write
	
	Repository
	
	Snapshot
	
	User account
	

Icon	Entity
	GUI task for which execution privilege may be granted or revoked
	Privileged command executed by a given GUI task

Table 11-4 Key to States

Icon	State
	(grey icon) Inactive, unknown, offline — CXFS services may not be active
	(blue icon) Enabled for mount — CXFS services may not be active
	(blue icon) Online, ready for use, up, or mounted without error
	(green swatch) Open, in use

Icon	State
	(blinking orange arrow) Mirror reviving
	(red icon) Error detected, down or mounted with error

Guided Configuration Tasks

This section discusses the following guided configuration tasks:

- "Set Up an Existing FailSafe Cluster for CXFS with the GUI" on page 205
- "Make Changes to Existing Cluster" on page 206
- "Fix or Upgrade Cluster Nodes" on page 207

Also see "Set Up a New Cluster with the GUI" on page 169, "Set Up a New CXFS Filesystem with the GUI" on page 171, and "Check Cluster Status with the GUI" on page 402. For information about XVM guided configuration tasks, see the *XVM Volume Manager Administrator's Guide*.

Set Up an Existing FailSafe Cluster for CXFS with the GUI

Note: Within the CXFS tasks, you can click 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 **OK**.

The **Set Up an Existing FailSafe Cluster for use with CXFS** task leads you through the steps required to convert existing IRIS FailSafe nodes and cluster to CXFS. It encompasses tasks that are detailed elsewhere. This task appears on the CXFS GUI **only** if you also have FailSafe installed.

There is a single database for FailSafe and CXFS. If a given node applies to both products, ensure that any modifications you make are appropriate for both products.

Do the following:

1. Click **Convert a FailSafe Cluster for use with CXFS**. This will change the cluster type to CXFS and FailSafe. See "Convert a FailSafe Cluster for use with CXFS with the GUI" on page 226.
2. Stop high availability (HA) services on the nodes to be converted using the FailSafe GUI. See the *FailSafe Administrator's Guide for SGI InfiniteStorage*.
3. Add the second heartbeat and control NIC (for FailSafe use) to the node definitions using the CXFS GUI. See "Modify a Node Definition with the GUI" on page 217.
4. Click **Convert a FailSafe Node for use with CXFS** to convert the local node (the node to which you are connected). A converted node will be of type CXFS and FailSafe or CXFS. See "Convert a FailSafe Node for use with CXFS with the GUI" on page 221.
5. Click **Convert a FailSafe Node for use with CXFS** to convert another node. Repeat this step for each node you want to convert.
6. Click **Start CXFS Services**.

Make Changes to Existing Cluster

This task lists different ways to edit an existing cluster. You can make changes while the CXFS services are active, such as changing the way the cluster administrator is notified of events; however, you must first stop CXFS services before testing connectivity. You must unmount a file system before making changes to it.

See the following:

- "Modify a Cluster Definition with the GUI" on page 225
- "Set Up a New CXFS Filesystem with the GUI" on page 171
- "Modify a CXFS Filesystem with the GUI" on page 244
- "Define a Node with the GUI" on page 208
- "Test Node Connectivity with the GUI" on page 223
- "Add or Remove Nodes in the Cluster with the GUI" on page 216

Fix or Upgrade Cluster Nodes

This task leads you through the steps required to remove an administration node from a cluster. It covers the following steps:

- "Stop CXFS Services with the GUI" on page 228.
- Perform the necessary maintenance on the node. Only if required, see "Reset a Node with the GUI " on page 217.
- "Start CXFS Services with the GUI" on page 227.
- Monitor the state of the cluster components in the view area. See "Check Cluster Status with the GUI" on page 402.

When shutting down, resetting, or restarting a CXFS client-only node, do not stop CXFS services on the node. (Stopping CXFS services is more intrusive on other nodes in the cluster because it updates the cluster database. Stopping CXFS services is appropriate only for a CXFS administration node.) Rather, let the CXFS shutdown scripts on the node stop CXFS when the client-only node is shut down or restarted.

Node Tasks with the GUI

This section discusses the following:

- "Define a Node with the GUI" on page 208
- "Examples of Defining a Node with the GUI" on page 214
- "Add or Remove Nodes in the Cluster with the GUI" on page 216
- "Reset a Node with the GUI " on page 217
- "Modify a Node Definition with the GUI" on page 217
- "Convert a FailSafe Node for use with CXFS with the GUI" on page 221
- "Delete a Node with the GUI" on page 222
- "Test Node Connectivity with the GUI" on page 223
- "Display a Node with the GUI" on page 223

Note: The **Set Up a New Cluster** guided configuration task leads you through the tasks required to set up the cluster and nodes. See "Set Up a New Cluster with the GUI" on page 169.

Define a Node with the GUI

Note: Within the CXFS tasks, you can click 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 **OK**.

To define a node, do the following:

1. **Hostname:** Enter the hostname of the node you are defining. You can use a simple hostname, such as `lilly`, if it can be resolved by the name server or `/etc/hosts` on all nodes in the cluster; otherwise, use a fully qualified domain name such as `lilly.example.com`. Use the `ping` command to display the fully qualified hostname. Do not enter an IP address.

If you attempt to define a cluster or other object before the local node has been defined, you will get an error message that says:

```
No nodes are registered on servername. You cannot define a cluster until you define the node to which the GUI is connected. To do so, click "Continue" to launch the "Set Up a New Cluster" task.
```

2. **Logical Name:** Enter the simple hostname (such as `lilly`) or an entirely different name (such as `nodeA`). If you entered in the simple hostname for the **Hostname** field, the same name will be entered into the **Logical Name** field by default. Logical names cannot begin with an underscore (`_`) or include any whitespace characters, and can be at most 255 characters.

Note: To rename a node, you must delete it and then define a new node.

3. **Operating System:** Choose the name of the operating system that is running on the node being defined. Choose **Windows** for Windows 2000, Windows 2003, or Windows XP. Choose **Linux 64** when defining an `x86_64` or `ia64` architecture. (Use the `uname -i` command to determine the architecture type.)

An IRIX node or an SGI ProPack node can be a server-capable administration node or a CXFS client-only node, depending upon the node function selected and the software installed. AIX, Linux third-party, Mac OS X, Solaris, and Windows nodes are always CXFS client-only nodes. (An IRIX client administration node is only supported when used for FailSafe coexecution.)

If you select a fail action that includes reset, you will be given an opportunity to provide reset information on a second page. Any potential metadata server should include reset in its fail action hierarchy.

You cannot later modify the operating system for a defined node. To change the operating system, you would have to delete the node and then define a new node with the new name.

4. **Node Function:** Select one of the following:

- **Server-capable Admin** is an IRIX or SGI ProPack node on which you will execute cluster administration commands and that you also want to be a CXFS metadata server. (You will use the **Define a CXFS Filesystem** task to define the specific filesystem for which this node can be a metadata servers.) Use this node function only if the node will be a metadata servers. You must install the `cluster_admin` product on this node.
- **Client Admin** is an IRIX node on which you will execute cluster administration commands but that you do not want to use as a CXFS metadata server. Use this node function only if the node will run FailSafe but you do not want it to be a metadata server. You must install the `cluster_admin` product on this node.
- **Client-only** is a node that shares CXFS filesystems but on which you will not execute cluster administration commands and that will not be a CXFS metadata server. Use this node function for all nodes other than those that will be metadata servers, or those that will run FailSafe without being a metadata server. You must install the product on this node. This node can run AIX, IRIX, HP-UX, Linux third-party, SGI ProPack, Mac OS X, Solaris, or Windows. (Nodes other than IRIX and SGI ProPack are **required** to be client-only nodes.)

5. **Networks for Incoming Cluster Messages:** Do the following:

- **Network:** Enter the IP address or hostname of the NIC. (The hostname must be resolved in the `/etc/hosts` file.) The priorities of the NICs must be the same for each node in the cluster. For information about why a private network is required, see "Private Network" on page 19.

FailSafe requires at least two NICs.

- **Messages to Accept:** Select **Heartbeat and Control**.

You can use the **None** setting if you want to temporarily define a NIC but do not want it to accept messages. For more information, see "Cluster Environment" on page 9.

- Click **Add** to add the NIC to the list.

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

To delete a NIC from the list, click the NIC in the list to select it, then click **Delete**.

By default, the priority 1 NICs are used as the private network; they must be on the same subnet. To allow one network to fail over to another, you must group the NICs into failover networks manually by using `cxfs_admin`. See Chapter 12, "Reference to `cxfs_admin` Tasks" on page 255.

6. **Node ID:** (*Optional for administration nodes*) An integer in the range 1 through 32767 that is unique among the nodes in the pool. If you do not specify a number for an administration node, CXFS will calculate an ID for you.

For administration nodes, the default ID is a 5-digit number based on the machine's serial number and other machine-specific information; it is not sequential. For client-only nodes, you must supply the node ID.

You must not change the node ID number after the node has been defined. (There is no default CXFS tiebreaker; for more information, see "CXFS Kernel Membership, Quorum, and Tiebreaker" on page 512.)

7. **Partition ID:** (*Optional*) Uniquely defines a partition in a partitioned Origin 3000 system, Altix 3000 series system, or Altix 4700 system. If your system is not partitioned, leave this field empty. Use the IRIX `mkpart` command or the SGI ProPack `proc` command to determine the partition ID value (see below).

Click **Next** to move to the next screen.

8. **Fail Action:** Specify the set of actions that determines what happens to a failed node: the second action will be followed only if the first action fails; the third action will be followed only if the first and second fail.

The available actions depend upon the operating system value selected for the node:

- **Fence:** disables access to the SAN from the problem node. Fencing provides faster recovery of the CXFS kernel membership than reset.
- **FenceReset:** performs a fence and then, if the node is successfully fenced, **also** performs an asynchronous reset of the node via a system controller (according to the chosen reset method); recovery begins without waiting for reset acknowledgement.

Note: A server-capable node should also include **Reset** in its fail action hierarchy (unless it is the only server-capable node in the cluster).

- **Reset:** performs a system reset via a system controller. A server-capable node should include **Reset** in its fail action hierarchy.
- **Shutdown:** tells the other nodes in the cluster to wait for a period of time (long enough for the node to shut itself down) before reforming the CXFS kernel membership. (However, there is no notification that the node's shutdown has actually taken place.) The default fail action hierarchy for IRIX or SGI ProPack nodes is **Reset, Shutdown**. The default for other nodes is **Shutdown**.



Caution: There are issues when using **Shutdown** with server-capable nodes; for more information and for a list of valid failure policy sets, see "Isolating Failed Nodes: Failure Policies" on page 28. If you are using dynamic heartbeat monitoring, you must not use the **Shutdown** setting on a client-only node. For information about heartbeat monitoring, see "Heartbeat Monitoring" on page 27. To specify a fail policy without **Shutdown** you must define or modify the node with `cxfs_admin`. See Chapter 12, "Reference to `cxfs_admin` Tasks" on page 255.

Note: If the failure hierarchy contains **Reset** or **FenceReset**, the reset might be performed before the system kernel core-dump can complete, resulting in an incomplete core-dump.

9. If you have chosen a failure hierarchy that includes **Reset** or **FenceReset**, provide the following information.

- This node:
 - **Port Type:** select one of the following:
 - **L1** (Origin/Onyx 300/350, Origin/Onyx 3200C)
 - **L2** (Any Altix with an L2, Prism, Origin/Onyx 3000 series, Origin 300/350 over a direct-connect serial line)
 - **MSC** (Origin 200, Onyx2 Deskside, SGI 2100/2200 deskside systems)
 - **MMSC** (Rackmount SGI 2400/2800, Onyx2).

Note: Altix XE systems use baseboard management controller (BMC) for reset. To configure reset via BMC or L2 over the network, use the `cxfs_admin` configuration tool. See "Create or Modify a Node with `cxfs_admin`" on page 270.

- **Reset Method:** The type of reset to be performed:
 - **Power Cycle** shuts off power to the node and then restarts it
 - **Reset** simulates the pressing of the reset button on the front of the machine
 - **NMI** (nonmaskable interrupt) performs a core-dump of the operating system kernel, which may be useful when debugging a faulty machine

Note: NMI is not available on systems containing a BMC.

- **Port Password:** The password for the system controller port, **not** the node's `root` password or PROM password. On some nodes, 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.

- **Temporarily Disable Port:** If you want to provide reset information now but do not want to allow the reset capability at this time, check this box. If this box is checked, CXFS cannot reset the node.
- Owner (node that sends the reset command):
 - **Logical Name:** Name of the node that sends the reset command. If you use serial cables, they 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:** Name of the terminal port (TTY) on the owner node to which the system controller is connected (the node being reset). `/dev/ttyd2` is the most commonly used port, except on Origin 300 and Origin 350 systems (where `/dev/ttyd4` is commonly used) and Altix 350 systems (where `/dev/ttyIOC0` is commonly used). The other end of the cable connects to this node's (the node being reset) system controller port, so the node can be controlled remotely by the owner node. Check the owner node's specific hardware configuration to verify which tty device to use.

Note: To specify reset method that uses Intelligent Platform Management Interface (IPMI) or L2 over the network, use the `cxfs_admin` configuration tool. See "Create or Modify a Node with `cxfs_admin`" on page 270.

10. Click **OK**.

Note: Do not add a second node until the first node icon appears in the view area. The entire cluster status information is sent to each CXFS administration node each time a change is made to the cluster database; therefore, the more CXFS administration nodes in a configuration, the longer it will take.

You can use the IRIX `mkpart` command to determine the partition ID:

- The `-n` option lists the partition ID (which is 0 if the system is not partitioned).
- The `-l` option lists the bricks in the various partitions (use `rack#.slot#` format in the GUI).

On SGI ProPack, you can find the partition ID by reading the `proc` file. For example:

```
[root@linux64 root]# cat /proc/sgi_sn/partition_id
0
```

The 0 indicates that the system is not partitioned. If the system is partitioned, the number of partitions (such as 1, 2, etc.) is displayed.

For example (output truncated here for readability):

```
# mkpart -n
Partition id = 1
# mkpart -l
partition: 3 = brick: 003c10 003c13 003c16 003c21 ...
partition: 1 = brick: 001c10 001c13 001c16 001c21 ...
```

You could enter one of the following for the **Partition ID** field:

```
1
001.10
```

Examples of Defining a Node with the GUI

The following figures show an example of defining a new node.

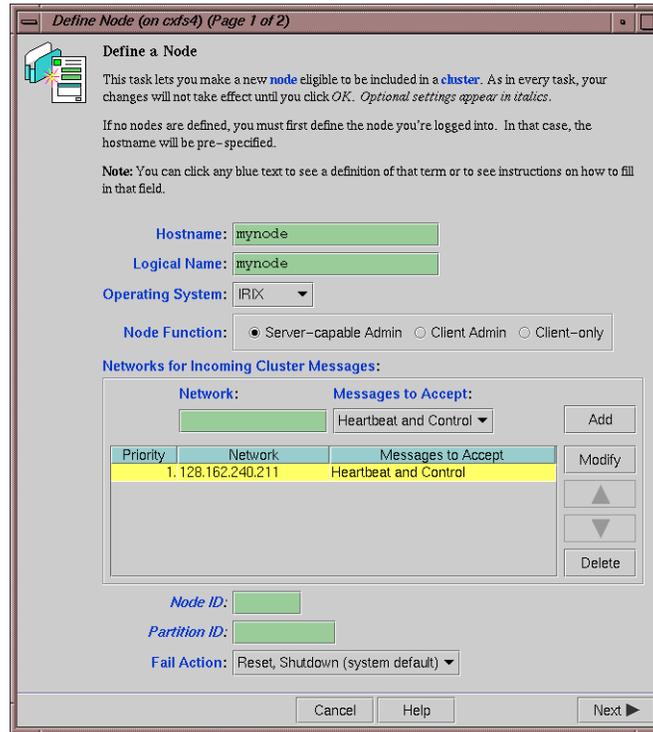


Figure 11-3 Example Node Definition

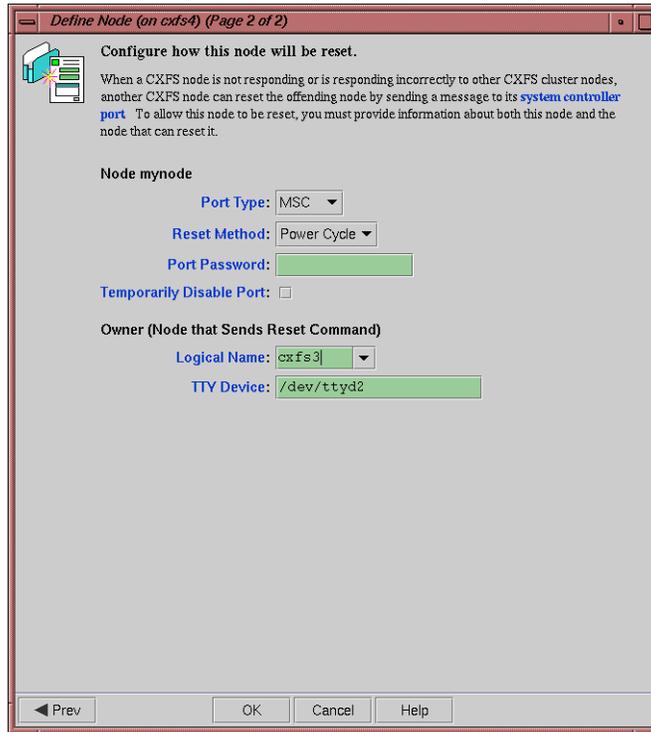


Figure 11-4 Example System Reset Settings

Add or 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 view area; set the **View** selection to **Nodes and Cluster**.

Do the following:

1. Add or remove the desired nodes:
 - To add a node, select its logical name from the **Available Nodes** pull-down menu and click **Add**. The node name will appear in the **Nodes to Go into Cluster** list. To select all of the available nodes, click **Add All**.
 - To delete a node, click its logical name in the **Nodes to Go into Cluster** screen. (The logical name will be highlighted.) Then click **Remove**.
2. Click **OK**.

Reset a Node with the GUI

You can use the GUI to reset IRIX or SGI ProPack 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 CXFS kernel membership.

To reset a node, do the following:

1. **Node to Reset:** Choose the node to be reset from the pull-down list.
2. Click **OK**.

Modify a Node Definition with the GUI

To rename a node or change its operating system, you must delete it and then define a new node.

To modify other information about a node, do the following:

1. **Logical Name:** 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.

2. **Networks for Incoming Cluster Messages:** The priorities of the NICs must be the same for each node in the cluster.
 - **Network:** To add a NIC for incoming cluster messages, enter the IP address or hostname into the **Network** text field and click **Add**.
 - To modify a NIC that is already in the list, click the network in the list in order to select it. Then click **Modify**. This moves the NIC out of the list and into the text entry area. You can then change it. To add it back into the list, click **Add**.
 - To delete a NIC, click the NIC in the priority list in order to select it. Then click **Delete**.
 - To change the priority of a NIC, click the NIC in the priority list in order to select it. Then click the up and down arrows in order to move it to a different position in the list.

You can use the **None** setting if you want to temporarily define a NIC but do not want it to accept messages. For more information, see "Cluster Environment" on page 9.

By default, the priority 1 NICs are used as the private network; they must be on the same subnet. To allow the one network to fail over to another, you must group the NICs into networks manually by using `cxfs_admin`. See Chapter 12, "Reference to `cxfs_admin` Tasks" on page 255.

Click **Next** to move to the next page.

3. **Partition ID:** (*Optional*) Uniquely defines a partition in a partitioned Origin 3000 system, Altix 3000 series system, or Altix 4700 system. If your system is not partitioned, leave this field empty. You can use the IRIX `mkpart` command or the SGI ProPack `proc` command to determine the partition ID value; see below.
4. **Fail Action:** Specify the set of actions that determines what happens to a failed node: the second action will be followed only if the first action fails; the third action will be followed only if the first and second fail.

The available actions depend upon the operating system value selected for the node:

- **Fence:** disables access to the SAN from the problem node. Fencing provides faster recovery of the CXFS kernel membership than reset.
- **FenceReset:** performs a fence and then, if the node is successfully fenced, **also** performs an asynchronous reset of the node via a system controller (according

to the chosen reset method); recovery begins without waiting for reset acknowledgement.

Note: A server-capable node should also include **Reset** in its fail action hierarchy (unless it is the only server-capable node in the cluster).

- **Reset:** performs a system reset via a system controller. A server-capable node should include **Reset** in its fail action hierarchy.
- **Shutdown:** tells the other nodes in the cluster to wait for a period of time (long enough for the node to shut itself down) before reforming the CXFS kernel membership. (However, there is no notification that the node's shutdown has actually taken place.) The default fail action hierarchy for IRIX or SGI ProPack nodes is **Reset, Shutdown**. The default for other nodes is **Shutdown**.



Caution: There are issues when using **Shutdown** with server-capable nodes; for more information and for a list of valid failure policy sets, see "Isolating Failed Nodes: Failure Policies" on page 28. If you are using dynamic heartbeat monitoring, you must not use the **Shutdown** setting on a client-only node. For information about heartbeat monitoring, see "Heartbeat Monitoring" on page 27. To specify a fail policy without **Shutdown** you must define or modify the node with `cxfs_admin`. See Chapter 12, "Reference to `cxfs_admin` Tasks" on page 255 .

5. If you have chosen a failure hierarchy that includes **Reset** or **FenceReset**, provide the following information.
 - This node:
 - **Port Type:** select one of the following:
 - **L1** (Origin/Onyx 300/350, Origin/Onyx 3200C)
 - **L2** (Any Altix with an L2, Prism, Origin/Onyx 3000 series, Origin 300/350 over a direct-connect serial line)
 - **MSC** (Origin 200, Onyx2 Deskside, SGI 2100/2200 deskside systems)
 - **MMSC** (Rackmount SGI 2400/2800, Onyx2).

Note: Altix XE systems use baseboard management controller (BMC) for reset. To configure reset via BMC or L2 over the network, use the `cxfs_admin` configuration tool. See "Create or Modify a Node with `cxfs_admin`" on page 270.

- **Reset Method:** The type of reset to be performed:
 - **Power Cycle** shuts off power to the node and then restarts it
 - **Reset** simulates the pressing of the reset button on the front of the machine
 - **NMI** (nonmaskable interrupt) performs a core-dump of the operating system kernel, which may be useful when debugging a faulty machine

Note: NMI is not available on systems containing a BMC.

- **Port Password:** The password for the system controller port, **not** the node's root password or PROM password. On some nodes, 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.
- **Temporarily Disable Port:** If you want to provide reset information now but do not want to allow the reset capability at this time, check this box. If this box is checked, CXFS cannot reset the node.
- Owner (node that sends the reset command):
 - **Logical Name:** Name of the node that sends the 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:** Name of the terminal port (TTY) on the owner node to which the system controller is connected. `/dev/ttyd2` is the most commonly used port, except on Origin 300 and Origin 350 systems (where

/dev/ttyd4 is commonly used) and Altix 350 systems (where /dev/ttyIOC0 is commonly used). 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.

Note: To specify reset method that uses Intelligent Platform Management Interface (IPMI) or L2 over the network, use the `cxfs_admin` configuration tool. See "Create or Modify a Node with `cxfs_admin`" on page 270.

6. Click **OK**.

You can use the IRIX `mkpart` command to determine the partition ID value:

- The `-n` option lists the partition ID (which is 0 if the system is not partitioned).
- The `-l` option lists the bricks in the various partitions (use `rack#.slot#` format in the GUI).

For example (output truncated here for readability):

```
# mkpart -n
Partition id = 1
# mkpart -l
partition: 3 = brick: 003c10 003c13 003c16 003c21 ...
partition: 1 = brick: 001c10 001c13 001c16 001c21 ...
```

You could enter one of the following for the **Partition ID** field:

```
1
001.10
```

On SGI ProPack, you can find the partition ID by reading the `proc` file. For example:

```
[root@linux64 root]# cat /proc/sgi_sn/partition_id
0
```

The 0 indicates that the system is not partitioned. If the system is partitioned, the number of partitions (such as 1, 2, etc.) is displayed.

Convert a FailSafe Node for use with CXFS with the GUI

This task appears on the CXFS GUI only if you also have FailSafe installed. It applies only to CXFS administration nodes.

You can convert an existing FailSafe node (of type `FailSafe`) to either of the following types:

- `CXFS` and `FailSafe`
- `CXFS`

Do the following:

1. Stop HA services on the node to be converted using the FailSafe GUI. See the *FailSafe Administrator's Guide for SGI InfiniteStorage*.
2. Add the second **Heartbeat and Control** NIC (for FailSafe use) to the node definition using the CXFS GUI. See "Modify a Node Definition with the GUI" on page 217.
3. Enter the following information:
 - **Logical Name:** Choose the logical name of the node from the pull-down list.
 - **Keep FailSafe Settings:**
 - To convert to type `CXFS` and `FailSafe`, click the checkbox
 - To convert to type `CXFS`, leave the checkbox blank
 - Click **OK**.

Note: If you want to rename a node, you must delete it and then define a new node.

To change other parameters, see "Modify a Node Definition with the GUI" on page 217. Ensure that modifications you make are appropriate for both FailSafe and CXFS.

To convert a CXFS node so that it applies to FailSafe, use the FailSafe GUI. For information about the FailSafe GUI, see the *FailSafe Administrator's Guide for SGI InfiniteStorage*.

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 a Cluster Definition with the GUI" on page 225.

To delete a node, do the following:

1. **Node to Delete:** Select the logical name of the node to be deleted from the pull-down list.
2. Click **OK**.

Test Node Connectivity with the GUI

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

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

1. Choose whether to test by network or serial connectivity by clicking the appropriate radio button.
2. Choose a node to be tested from the pull-down list and add it to the test list by clicking **Add**.

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

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

Display a Node with the GUI

After you define nodes, you can use the **View** selection in the view area to display the following:

- **Nodes and Cluster** shows the nodes that are defined as part of a cluster or as part of the pool (but not in the cluster)

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

Cluster Tasks with the GUI

This section discusses the following:

- "Define a Cluster with the GUI" on page 224
 - "Modify a Cluster Definition with the GUI" on page 225
 - "Convert a FailSafe Cluster for use with CXFS with the GUI" on page 226
 - "Delete a Cluster with the GUI" on page 227
 - "Display a Cluster with the GUI" on page 227
-

Note: The **Set Up a New Cluster** guided configuration task leads you through the tasks required to set up the cluster and nodes. See "Set Up a New Cluster with the GUI" on page 169.

Define a Cluster with the GUI

A *cluster* is a collection of nodes coupled to each other by a private 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. Enter the following information:
 - **Cluster Name:** The logical name of the cluster. Logical names cannot begin with an underscore (`_`) or include any whitespace characters, and can be at most 255 characters. Clusters must have unique names.
 - **Cluster ID:** A unique number within your network in the range 1 through 255. The cluster ID is used by the operating system 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 must not change a cluster ID after the cluster has been defined. Clusters must have unique IDs.
 - **Cluster Mode:** Usually, you should set the cluster to the default `Normal` mode.

Setting the mode to `Experimental` turns off heartbeating in the CXFS kernel membership code so that you can debug the cluster without causing node failures. For example, this can be useful if you just want to disconnect the network for a short time (provided that there is no other cluster networking activity, which will also detect a failure even if there is no heartbeating) or if

you want to enter the kernel debugger (which stops heartbeat) on a CXFS node. You should only use `Experimental` mode when debugging.

- **Notify Administrator** (of cluster and node status changes):
 - **By e-mail:** This choice 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:** This choice requires that you specify the command to be run whenever the status changes for a node or cluster.
 - **Never:** This choice specifies that notification is not sent.

2. Click **OK**.

Modify a Cluster Definition with the GUI

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

1. Enter the following information:
 - **Cluster Name:** Choose from the pull-down list.
 - **Cluster Mode:** Usually, you should set the cluster to the default `Normal` mode. See "Define a Cluster with the GUI" on page 224, for information about `Experimental` mode.
 - **Notify Administrator:** Select the desired notification. For more information, see "Define a Cluster with the GUI" on page 224.
2. Click **OK**.

To modify the nodes that make up a cluster, see "Add or Remove Nodes in the Cluster with the GUI" on page 216.

Note: If you want to rename a cluster, you must delete it and then define a new cluster. If you have started CXFS services on the node, you must either reboot it or reuse the cluster ID number when renaming the cluster.

However, be aware that if you already have CXFS filesystems defined and then rename the cluster, CXFS will not be able to mount the filesystems. For more information, see "Cannot Mount Filesystems" on page 449.

Convert a FailSafe Cluster for use with CXFS with the GUI

This task appears on the CXFS GUI only if you also have FailSafe installed.

To convert the information from an existing IRIS FailSafe cluster (that is, of type `FailSafe`) to create a cluster that applies to CXFS (that is, of type `CXFS` and `FailSafe` or of type `CXFS`), do the following:

1. Enter the following information:
 - **Cluster Name:** Choose from the pull-down list.
 - **Cluster ID:** Enter a unique number within your network in the range 1 through 255. The cluster ID is used by the operating system 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 must not change a cluster ID after the cluster has been defined.
2. Click **OK**.

The cluster will apply to both IRIS FailSafe and CXFS. To modify the nodes that make up a cluster, see "Add or Remove Nodes in the Cluster with the GUI" on page 216.

Note: If you want to rename a cluster, you must delete it and then define a new cluster.

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 "Add or Remove Nodes in the Cluster with the GUI" on page 216.

To delete a cluster, do the following:

1. **Cluster to Delete:** The name of the cluster is selected for you.
2. Click **OK**.

Display a Cluster with the GUI

From the **View** selection, you can choose elements to examine. To view details of the cluster, click the cluster name or icon; status and configuration information will appear in the details area on the right.

Cluster Services Tasks with the GUI

This section discusses the following:

- "Start CXFS Services with the GUI" on page 227
- "Stop CXFS Services with the GUI" on page 228
- "Set Tiebreaker Node with the GUI" on page 229
- "Set Log Configuration with the GUI" on page 230
- "Revoke Membership of the Local Node with the GUI" on page 232
- "Allow Membership of the Local Node with the GUI" on page 232

Start CXFS Services with the GUI

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

1. **Node(s) to Activate:** Select **All Nodes** or the individual node on which you want to start CXFS services.

2. Click **OK**.

Stop CXFS Services with the GUI

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

To stop CXFS 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:

- IRIX:

```
# /etc/init.d/cluster stop (on an admin node)
# /etc/init.d/cxfs stop (on an admin node)

# /etc/init.d/cxfs_client stop (on a client-only node)
```

- SGI ProPack:

```
# /etc/init.d/cxfs_cluster stop (on an admin node)
# /etc/init.d/cxfs stop (on an admin node)

# /etc/init.d/cxfs_client stop (on a client-only node)
```

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

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

1. Enter the following information:
 - **Force:** If you want to forcibly stop CXFS services even if there are errors (which would normally prevent the stop operation), click the **Force** checkbox.
 - **Node(s) to Deactivate:** Select `All Nodes` or the individual node on which you want to stop CXFS services.

If you stop CXFS services on one node, that node will no longer have access to any filesystems. If that node was acting as the metadata server for a

filesystem, another node in the list of potential metadata servers will be chosen. Clients of the filesystem will experience a delay during this process.

2. Click **OK**. It may take a few minutes to complete the process.

After you have stopped CXFS services on a node, the node is no longer an active member of the cluster. CXFS services will not be restarted when the system reboots.



Caution: You should stop CXFS services before using the shutdown or reboot commands. If you execute shutdown or reboot when CXFS 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 Tiebreaker Node with the GUI

A CXFS *tiebreaker node* determines whether a CXFS kernel membership quorum is maintained when exactly half of the server-capable nodes are up and can communicate with each other. There is no default CXFS tiebreaker. For more information, see "CXFS Kernel Membership, Quorum, and Tiebreaker" on page 512.



Caution: If one of the server-capable nodes is the CXFS tiebreaker in a two server-capable cluster, failure of that node or stopping the CXFS services on that node will result in a cluster-wide forced shutdown. Therefore SGI recommends that you use client-only nodes as tiebreakers so that either server could fail but the cluster would remain operational via the other server.

To ensure data integrity, SGI recommends that you use system reset for all potential metadata servers and reset or or I/O fencing for all client-only nodes; reset is required for IRIS FailSafe.

The current CXFS tiebreaker node is shown in the detailed view of the cluster.

To set the CXFS tiebreaker node, do the following:

1. **Tie-Breaker Node:** Select the desired node from the list. If there currently is a CXFS tiebreaker, it is selected by default.

To unset the CXFS tiebreaker node, select None.

2. Click **OK**.

Set Log Configuration with the GUI

CXFS maintains logs for each of the CXFS daemons. CXFS logs both normal operations and critical errors to individual log files for each log group and the system log file:

- IRIX: `/var/adm/SYSLOG`
- SGI ProPack: `/var/log/messages`

You can customize the logs according to the level of logging you wish to maintain.



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

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

- **Log Group:** A set of one or more CXFS processes that use the same log configuration. A log group usually corresponds to one CXFS daemon, such as `crsd`.
- **Log Level:** A number controlling the amount of log messages that CXFS will write into an associated log group's log file.
- **Log File:** The file in which to log messages.

See also "Status in Log Files" on page 400.

Display Log Group Definitions with the GUI

To display log group definitions, do the following:

1. **Log Group:** 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.

2. Click **OK**.

Configure Log Groups with the GUI

To configure a log group, do the following in the **Set Log Configuration** task:

1. Enter the appropriate information:
 - **Log Group:** 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.
 - **Log Level:** Select the log level, which specifies the amount of logging.



Caution: The **Default** log level is quite verbose; using it could cause space issues on your disk. You may wish to select a lower log level. Also see "Log File Management" on page 339, "`cad.options` on CXFS Administration Nodes" on page 129, and "`fs2d.options` on CXFS Administration Nodes" on page 130.

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 IRIX `/var/adm/SYSLOG` and SGI ProPack `/var/log/messages` 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
2. **Log File:** Do not change this value.
 3. Click **OK**.

Revoke Membership of the Local Node with the GUI

You should revoke CXFS kernel membership of the local node only in the case of error, such as when you need to perform a forced CXFS shutdown (see "Shutdown of the Database and CXFS" on page 332).

To revoke CXFS kernel membership for the local node, do the following:

1. **Local Node:** Verify the name of the local node, which will be displayed in the pop-up window.
2. Click **OK** to complete the task.

This result of this task will be considered as a node failure by the rest of the cluster. The rest of the cluster may then fail due to a loss of CXFS kernel membership quorum, or it may decide to reset the failed node. To avoid the reset, you can modify the node definition to disable the system controller status.

Allow Membership of the Local Node with the GUI

You must allow CXFS kernel membership for the local node (the node to which the GUI is connected) after fixing the problems that required a forced CXFS shutdown; doing so allows the node to reapply for CXFS kernel membership in the cluster. A forced CXFS shutdown can be performed manually or can be triggered by the kernel. For more information, see "Shutdown of the Database and CXFS" on page 332.

You must actively allow CXFS kernel membership of the local node in the following situations:

- After a manual revocation as in "Revoke Membership of the Local Node with the GUI" on page 232.
- When instructed to by an error message on the console or in system log file:
 - IRIX: `/var/adm/SYSLOG`
 - SGI ProPack: `/var/log/messages`
- After a kernel-triggered revocation. This situation is indicated by the following message in system log file (IRIX `/var/adm/SYSLOG` or SGI ProPack `/var/log/messages`):

```
Membership lost - withdrawing from cluster
```

To allow CXFS kernel membership for the local node, do the following:

1. **Local Node:** Verify the name of the local node, which will be displayed in the pop-up window.
2. Click **OK** to complete the task.

Switches and I/O Fencing Tasks with the GUI

This section discusses the following:

- "Define a Switch with the GUI" on page 233
- "Modify a Switch Definition with the GUI" on page 235
- "Update Switch Port Information with the GUI" on page 236
- "Delete a Switch Definition with the GUI" on page 236
- "Raise the I/O Fence for a Node with the GUI" on page 237
- "Lower the I/O Fence for a Node with the GUI" on page 237

See the release notes for supported switches.

Note: Nodes without system controllers require I/O fencing to protect data integrity.

Define a Switch with the GUI

This task lets you define a new Brocade switch to support I/O fencing in a cluster.

Do the following:

1. Enter the following information:
 - **Switch Name:** Enter the hostname of the switch; this is used to determine the IP address of the switch.
 - **Username:** Enter the user name to use when sending a telnet message to the switch. By default, this value is admin.
 - **Password:** Enter the password for the specified **Username** field.
 - **Mask:** Enter one of the following:

- A list of ports in the switch that will never be fenced. The list has the following form, beginning with the # symbol and separating each port number with a comma:

`#port , port , port . . .`

Each *port* is a decimal integer in the range 0 through 1023. Use a hyphen to specify an inclusive range. For example, the following indicates that port numbers 2, 4, 5, 6, 7, and 23 will never be fenced:

`#2 , 4-7 , 23`

- A hexadecimal string that represents the list of ports in the switch that will never be fenced.

Ports are numbered from zero. If a given bit has a binary value of 0, the port that corresponds to that bit is eligible for fencing operations; if 1, then the port that corresponds to that bit will always be excluded from any fencing operations. For example, Figure 11-5 shows that a mask of FF03 for a 16-port switch indicates that only ports 2–7 are eligible for fencing (because they have binary values of 0). Similarly, it shows that a mask of A4 for an 8-port switch allows fencing only on ports 0, 1, 3, 4, and 6 (the port numbers corresponding to binary 0) — ports 2, 5, and 7 will never be fenced (the port numbers corresponding to the nonzero value).

16-port Switch (1= never fence, 0= may fence)

Port #	15 14 13 12	11 10 9 8	7 6 5 4	3 2 1 0
Binary	1 1 1 1	1 1 1 1	0 0 0 0	0 0 1 1
Hexadecimal	F	F	0	3

8-port Switch

Port #	7 6 5 4	3 2 1 0
Binary	1 0 1 0	0 1 0 0
Hexadecimal	A	4

Figure 11-5 Bit Mask Representation for I/O Fencing

CXFS administration nodes automatically discover the available HBAs and, when fencing is triggered, will fence off all of the Fibre Channel HBAs when the **Fence** or **FenceReset** fail action is selected. However, masked HBAs will not be fenced. Masking allows you to prevent the fencing of devices that are attached to the SAN but are not shared with the cluster, to ensure that they remain available regardless of CXFS status. You would want to mask HBAs used for access to tape storage, or HBAs that are only ever used to access local (nonclustered) devices.

- **Vendor:** Select the name of the switch vendor or enter the vendor name manually if not found in the list.
2. Click **OK** to complete the task.

Modify a Switch Definition with the GUI

This task lets you modify an existing Brocade switch definition.

Note: To modify the definition of another type of switch, such as QLogic, you must use the `hafence(1M)` command. See "Using `hafence` to Manipulate a Switch" on page 318.

Do the following:

1. Enter the following information:
 - **Switch Name:** Select the hostname of the switch to be modified.
 - **Username:** Enter the user name to use when sending a `telnet` message to the switch. By default, this value is `admin`.
 - **Password:** Enter the password for the specified **Username** field.
 - **Mask:** Enter a list of port numbers or a hexadecimal string that represents the list of ports in the switch that will not be fenced. For more information, see "Define a Switch with the GUI" on page 233.
2. Click **OK** to complete the task.

Note: You cannot modify the vendor name for a switch. To use a different vendor, delete the switch and redefine it.

Update Switch Port Information with the GUI

This task lets you update the mappings between the host bus adapters (HBAs) and switch ports. You should run this command if you reconfigure any switch or add ports. Click **OK** to complete the task.

Delete a Switch Definition with the GUI

This task lets you delete an existing switch definition. Do the following:

1. **Switch Name:** Select the hostname of the Fibre Channel switch to be deleted.
2. Click **OK** to complete the task.

Raise the I/O Fence for a Node with the GUI

This task lets you raise the I/O fence for a node. Raising an I/O fence isolates the node from the SAN; CXFS sends a messages via the telnet protocol to the switch and disables the port. After the node is isolated, it cannot corrupt data in the shared CXFS filesystem.

Do the following:

1. **Raise Fence for Node:** Select the name of the node you want to isolate. Only nodes that have been configured with a **Fence** or **FenceReset** fail action can be selected.
2. Click **OK** to complete the task.

Lower the I/O Fence for a Node with the GUI

This task lets you lower the I/O fence for a given node by reenabling the port. Lowering an I/O fence allows the node to reconnect to the SAN and access the shared CXFS filesystem.

Do the following:

1. **Lower Fence for Node:** Select the node you want to reconnect. Only nodes that have been configured with a **Fence** or **FenceReset** fail action can be selected.
2. Click **OK** to complete the task.

Filesystem Tasks with the GUI

The following tasks let you configure CXFS 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 CXFS Filesystem** guided configuration task leads you through the steps required to set up a new CXFS filesystem. See "Set Up a New CXFS Filesystem with the GUI" on page 171.

This section discusses the following:

- "Make Filesystems with the GUI" on page 238
- "Grow a Filesystem with the GUI" on page 240
- "Define CXFS Filesystems with the GUI" on page 241
- "Modify a CXFS Filesystem with the GUI" on page 244
- "Mount CXFS Filesystems with the GUI" on page 245
- "Unmount CXFS Filesystems with the GUI" on page 246
- "Mount a Filesystem Locally" on page 246
- "Unmount a Local Filesystem" on page 247
- "Delete a CXFS Filesystem with the GUI" on page 247
- "Remove Filesystem Mount Information" on page 247
- "Relocate a Metadata Server for a CXFS Filesystem with the GUI" on page 248

Make Filesystems with the GUI

This task lets you create a filesystem on a volume that is online but not open. To create filesystems on multiple volume elements, use the **Browse** button.



Caution: Clicking OK will erase all data that exists on the target volume.

To make a filesystem, do the following:

1. Enter the following information:
 - **Domain:** Select the domain that will own the volume element to be created. Choose **Local** if the volume element or disk is defined for use only on the node to which the GUI is connected, or choose **Cluster** if it is defined for use on multiple nodes in the cluster.
 - **Volume Element:** Select the volumes on which to create the filesystem or select the volume elements whose parent volumes will be used for the filesystems. The menu lists only those volume elements that are available. (When volume elements other than volumes are created or detached, the system automatically creates a volume and a subvolume that are associated

with the volume element. If you did not explicitly name an automatically generated volume, the GUI will display its children only.)

- **Specify Sizes:** Check this box to modify the default options for the filesystem, including data region size, log size, and real-time section size.

By default, the filesystem will be created with the data region size equal to the size of the data subvolume. If the volume contains a log subvolume, the log size will be set to the size of the log subvolume. If the volume contains a real-time subvolume, the real-time section size will be set to the size of the real-time subvolume.

2. If you checked the **Specify Sizes** box, click **Next** to move to page 2. On page 2, enter the following information. For more information about these fields, see the IRIX `mkfs_xfs` or SGI ProPack `mkfs.xfs` man page.

- **Block Size:** Select the fundamental block size of the filesystem in bytes.
- **Directory Block Size:** Select the size of the naming (directory) area of the filesystem in bytes.
- **Inode Size:** Enter the number of blocks to be used for inode allocation, in bytes. The inode size cannot exceed one half of the **Block Size** value.
- **Maximum Inode Space:** Enter the maximum percentage of space in the filesystem that can be allocated to inodes. The default is 25%. (Setting the value to 0 means that the entire filesystem can become inode blocks.)
- **Flag Unwritten Extents:** Check this box to flag unwritten extents. If unwritten extents are flagged, filesystem write performance will be negatively affected for preallocated file extents because extra filesystem transactions are required to convert extent flags for the range of the file.

You should disable this feature (by unchecking the box) if the filesystem must be used on operating system versions that do not support the flagging capability.

- **Data Region Size:** Enter the size of the data region of the filesystem as a number of 512-byte blocks. This number is usually equal to the size of the data subvolume. You should specify a size other than 0 only if the filesystem should occupy less space than the size of the data subvolume.
- **Use Log Subvolume for Log:** Check this box to specify that the log section of the filesystem should be written to the log subvolume of the XVM logical

volume. If the volume does not contain a log subvolume, the log section will be a piece of the data section on the data subvolume.

- **Log Size:** Enter the size of the log section of the filesystem as a number of 512-byte blocks. You should specify a size other than 0 only if the log should occupy less space than the size of the log subvolume.
- **Real-Time Section Size:** Enter the size of the real-time section of the filesystem as a number of 512-byte blocks. This value is usually equal to the size of the real-time subvolume, if there is one. You should specify a size other than 0 only if the real-time section should occupy less space than the size of the real-time subvolume.

Note: XVM on SGI ProPack does not support real-time subvolumes.

3. Click **OK**.

Grow a Filesystem with the GUI

This task lets you grow a mounted filesystem.

Note: In order to grow a filesystem, you must first increase the size of the logical volume on which the filesystem is mounted. For information on modifying XVM volumes, see the *XVM Volume Manager Administrator's Guide*.

To grow a filesystem, do the following:

1. Enter the following information:
 - **Filesystem:** Select the name of the filesystem you want to grow. The list of available filesystems is determined by looking for block devices containing XFS superblocks.
 - **Specify Sizes:** Check this option to modify the default options for the filesystem, including data region size and (if already present for the filesystem) log size and real-time section size.

By default, the filesystem will be created with the data region size equal to the size of the data subvolume. If the volume contains a log subvolume, the log size will be set to the size of the log subvolume. If the volume contains a

real-time subvolume, the real-time section size will be set to the size of the real-time subvolume.

2. If you checked the **Specify Sizes** box, click **Next** to move to page 2. For more information about these fields, see the IRIX `mkfs_xfs` or SGI ProPack `mkfs.xfs` man page.
 - **Data Region Size:** Enter the size of the data region of the filesystem as a number of 512-byte blocks. This number is usually equal to the size of the data subvolume. You should specify a size other than 0 only if the filesystem should occupy less space than the size of the data subvolume.
 - **Log Size:** Enter the size of the log section of the filesystem as a number of 512-byte blocks. You should specify a size other than 0 only if the log should occupy less space than the size of the log subvolume. This option only appears if the filesystem has a log subvolume.
 - **Real-Time Section Size:** Enter the size of the real-time section of the filesystem as a number of 512-byte blocks. This value is usually equal to the size of the real-time subvolume, if there is one. You should specify a size other than 0 only if the real-time section should occupy less space than the size of the real-time subvolume. This option only appears if the filesystem has a real-time subvolume.

Note: XVM on SGI ProPack does not support real-time subvolumes.

3. Click **OK**.

Define CXFS Filesystems with the GUI

This task lets you define one or more CXFS filesystems having the same ordered list of potential metadata servers and the same list of client nodes.

Note: If you select multiple device names, the path you enter for the mount point will be used as a prefix to construct the actual mount point for each filesystem.

This task assumes that you have created volume headers on your disk drives, created the XVM logical volumes, and made the filesystems. "Initial Setup with the CXFS GUI" on page 167.

To define filesystems, do the following:

1. Enter the following information:

- **Device Name:** Select the device names of the XVM volumes on which the filesystems will reside.
- **Mount Point:** The directory on which the specified filesystem will be mounted. This directory name must begin with a slash (/). The same mount point will be used on all the nodes in the cluster. For example, if you select the device name `/dev/cxvm/cxfs1` and want to mount it at `/mount/cxfs1`, you would enter `/mount/cxfs1` for the **Mount Point** value.

If you selected multiple device names in order to define multiple CXFS filesystems, the mount point path will be constructed using the mount point you enter as a **prefix** and the name of each device name (not including the `/dev/cxvm` portion) as the suffix. For example, if you select two volume device names (`/dev/cxvm/cxfs1` and `/dev/cxvm/cxfs2`) and enter a mount point of `/mount/`, then the CXFS filesystems will be mounted as `/mount/cxfs1` and `/mount/cxfs2`, respectively. If instead you had entered `/mount` for the mount point, the filesystems would be mounted as `/mountcxfs1` and `/mountcxfs2`.

For more information, see the `mount` man page.

- *(Optional)* **Mount Options:** These options are passed to the `mount` command and are used to control access to the specified XVM volume. Separate multiple options with a comma. For a list of the available options, see the `fstab` man page.
- **Force Unmount:** Select the default behavior for the filesystem. This option controls what action CXFS takes if there are processes that have open files or current directories in the filesystems that is to be unmounted. If you select **On**, the processes will be killed and the unmount will occur. If you select **Off**, the processes will not be killed and the filesystem will not be unmounted. SGI recommends that you set **Force Unmount** to **On** in order to improve the stability of the CXFS cluster. This value can be overridden when you perform a manual unmount; see "Unmount CXFS Filesystems with the GUI" on page 246.
- **Metadata Servers:** A list of administration nodes that are able to act as metadata servers. All potential metadata servers within a cluster must run the same type of operating system (that is, all IRIX or all SGI ProPack).

To add a CXFS administration node to the list of servers, choose a name from the pull-down node list and click **Add**. To select all nodes listed, click **Add All**.

Note: Relocation is disabled by default. Recovery and relocation are supported only when using standby nodes. Therefore, you should only define multiple metadata servers for a given filesystem if you are using the standby node model. See "Relocation" on page 20.

To remove a node from the list of servers, click the name in the list to select it and then click **Remove**.

Note: The order of servers is significant. The first node listed is the preferred metadata server. Click a logical name to select it and then click 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 during the boot-up cycle because of network latencies and other unpredictable delays. The first available node in the list will be used as the active metadata server.

- **Enable Mount on:** A choice of either all nodes in the cluster or a list of selected CXFS administration nodes that you specify on a second page. (The filesystem is always mounted on the current metadata server.)
- **If Nodes are Added to the Cluster Later:** This option permits the filesystem to be mounted on all nodes that might be added to the cluster at some later date. This option is selected by default.
- If you chose **Only Selected Nodes** above, click **Next** to move to the second page of the task.

Selected Nodes: You can select the desired nodes from the **Node** list. You can also click **Add All** to select all nodes, which is the same as selecting **All Nodes Currently in Cluster**.

2. Click **OK**.

After defining the filesystems, you can mount them on the specified client nodes in the cluster by running the **Mount CXFS Filesystems** task.

Note: After a filesystem has been defined in CXFS, running `mkfs` on it (or using the "Make Filesystems with the GUI" on page 238 task) will cause errors to appear in the system log file. To avoid these errors, run `mkfs` before defining the filesystem in CXFS, or delete the CXFS filesystem before running `mkfs`. See "Delete a CXFS Filesystem with the GUI" on page 247.

Modify a CXFS Filesystem with the GUI

Note: You cannot modify a mounted filesystem.

To modify an existing filesystem, do the following:

1. Enter the following information:
 - **Filesystem to Modify:** Choose a filesystem from the pull-down menu. This displays information for that filesystem in the various fields.
 - **Mount Point** and **Mount Options:** 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.
 - *(Optional)* **Mount Options:** These options are passed to the `mount` command and are used to control access to the specified XVM volume. For a list of the available options, see the `fstab` man page.
 - **Metadata Servers:**
 - To delete a node from the list of servers, click its name and then click **Delete**.
 - To add a new CXFS administration node to the list of servers, select it from the pull-down list and click **Add**. To select all CXFS administration nodes, select **Add All**. The list for a given filesystem must consist of nodes running the same operating system.
 - To rearrange the priority of a server, select it by clicking its name and then click the arrow buttons as needed.

- **Enable Mount on:** A choice of either all nodes in the cluster or a list of selected nodes that you specify on a second page. (The filesystem is always mounted on the current metadata server.)
- **If Nodes are Added to the Cluster Later:** This option permits the filesystem to be mounted on all nodes that might be added to the cluster at some later date. This option is selected by default.
- If you chose **Only Selected Nodes** above, click **Next** to move to the second page of the task.

Selected Nodes: You can select the desired nodes from the **Node** list. You can also click **Add All** to select all nodes, which is the same as selecting **All Nodes Currently in Cluster**.

2. Click **OK**.

Mount CXFS Filesystems with the GUI

To mount existing filesystems on all of their client nodes, do the following:

1. **Filesystem to Mount:** Choose the filesystem to be mounted.
2. Click **OK**.

If CXFS 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 a Filesystem** task. The filesystem will not actually be mounted until you have started CXFS services. For information, see "Start CXFS Services with the GUI" on page 227.

Unmount CXFS Filesystems with the GUI

To unmount filesystems from all of their client nodes, do the following:

1. Enter the following information:
 - **Filesystem to Unmount:** Choose the filesystems to be unmounted.
 - **Force Unmount :** Click **On** to force an unmount for all selected filesystems (no matter how they have been defined) or **Default** to force an unmount for those filesystems that have the forced unmount option set in their definition.

This option controls what action CXFS takes if there are processes that have open files or current directories in the filesystems that are to be unmounted. If forced is used (by selecting **On** or by selecting **Default** if force is the default behavior), the processes will be killed and the unmount will occur. If you select **Off**, the processes will not be killed and the filesystem will not be unmounted. The option is set to **Default** by default.

2. Click **OK**.

Mount a Filesystem Locally

This task lets you mount a filesystem only on the node to which the GUI is connected (the local node).

To mount a filesystem locally, do the following:

1. Enter the following information:
 - **Filesystem to Mount:** Select the filesystem you wish to mount. The list of available filesystems is determined by looking for block devices containing XFS superblocks.
 - **Mount Point:** Specify the directory on which the selected filesystem will be mounted.
 - *(Optional)* **Mount Options:** Specify the options that should be passed to the mount command. For more information about available options, see the `fstab` man page.
2. By default, the filesystem will remount every time the system starts. However, if you uncheck the box, the mount will take place only when you explicitly use this task.

3. Click **OK**.

For more information, see the `mount` man page.

Unmount a Local Filesystem

To unmount a filesystem from the local node, do the following:

1. Enter the following information:
 - **Filesystem to Unmount:** Choose the filesystem to be unmounted.
 - **Remove Mount Information:** Click the check box to remove the mount point from the `/etc/fstab` file, which will ensure that the filesystem will remain unmounted after the next reboot. This item is available only if the mount point is currently saved in `/etc/fstab`.
2. Click **OK**.

Delete a CXFS Filesystem with the GUI

You cannot delete a filesystem that is currently mounted. To unmount a filesystem, see "Unmount CXFS Filesystems with the GUI" on page 246.

To permanently delete an unmounted filesystem, do the following:

1. **Filesystem to Delete:** Choose the name of the filesystem from the pull-down list.
2. Click **OK**.

Remove Filesystem Mount Information

This task lets you delete a local filesystem's mount information in `/etc/fstab`.

Note: The filesystem will still be present on the volume.

Do the following:

1. **Filesystem Name:** Select the filesystem for which you want to remove mount information. The list of available filesystems is determined by looking for block devices containing XFS superblocks.
2. Click **OK**.

Relocate a Metadata Server for a CXFS Filesystem with the GUI

If relocation is explicitly enabled in the kernel with the `cxfs_relocation_ok` systune, you can relocate the metadata server for a filesystem to any other potential metadata server in the list (see "Relocation" on page 20). The filesystem must be mounted on the system to which the GUI is connected.

1. Enter the following information:
 - **Filesystem:** Select the desired filesystem from the list.
 - **Current Metadata Server:** The current metadata server will be displayed for you.
 - **New Metadata Server:** Select the desired node from the list.

The selected server will assume responsibility for moderating access to the selected filesystem **after** you run the **Start CXFS Services** task; see "Start CXFS Services with the GUI" on page 227.

2. Click **OK** to complete the task.

CXFS kernel membership is not affected by relocation. However, users may experience a degradation in filesystem performance while the metadata server is relocating.

Privileges Tasks with the GUI

The privileges tasks let you grant specific users the ability to perform specific tasks and to revoke those privileges.

Note: You cannot grant or revoke tasks for users with a user ID of 0.

This section discusses the following:

- "Grant Task Access to a User or Users" on page 249
- "Revoke Task Access from a User or Users" on page 252

Grant Task Access to a User or Users

You can grant access to a specific task to one or more users at a time.

Note: Access to the task is only allowed on the node to which the GUI is connected; if you want to allow access on another node in the pool, you must connect the GUI to that node and grant access again.

Do the following:

1. Select the user or users for whom you want to grant access. You can use the following methods to select users:
 - Click to select one user at a time
 - Shift+click to select a block of users
 - Ctrl+click to toggle the selection of any one user, which allows you to select multiple users that are not contiguous
 - Click **Select All** to select all usersClick **Next** to move to the next page.
2. Select the task or tasks to grant access to, using the above selection methods. Click **Next** to move to the next page.
3. Confirm your choices by clicking **OK**.

Note: If more tasks than you selected are shown, then the selected tasks run the same underlying privileged commands as other tasks, such that access to the tasks you specified cannot be granted without also granting access to these additional tasks.

To see which tasks a specific user can access, select **View: Users**. Select a specific user to see details about the tasks available to that user.

To see which users can access a specific task, select **View: Task Privileges**. Select a specific task to see details about the users who can access it and the privileged commands it requires.

Granting Access to a Few Tasks

Suppose you wanted to grant user `diag` permission to define, modify, and mount CXFS filesystems. You would do the following:

1. Select `diag` and click **Next** to move to the next page.
2. Select the tasks you want `diag` to be able to execute:
 - a. **Ctrl+click Define CXFS Filesystem**
 - b. **Ctrl+click Modify CXFS Filesystem**
 - c. **Ctrl+click Mount CXFS Filesystem**

Click **Next** to move to the next page.

3. Confirm your choices by clicking **OK**.

Figure 11-6 shows the tasks that `diag` can now execute. This screen is displayed when you select **View: Users** and click `diag` to display information in the details area of the GUI window. The privileged commands listed are the underlying commands executed by the GUI tasks.

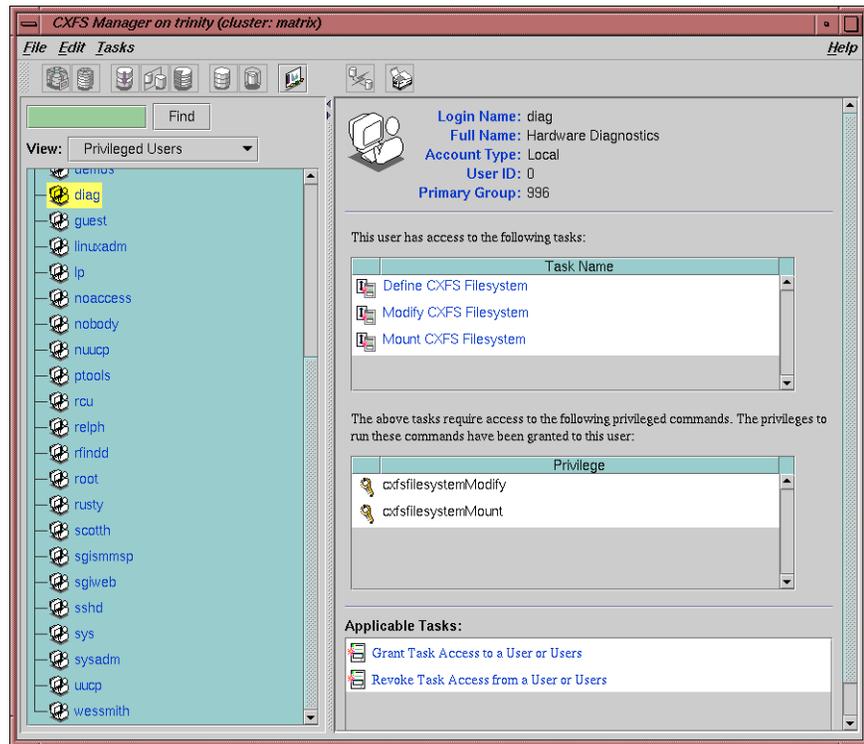


Figure 11-6 Task Privileges for a Specific User

Granting Access to Most Tasks

Suppose you wanted to give user `sys` access to all tasks **except** changing the cluster contents (which also implies that `sys` cannot delete the nodes in the cluster, nor the cluster itself). The easiest way to do this is to select all of the tasks and then deselect the few you want to restrict. You would do the following:

1. Select `sys` and click **Next** to move to the next page.
2. Select the tasks you want `sys` to be able to execute:
 - a. Click **Select All** to highlight all tasks.
 - b. Deselect the task to which you want to restrict access. **Ctrl+click Add/Remove Nodes in Cluster.**

Click **Next** to move to the next page.

3. Confirm your choices by clicking **OK**.

Revoke Task Access from a User or Users

You can revoke task access from one or more users at a time.

Note: Access to the task is only revoked on the node to which the GUI is connected; if a user has access to the task on multiple nodes in the pool, you must connect the GUI to those other nodes and revoke access again.

Do the following:

1. Select the user or users from whom you want to revoke task access. You can use the following methods to select users:
 - Click to select one user at a time
 - Shift+click to select a block of users
 - Ctrl+click to toggle the selection of any one user, which allows you to select multiple users that are not contiguous
 - Click **Select All** to select all users

Click **Next** to move to the next page.

2. Select the task or tasks to revoke access to, using the above selection methods. Click **Next** to move to the next page.
 3. Confirm your choices by clicking **OK**.
-

Note: If more tasks than you selected are shown, then the selected tasks run the same underlying privileged commands as other tasks, such that access to the tasks you specified cannot be revoked without also revoking access to these additional tasks.

To see which tasks a specific user can access, select **View: Users**. Select a specific user to see details about the tasks available to that user.

To see which users can access a specific task, select **View: Task Privileges**. Select a specific task to see details about the users who can access it.

Reference to `cxfs_admin` Tasks

For an overview of the tasks that must be performed to configure a cluster, see "Initial Setup with the `cxfs_admin` Command" on page 172.

You can also use the `clconf_info` tool to view status. See Chapter 17, "Monitoring Status" on page 399.

For help with error messages, see "`cxfs_admin` Errors" on page 482.

This chapter discusses the following:

- "`cxfs_admin` Overview" on page 255
- "Node Tasks with `cxfs_admin`" on page 270
- "Cluster Tasks with `cxfs_admin`" on page 282
- "CXFS Filesystem Tasks with `cxfs_admin`" on page 287
- "Network Failover Tasks with `cxfs_admin`" on page 295
- "Switch Tasks with `cxfs_admin`" on page 295
- "Saving and Recreating the Current Configuration with `cxfs_admin`" on page 300

Note: The example output might not exactly match the output of your system.

See also Appendix H, "Migration from `cmgr` to `cxfs_admin`" on page 659.

`cxfs_admin` Overview

To use the `cxfs_admin` command, you must be logged in as `root` on a server-capable node that has permission to access the CXFS cluster database. See "Setting `cxfs_admin` Access Permissions" on page 268.

Note: For the steps to create a cluster for the first time, see "Initial Setup with the `cxfs_admin` Command" on page 172.

This section discusses the following:

- "Command Syntax Overview" on page 256
- "Getting Help" on page 258
- "Making Changes Safely" on page 259
- "Basic and Advanced Mode" on page 260
- "Using Prompting Mode" on page 262
- "Command History" on page 263
- "Waiting for Commands to Complete" on page 264
- "Entering `cxfs_admin` Commands on the Command Line" on page 264
- "Using Script Files" on page 265
- "Setting `cxfs_admin` Defaults" on page 266
- "Setting `cxfs_admin` Access Permissions" on page 268
- "Exiting from `cxfs_admin`" on page 269

Command Syntax Overview

Some `cxfs_admin` commands affect the `cxfs_admin` operating environment itself, some display status information, and others affect objects or classes. Within `cxfs_admin`, an *object* is a specific item that is configured in the CXFS cluster and a *class* contains a group of similar objects. For example, the filesystem names `fs1` and `fs2` would both be objects within the `filesystem` class.

Within a class, all objects must have unique names. If all objects in the cluster have unique names, you can abbreviate some commands by omitting the class name. However, if two or more objects in the cluster database have the same name, you must specify the class in order to uniquely identify the object.

The basic command syntax is:

```
command [[class:object] [attributes]
```

where *attributes* takes a number of forms depending on the context:

```
attribute  
attribute=value  
attribute=value1,value2,value3...
```

The actual syntax components for any given command varies, based on the needs of the command. For example, the following command requires no parameters to see a summary of the cluster:

```
cxfs_admin:mycluster> show
```

If an object name is unique within the cluster database, you can omit its class name. For example, if the name `nodeA` is unique within the database:

```
cxfs_admin:mycluster> show nodeA
```

However, if there were multiple objects named `production`, you must enter the class name:

```
cxfs_admin:mycluster> show node:production
```

Classes and objects may include the following shell-style wildcard characters:

```
*
?
[...]
```

Command names and attribute names are not case-sensitive. However, all attribute values other the node name (in `create node nodename`) and the switch name (in `create switch switchname`) are case sensitive.

You can see possible attributes by pressing the <TAB> key after entering the command or object. For example:

```
cxfs_admin:mycluster> create filesystem <TAB>
Required attributes:
  name= : A string

Optional attributes:
  forced_unmount= : True/false or enabled/disabled (default is "false")
  mounted=       : True/false or enabled/disabled (default is "true")
  mountpoint=    : A pathname
  options=       : Nothing, one or more strings (can be empty)
```

The required attributes are listed first followed by optional attributes. The list of attributes will differ depending upon whether you are in basic or advanced mode; see "Basic and Advanced Mode" on page 260.

Partially typing in the attribute name and pressing <TAB> will complete the attribute name if unique, or show a list of matching attribute names. To see what kind of values are required for an attribute, press <TAB> after the = sign. For example:

```
cxfs_admin:mycluster> create node os=<TAB>
AIX          IRIX      Linux      MacOSX    Solaris    Unknown    Windows
```

Use \$ to refer to the object in the last command.

For example, to delete nodeA, if it has a unique name within the cluster database:

```
cxfs_admin:mycluster> disable nodeA
cxfs_admin:mycluster> delete $
```

To specify multiple objects, separate them with a comma. For example:

```
cxfs_admin:mycluster> show nodeA,nodeB
```

With the show and config commands, you can use a * character in place of the *object* to apply the command to entire cluster. If you do not specify any attributes, you can omit the * character.

You can abbreviate commands, objects, and attributes by entering in the first character or two followed by pressing the <TAB> key. If more than one match is available, cxfs_admin shows a list of the possible matches.

Getting Help

At any time, you can enter help or ? to see help text.

To see help for a given topic:

```
help topicname
```

The list of general topics includes the following:

```
attributes
commandline
commands
cxfs
objects
overview
setup
syntax
```

```
tasks
waiting
```

There is also help for each `cxfs_admin` command. For example, to see help about the `create` command:

```
cxfs_admin:mycluster> help create
```

To see all of the available help topics, press the <TAB> key:

```
cxfs_admin:mycluster> help <TAB>
```

To see a list of available commands for an object, such as a class like `filesystem` or a specific instance of a class like the `filesystem myfs`, use the `ops` command:

```
ops object
```

For example:

```
cxfs_admin:mycluster> ops filesystem
Commands for "filesystem":
    config, create, ops, show
cxfs_admin:mycluster> ops myfs
Commands for "filesystem:myfs":
    delete, modify, mount, ops, show, unmount
```

Making Changes Safely

The `cxfs_admin` tool only allows one user to use `cxfs_admin` to make changes to the cluster database at a time. If you are the first person to invoke `cxfs_admin`, you automatically get the lock. If someone else already has the lock, you will enter in read-only mode. If you are in read-only mode, it is reflected in the `cxfs_admin` prompt.

To forcefully obtain the lock from someone else, you can use the `steal` attribute with the `lock` command. For example:

```
cxfs_admin:mycluster (read only) > lock
The administration lock is already held by root@node2 (pid=48449)
cxfs_admin:mycluster (read only) > lock steal=true
The administration lock has been stolen from root@node2 (pid=48449)
cxfs_admin:mycluster>
```

If someone holds the lock while you are using `cxfs_admin` but later drops it, there is no need to steal the lock.

If you want to manually enter read-only mode, use the `unlock` command. For example:

```
cxfs_admin:mycluster> unlock
cxfs_admin:mycluster (read only) >
```



Caution: The `cxfs_admin` lock does not prevent other users from using the CXFS GUI or the `cmgr` command while `cxfs_admin` is running. You should make database changes with only one instance of the CXFS GUI, `cmgr`, or locked `cxfs_admin` commands at any one time.

Basic and Advanced Mode

The `cxfs_admin` operates in two modes:

- Basic, which only shows the common options and attributes in `show` output, `<TAB>` key completion, and prompting mode.
- Advanced, which allows `<TAB>` key completion, prompts for all possible fields, displays all attributes, and includes debugging information in output. Advanced-mode commands and attributes are not included in prompts or `<TAB>` key completion when you are in basic mode. However, you can still manually enter an advanced attribute if you know it, even in basic mode. The advanced commands and attributes are noted in their help topics.

Note: You should only use the advanced-mode commands and attributes at the advice of SGI support. Using the advanced mode commands or changing advanced mode attributes may induce unexpected behavior.

You can enter advanced mode by using `cxfs_admin -a` on the command line or by entering the following `cxfs_admin` command:

```
cxfs_admin:mycluster> set mode=advanced
```

To return to basic mode:

```
cxfs_admin:mycluster> set mode=basic
```

For example, the following output shows only the basic-mode information:

```
cxfs_admin:mycluster> set mode=basic
Mode is set to basic
cxfs_admin:mycluster> show node2
node:node2:
  cellid=2
  enabled=true
  os=Linux
  private_net:
    192.168.0.130, 192.0.34.130
  status:
    client=stable
    connected=true
    fencing=Stable
    filesystems=up
    license:
      cpu_count=4
      have_license=true
      oem=none
      os=Linux64
      version=
    membership=up
    summary=Stable
    version=4.2.0.1
    wwns:
      210000e08b0e2359
    xvm=up
  type=client_only
```

For example, the following output shows all information that is available in advanced mode for node2:

```
cxfs_admin:cfi> set mode=advanced
Mode is set to advanced
cxfs_admin:cfi> show node2
node:node2:
  admin_node=false
  cellid=2
  clustername=cfi
  enabled=true
  failpolicy=Fence
```

```
hostname=node.example.com
nodeid=130
os=Linux
private_net:
    192.168.0.130, 192.0.34.130
server_capable=false
status:
    build=18:56:44 Mar 25 2007
    client=stable
    connected=true
    fencing=Stable
    filesystems=up
    license:
        cpu_count=4
        have_license=true
        oem=none
        os=Linux64
        version=
    member=true
    membership=up
    stable=true
    summary=Stable
    version=4.2.0.1
    wwns:
        210000e08b0e2359
    xvm=up
type=client_only
```

Note: If a client is not connected to the cluster, the `build` and `version` fields will not display because the node cannot respond to for requests for this information.

Using Prompting Mode

Some `cxfs_admin` commands will prompt you for required attributes if you press `ENTER` after the command name. To see information about the legal values for an attribute, press `<TAB>` after a question.

For example:

```
cxfs_admin:mycluster> create
What do you want to create? The following can be used:
    failover_net, filesystem, node, switch
create what? node
Specify the attributes for create node:
name? mynode
os? <TAB>
    AIX          IRIX          Linux          MacOSX        Solaris        Unknown        Windows
os?
```

In basic mode, you are only prompted for required parameters. To be prompted for all possible parameters, use advanced mode. See "Basic and Advanced Mode" on page 260.

Depending upon the context, `cxfs_admin` prompts will vary based upon your answers to previous prompts. For example, if you specify that a node's `os` value is `MacOSX`, `cxfs_admin` will not prompt you for the `type` because Mac OS X nodes are required to be client-only nodes.

To exit from prompt mode, sent an interrupt signal (typically, press `Ctrl-C`).

Command History

The `history` command displays a list of commands that have been used in `cxfs_admin` since it was started:

- Display all of the commands (up to the previous 1000 commands):

```
history
```

- Limit the commands to the last specified number of items:

```
history num=number_of_items
```

For example, to display only the last 10 commands:

```
cxfs_admin:mycluster> history num=10
```

- Clear the history:

```
history clear
```

- Send the history to a file (you must enter the full pathname of the file):

```
history output=full_pathname
```

For example, to send the history output to the file /tmp/myhistory:

```
cxfs_admin:mycluster> history output=/tmp/myhistory
```

Waiting for Commands to Complete

Some commands in `cxfs_admin` take a noticeable period of time to complete. `cxfs_admin` displays informational updates as a command progresses or a period character if nothing has changed within 2 seconds.

After 1 minute without change, a command will terminate. This may happen when there is a problem in creating or modifying a node or filesystem. The update message shows the problem status.

To interrupt a command, send an interrupt signal (usually `Ctrl-C`).

Entering cxfs_admin Commands on the Command Line

You can enter `cxfs_admin` commands directly from the `cxfs_admin` command line by using the following format:

```
# cxfs_admin -c "cxfs_admin_commands"
```

For example, to display information about the cluster:

```
# /usr/cluster/bin/cxfs_admin -c "show cluster"
cxfs:cluster:
  mycluster:
    access:
      admin=server
      monitor=
    failover_net:
      10.11.0.0, 128.162.242.0
    filesystem:
      concatfs, mirrorfs, stripefs
    node:
      aix-client, irix-client, linux-client, mac-client, mds1, mds2,
      solaris-client, windows-client,
```

```
status:
  filesystems:
    summary=concatfs: mac-client trying to mount
    mirrorfs: mac-client trying to mount
    stripefs: mac-client trying to mount
  licenses:
    cxfs_client
  nodes:
    summary=aix-client: Inactive
    mac-client: Mounted 0 of 3 filesystems
    solaris-client: Inactive
    windows-client: Inactive
    summary=node(s) not stable, filesystem(s) not stable
  switch:
    fcswitch12, fcswitch13
  tiebreaker=irix-client
```

Using Script Files

You can execute a series of `cxfs_admin` commands by using the `-f` option and specifying an input file:

```
# cxfs_admin -f command_file
```

For example, suppose the file `/tmp/showme` contains the following:

```
cxfs6# more /tmp/showme
show cluster
show filesystem
```

You can execute the following command, which will yield the indicated output:

```
# /usr/cluster/bin/cxfs_admin -f /tmp/showme
Connecting to the CXFS server for the "mycluster" cluster...
cxfs:cluster:
  mycluster:
    access:
      admin=server
      monitor=
    failover_net:
      (none)
    filesystem:
```

```
    thump
node:
  node1, node2
status:
  filesystems:
    summary=Stable
  licenses:
    cxfs_client
  nodes:
    summary=node2: Fencing issue: No switches configured
    summary=node(s) not stable
switch:
  (none)
tiebreaker=
filesystem:
  myfs:
    forced_unmount=false
    mount=false
    mountpoint=/myfs
    nodes:
      node2, node1
    options=
    servers:
      node1
    status:
      nodes:
        node2=unmounted
        node1=unmounted
      summary=Unmounted
```

Setting cxfs_admin Defaults

You can use one of the following methods to set the defaults for the way `cxfs_admin` behaves and the editor to use within `cxfs_admin`, in the following order of precedence:

1. Use the `set` command within `cxfs_admin`:

```
set
  [editor=emacs|vi]           (emacs)
  [line_wrap=true|false]     (true)
```

```
[mode=basic|advanced]      (basic)
[stop_on_error=true|false] (true)
```

For example, to change to vi:

```
cxfs_admin:mycluster> set editor=vi
```

Usage notes:

- `editor` specifies the editor style (`emacs` or `vi`). The default is `emacs`.
- `line_wrap` specifies the ability to wrap a line at the edge of the current window (`true`) or no line wrap (`false`). The default is `true`.
- `mode` determines whether all values (`advanced`) or only those values that are required (`basic`). The default is `basic`. See "Basic and Advanced Mode" on page 260.
- `stop_on_error` will abort a command upon encountering an error (`true`) or keep going (`false`). The default is `true`.

2. Set the following environment variables:

Environment Variable	Values
<code>CXFS_ADMIN_CLUSTER_NAME</code>	<i>clustername</i> . Setting this value lets you bypass using the <code>-i</code> option if you have multiple clusters using the same public network as the backup CXFS metadata network. There is no default.
<code>CXFS_ADMIN_EDITOR</code>	<code>emacs</code> (default) or <code>vi</code>
<code>CXFS_ADMIN_LINE_WRAP</code>	<code>true</code> (default) or <code>false</code>
<code>CXFS_ADMIN_MODE</code>	<code>basic</code> (default) or <code>advanced</code>
<code>CXFS_ADMIN_STOP_ON_ERROR</code>	<code>true</code> (default) or <code>false</code>

3. Use the `.cxfs_admin` file in your home directory (as defined by the `$HOME` environment variable) to set the following:

```
mode=basic|advanced
cluster_name=clustername
editor=emacs|vi
stop_on_error=true|false
line_wrap=true|false
```

Lines within the `.cxfs_admin` file that begin with the `#` character or a space are ignored, as are lines that do not contain the `=` character.

For example, to use the `mycluster` cluster in advanced mode and the `vi` editor:

```
# My settings for cxfs_admin:
cluster=mycluster
mode=advanced
editor=vi
```

Setting `cxfs_admin` Access Permissions

The `access` command allows you to specify hosts that have permission to modify the cluster configuration and hosts that have permission to monitor the cluster state:

```
access
  allow=hostname_or_IPaddress_list
    permission=admin|monitor           (monitor, only available with allow)
  deny=server_name
```

By default, all server-capable nodes in the cluster are granted `admin` access (without using the `access` command).

For example, to grant `remotehostA` and `remotehostB` permission to modify the cluster configuration:

```
cxfs_admin:mycluster> access allow=remotehostA,remotehostB permission=admin
```

To grant read-only rights in order to monitor to the cluster configuration and status (`monitor` is the default access level):

```
cxfs_admin:mycluster> access allow=remotehostA
```

To revoke all access to the cluster database for a host that was previously granted some level of access, use the following command:

```
cxfs_admin:mycluster> access deny=remotehostA,remotehostB
```

To view the current access rights, use the following command:

```
show access
```

For example:

```
cxfs_admin:mycluster> show access
access:
```

```
admin=server  
monitor=cluster
```

Usage notes:

- `allow` specifies the hosts to be granted the specified permission. These hosts must be on the same private network as the cluster nodes. To specify multiple hosts, use a comma-separated list. There are three reserved hostnames:
 - `cluster` denotes any node defined in the cluster
 - `server` denotes any server-capable node, even one that is disabled from CXFS membership (see "Disable a Node with `cxfs_admin`" on page 279)
 - `any` denotes any system that is on the private network
- `permission` specifies read/write access (`admin`) or read-only access (`monitor`). The default is `monitor`.
- `deny` specifies the hosts to be denied all access to the cluster database (for hosts that were previously granted some level of access). To specify multiple hosts, use a comma-separated list. The same reserved hostnames as `allow` apply.

Accessing the Correct Cluster at a Multiple-Cluster Site

If you have multiple clusters using the same public network as the backup CXFS metadata network, use the `-i` option to identify the cluster name. For example:

```
# cxfs_admin -i mycluster
```

Exiting from `cxfs_admin`

To exit from prompt mode, sent an interrupt signal (typically, press `Ctrl-C`).

To exit out of the `cxfs_admin` session, enter `exit` or `quit` at the `cxfs_admin` command line:

```
cxfs_admin:mycluster> exit
```

Node Tasks with `cxfs_admin`

This section discusses the following:

- "Create or Modify a Node with `cxfs_admin`" on page 270
- "Delete a Node with `cxfs_admin`" on page 279
- "Enable a Node with `cxfs_admin`" on page 279
- "Disable a Node with `cxfs_admin`" on page 279
- "Show Node Information with `cxfs_admin`" on page 280

Note: The entire cluster status information is sent to each CXFS administration node each time a change is made to the cluster database; therefore, the more CXFS administration nodes in a configuration, the longer it will take.

Create or Modify a Node with `cxfs_admin`

To define a node, use the following command and attributes (line breaks shown here for readability, defaults in parentheses):

```
create node
  name=nodename
  os=AIX|IRIX|Linux|MacOSX|Solaris|Windows|Unknown
  private_net private_network_IPaddress_list|hostname_list
  type=client_admin|client_only|server_admin           (client_only)
Advanced-mode:
  enabled=true|false                                  (true)
  failpolicy=FenceReset,Fence,Reset,Shutdown         (Fence,Shutdown)
  hostname=logical_hostname                        (fully_qualified_domain_name_of_nodename)
  nodeid=nodeID                                     (assigned by cxfs_admin)
  partition_id=partition_number
  reset_method=nmi|powerCycle|reset                  (powerCycle)
  reset_port=11|12|bmc|mhc|mssc
  reset_password=password
  reset_status=enabled|disabled                       (enabled)
  reset_node=node_sending_reset_command
  reset_comms=tty|network|ipmi
  reset_device=port|IP_address_or_hostname_of_device
```

When you create a client-only node, it will by default automatically be enabled and join the cluster. When adding the first server-capable node, you must restart it or restart CXFS services and cluster services on the node:

- IRIX:

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

- SGI ProPack:

```
# /etc/init.d/cxfs stop
# /etc/init.d/cxfs_cluster stop
# /etc/init.d/cxfs_cluster start
# /etc/init.d/cxfs start
```

To use prompting mode, press <ENTER>. To obtain information about legal values, press <TAB>.

For example, to create a client-only node, you could do the following, pressing the <TAB> key to see the list of operating system values:

```
cxfs_admin:mycluster> create node
Specify the attributes for create node:
name? newnode
os? <TAB>
  AIX      IRIX      Linux      MacOSX     Solaris    Unknown    Windows
os? irix
private_net? 192.168.0.178
type? client_only
Node "newnode" has been created, waiting for it to join the cluster...
Waiting for node newnode, current status: Inactive
Waiting for node newnode, current status: Establishing membership
Waiting for node newnode, current status: Probing XVM volumes
Operation completed successfully
```

Note: A cluster can have server-capable nodes of only one operating system type: either all IRIX or all Linux. Therefore, if you are adding a new node to an existing cluster that already contains Linux server-capable nodes, `cxfs_admin` assumes that any IRIX node added must be a client-only node. Similarly, if the cluster already contained IRIX server-capable nodes, then `cxfs_admin` assumes that a new Linux node must be a client-only node, but it must prompt you for the `type` value for a new IRIX node, because that node could be server-capable or client-only, as shown above.

To create a server-capable administration node using the defaults, you must delete the `client_only` default for `type` and enter `server_admin`. For example:

```
cxfs_admin:mycluster> create node
Specify the attributes for create node:
name? newnode
os? irix
private_net? 192.168.0.178
type? server_admin
Node "newnode" has been created, waiting for it to join the cluster...
Please restart all cxfs and cluster services on the server "newnode" to make it
join the cluster.
```

To create a server-capable administration node in advanced mode, which can prompt you to set additional values, such as for `reset_method` and `failpolicy`:

```
cxfs_admin:mycluster> set mode=advanced
cxfs_admin:mycluster> create node
Specify the attributes for create node:
name? newnode
os? irix
private_net? 192.168.0.178
type? server_admin
hostname? newnode.example.com
enabled? true
failpolicy? Reset,Shutdown
nodeid? 1
partition_id?
reset_method? reset
reset_port? 12
reset_password?
reset_status? enabled
```

```
reset_node? node2  
reset_comms? network  
reset_device? newnode-12.mycompany.com
```

Node "newnode" has been created, waiting for it to join the cluster...
Please restart all cxfs and cluster services on the server "newnode" to make it join the cluster.

To modify the failpolicy to eliminate Shutdown:

```
cxfs_admin:mycluster> modify newnode failpolicy=Reset,Fence
```

Basic-mode usage notes:

- name is a simple hostname (such as lilly) or a fully qualified domain name (such as lilly.example.com) or an entirely different name (such as node1). It cannot begin with a number or an underscore (_), or include any whitespace characters, and can be at most 255 characters.

- os is one of the following:

```
AIX  
IRIX  
Linux (SGI ProPack or Linux third-party)  
MacOSX  
Solaris  
Windows  
Unknown
```

- private_net is the IP address or hostname of the private network. (The hostname must be resolved in the /etc/hosts file.) SGI requires that this network be private; see "Private Network" on page 19.

There can be up to 8 network interfaces. There is no default.

For more information about using the hostname, see "Hostname Resolution and Network Configuration Rules" on page 93.

- type specifies the function of the node. Enter one of the following:
 - client_only is a node that shares CXFS filesystems but will never be a CXFS metadata server. Most nodes should be client-only nodes. AIX, MacOSX, Solaris, Windows and Unknown nodes are automatically specified as client-only and you will not be prompted for this value for these operating systems.

- `server_admin` is an IRIX or Linux node that is a potential CXFS metadata server. (You will use the `create filesystem` command to define the specific filesystem for which this node can be a metadata server.)
- `client_admin` is an IRIX node that runs FailSafe but that will never be a CXFS metadata server. This is provided for backwards compatibility purposes for FailSafe administration nodes.

Advanced-mode usage notes:

- `enabled` determines if a node will be able to obtain CXFS membership (`true`) or not (`false`). By default, the new node is enabled (`true`). To enable a command created with `enabled=false`, use the `enable` command. See "Enable a Node with `cxfs_admin`" on page 279.
- `failpolicy` determines what happens to a failed node. You can specify up to three methods. The second method will be completed only if the first method fails; the third method will be completed only if both the first and second options fail. Separate options by commas (not whitespace). The option choices are as follows:
 - `Fence` disables access to the SAN from the problem node. Fencing provides faster recovery of the CXFS kernel membership than `reset`.
 - `FenceReset` performs a fence and then, if the node is successfully fenced, **also** performs an asynchronous reset of the node via a system controller; recovery begins without waiting for reset acknowledgement.

Note: SGI recommends that a server-capable node include `Reset` in its `failpolicy` (unless it is the only server-capable node in the cluster). See "Isolating Failed Nodes: Failure Policies" on page 28.

The `FenceReset` and `Fence` policies are mutually exclusive.

- `Reset` performs a system reset via a system controller. This action requires a `reset_method` value; see "Requirements" on page 38.
- `Shutdown` tells the other nodes in the cluster to wait for a period of time (long enough for the node to shut itself down) before reforming the CXFS kernel membership. (However, there is no notification that the node's shutdown has actually taken place.)



Caution: Because there is no notification that a shutdown has occurred, if you have a cluster with no tiebreaker, you must not use the `shutdown` setting for any server-capable node in order to avoid multiple clusters being formed. See "Shutdown" on page 35.

You should not use the Shutdown failure policy on client nodes if you choose dynamic heartbeat monitoring for the cluster.

Note: If the failure hierarchy contains `reset` or `fencerereset`, the reset might be performed before the system kernel core-dump can complete, resulting in an incomplete core-dump.

For a list of valid failpolicy sets, see "Isolating Failed Nodes: Failure Policies" on page 28.

For example, to perform a reset only if a fencing action fails, specify the following:

```
failpolicy=Fence,Reset
```

Note: If you do not specify `Shutdown` and all of the other methods fail, the node attempting to deliver the CXFS kernel membership will stall delivering the membership until either the failed node attempts to re-enter the cluster or the system administrator intervenes using `cms_intervene`. Objects held by the failed node stall until membership finally transitions and initiates recovery.

To perform a fence and an asynchronous reset, specify the following:

```
failpolicy=FenceReset
```

- `hostname` is by the fully qualified hostname. Use the `ping` to display the fully qualified hostname. Do not enter an IP address. The default for `hostname` is the fully qualified domain name for the value of `name`.
- `nodeid` is an integer in the range 1 through 32767 that is unique among the nodes in the cluster. If you change this value after a node has been defined, you must reboot the affected node. You do not normally need to specify this attribute because `cxfs_admin` will calculate an ID for you.

- `partition_id` uniquely defines a partition in a partitioned Origin 3000 system, Altix 3000 series system, or Altix 4700 system. For a non-partitioned system, this attribute is not required (the default unassigned).

Note: For an Origin 3000 series system, use the `mkpart` command to determine this value:

- The `-n` option lists the partition ID (which is 0 if the system is not partitioned).
- The `-l` option lists the bricks in the various partitions (use `rack#.slot#` format in `cxfs_admin`)

For example (output truncated here for readability):

```
# mkpart -n
Partition id = 1
# mkpart -l
partition: 3 = brick: 003c10 003c13 003c16 003c21 003c24 003c29 ...
partition: 1 = brick: 001c10 001c13 001c16 001c21 001c24 001c29 ...
```

To unset the partition ID, use a value of 0.

For an Altix 3000, you can find the partition ID by reading the `proc` file. For example:

```
[root@linux root]# cat /proc/sgi_sn/partition_id
0
```

The 0 indicates that the system is not partitioned. If the system is partitioned, the number of partitions (such as 1, 2, etc.) is displayed.

- `reset_method` can be one of the following:
 - `powerCycle` shuts off power to the node and then restarts it
 - `reset` simulates the pressing of the reset button on the front of the machine
 - `nmi` (nonmaskable interrupt) performs a core-dump of the operating system kernel, which may be useful when debugging a faulty machine

Note: `nmi` is not available on systems containing a baseboard management controller (BMC).

The default is `powerCycle`.

- `reset_port` is the system controller port type based on the node hardware, as show in Table 12-1 on page 278.
- `reset_password` is the password for the node's system controller port (not the node's `root` password or PROM password). On some nodes, the system administrator may not have set this password. If you wish to set or change the system controller password, consult the hardware manual for your node.
- `reset_status` specifies if the system reset capability is turned on (`enabled`) or turned off (`disabled`). Using `disabled` allows you to provide information about the system controller but temporarily disable reset (meaning that CXFS cannot reset the node). The default for nodes with system controllers is `enabled`, for nodes without system controllers, the default is `disabled`; see "Requirements" on page 38.
- `reset_node` specifies the name of the node that is to send the reset command. It can be the logical name, hostname, or fully qualified domain name. If you use `reset_comms=tty`, serial cables must physically connect the node being defined and the owner node through the system controller port. The node must be a member of the cluster.
- `reset_comms` is `tty` for TTY serial devices, `network` for network reset to systems with L2 system controllers, or `ipmi` for intelligent platform management interface (IPMI) network reset to systems with BMC system controllers.
- `reset_device` is one of the following:
 - For systems with serial ports (`reset_comms=tty`), this is the name of the terminal port (TTY) on the owner node (the node issuing the reset). A serial cable connects the terminal port on the owner node to the system controller of the node being reset. `/dev/ttyd2` is the most commonly used port, except on Origin 300 and Origin 350 systems (where `/dev/ttyd4` is commonly used) and Altix 350 systems (where `/dev/ttyIOC0` is commonly used).

Note: Check the owner node's specific hardware configuration to verify which tty device to use.

- For systems with network-attached L2 system controllers (`reset_comms=network`), this is the IP address or hostname of the L2 controller on the node being reset. For example, `reset_device=nodename-l2.mycompany.com`.

- For systems with network-attached BMC system controllers (reset_comms=ipmi), this is the IP address or hostname of the BMC controller on the node being reset. For example, reset_device=nodename-bmc.mycompany.com.

For example:

For an Origin 3000 series system:

```
reset_comms=tty reset_device=/dev/ttyd2
```

For an SGI Altix 3000 Bx2 system:

```
reset_comms=network reset_device=nodename-12.mycompany.com
```

For an Altix 350 system without an L2:

```
reset_comms=tty reset_device=/dev/ttyIOC0
```

For an Altix XE system with a BMC:

```
reset_comms=ipmi reset_device=nodename-bmc.mycompany.com
```

Table 12-1 System Controller Types

bmc	11	12	mmsc	msc
Any Altix XE	Origin/Onyx 300/350	Any Altix with an L2	Rackmount SGI 2400/2800	Origin 200
	Origin/Onyx 3200C	Prism	Onyx2	Onyx2 Deskside
		Origin/Onyx 3000 series		SGI 2100/2200 deskside systems
		Origin 300/350		

Delete a Node with `cxfs_admin`

To delete a node from the cluster and the cluster database, use the following command:

```
delete [node:]nodename
```

If the node is enabled (which is the default), you must disable it before you delete it. For example, if `mynode` is a unique name in the cluster database:

```
cxfs_admin:mycluster> disable mynode  
cxfs_admin:mycluster> delete mynode
```

Note: If you delete an active metadata server, `cxfs_admin` will enter read-only mode. You can use the `lock` or `lock steal=true` to reenter lock mode. For more information, see "Making Changes Safely" on page 259.

Enable a Node with `cxfs_admin`

To allow a disabled node to join the cluster, enter the following:

```
enable [node:]nodename
```

For example, if `node1` is a unique name in the cluster database:

```
cxfs_admin:mycluster> enable node1
```

Disable a Node with `cxfs_admin`

To prevent a node from joining the cluster, enter the following:

```
disable [node:]nodename
```

For example, if `node1` is a unique name in the cluster database:

```
cxfs_admin:mycluster> disable node1
```

Note: This procedure is only recommended as needed for a CXFS server-capable node because it updates the cluster database and is therefore intrusive to other nodes. When shutting down a CXFS client-only node, do not disable it. Rather, let the CXFS services stop by themselves when the client-only node is shut down.

After you have disabled a node, the node is no longer an active member of the cluster.



Caution: If you disable a node, it will be marked as `Disabled` and it will therefore not rejoin the cluster after a reboot. To allow a node to rejoin the cluster, you must enable the node. See "Enable a Node with `cxfs_admin`" on page 279.

Show Node Information with `cxfs_admin`

You can display a node's parameters with the following command:

```
show [node:]nodename
```

For example, if `node1` is a unique name in the cluster database:

```
cxfs_admin:mycluster> show node1
node:node1:
  cellid=1
  enabled=true
  os=IRIX
  private_net:
    192.168.0.204
  status:
    client=stable
    connected=true
    fencing=Stable
    filesystems=up
    license:
      cpu_count=1
      have_license=true
      oem=none
      os=IRIX64
      version=4.0.0.2
    membership=up
    summary=Stable
    version=4.0.0.2
  wwns:
    210000e08b081f23
  xvm=up
  type=client_only
```

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

```
show node
```

For example:

```
cxfs_admin:mycluster> show node
node:
  node1:
    cellid=1
    enabled=true
    os=IRIX
    private_net:
      192.168.0.204
    status:
      client=stable
      connected=true
      fencing=Stable
      filesystems=up
      license:
        cpu_count=1
        have_license=true
        oem=none
        os=IRIX64
        version=4.0.0.2
      membership=up
      summary=Stable
      version=4.0.0.2
      wwns:
        210000e08b081f23
      xvm=up
      type=client_only
  node2:
    cellid=2
    enabled=true
    os=IRIX
    private_net:
      192.168.0.186
    status:
      connected=true
      fencing=Stable
```

```
        license:
            have_license=true
            summary=Stable
            version=4.0.0.1
            wwns:
                210000e08b0ead8c
            type=server_admin
    newnode:
        cellid=0
        enabled=true
        os=IRIX
        private_net:
            128.162.232.79
        status:
            connected=false
            summary=Inactive
        type=client_only
```

Cluster Tasks with `cxfs_admin`

This section discusses the following:

- "Create or Modify a Cluster with `cxfs_admin`" on page 282
- "Create a Tiebreaker with `cxfs_admin`" on page 284
- "Delete a Cluster with `cxfs_admin`" on page 285
- "Display a Cluster with `cxfs_admin`" on page 285
- "Show License Information with `cxfs_admin`" on page 286

Create or Modify a Cluster with `cxfs_admin`

To create the cluster, use the following command (line breaks shown here for readability, defaults in parentheses):

```
create cluster name=clustername
Advanced-mode:
    heartbeat_monitor=dynamic|static    (static)
    id=clusterID
```

For example:

```
cxfs_admin:> create cluster name=mycluster
```

You can use the `modify` command to add a tiebreaker node or change the heartbeat monitor type. (You cannot change the cluster's name or ID.)

```
modify clustername
  tiebreaker=client_only_nodename
Advanced-mode:
  heartbeat_monitor=dynamic|static
```

For example, if `mycluster` is a unique name in the cluster database, to make the client-only node `clientA` the CXFS tiebreaker:

```
cxfs_admin:mycluster> modify mycluster tiebreaker=clientA
```

Basic-mode usage notes:

- *clustername* is the logical name of the cluster. It cannot begin with a number or an underscore (`_`), or include any whitespace characters, and can be at most 255 characters.

Note: In basic mode, you are not prompted for a device name. Instead, `cxfs_admin` uses the value for `name` and prepends `/dev/cxvm/` to it.

- *tiebreaker* specifies the CXFS tiebreaker. See "Create a Tiebreaker with `cxfs_admin`" on page 284.

Advanced-mode usage notes:

- *heartbeat_monitor* specifies how cluster membership is monitored. All nodes send heartbeat messages once per second. If a node does not receive a heartbeat within a defined period, that node loses membership and is denied access to the cluster's filesystems. The defined period is one of the following:
 - `static`: Monitors constantly at 1-second intervals and declares a timeout after 5 consecutive missed seconds (default).
 - `dynamic`: Starts monitoring only when the node is processing a message from another node (such as for token recall or XVM multicast) or when the client monitors the server because it has a message pending (for example, a token

acquire or metadata operation). Once monitoring initiates, it monitors at 1-second intervals and declares a timeout after 5 consecutive missed seconds, just like static monitoring. Dynamic heartbeat monitoring is appropriate for clusters that have clients with heavy workloads; using it avoids inappropriate loss of membership. However, it may take longer to recover a client's tokens and other state information when there is an actual problem.

Note: You should not use the `Shutdown` failure policy on client nodes if you choose dynamic heartbeat monitoring for the cluster.

- `id` is a unique number within your network in the range 1 through 255. The cluster ID is used by the operating system 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. Clusters must have unique IDs.

Create a Tiebreaker with `cxfs_admin`

The CXFS tiebreaker node determines whether a CXFS kernel membership quorum is maintained when exactly half of the server-capable nodes can communicate with each other. There is no default CXFS tiebreaker.



Caution: SGI recommends that you use client-only nodes as tiebreakers to ensure that the cluster remains operational. `cxfs_admin` will only let you specify a server-capable node as a tiebreaker if the cluster contains four or more server-capable nodes, and an even number of server-capable nodes.

The reset capability or I/O fencing with switches is **mandatory** to ensure data integrity for all nodes. Clusters should have an odd number of server-capable nodes. If you have an even number of server-capable administration nodes, define a CXFS tiebreaker node. (See "CXFS Recovery Issues in a Cluster with Only Two Server-Capable Nodes " on page 526.)

To set the CXFS tiebreaker node, use the `modify` command as follows:

```
modify [cluster:]clustername tiebreaker=client_nodename
```

For example:

```
cxfs_admin:mycluster> modify mycluster tiebreaker=myclient
```

To unset the CXFS tiebreaker node, do not supply a value for tiebreaker. For example:

```
cxfs_admin:mycluster> modify mycluster tiebreaker=
```

Delete a Cluster with `cxfs_admin`

To delete a cluster, use the following command:

```
delete [cluster:]clustername
```

For example, if `mycluster` is a unique name in the cluster database:

```
cxfs_admin:mycluster> delete mycluster
```

However, you cannot delete an active cluster; you must first unmount and delete the filesystems, disable and delete the nodes, and so on.

Display a Cluster with `cxfs_admin`

To display the cluster, use the following command:

```
show cluster
```

For example:

```
cxfs_admin:mycluster> show cluster
cxfs:cluster:
  mycluster:
    access:
      admin=server
      monitor=
    failover_net:
      10.11.0.0, 128.162.242.0
    filesystem:
      concatfs, mirrorfs, stripefs
    node:
      aix-client, irix-client, mac-client, linux-client, mds1, mds2, solaris-client, windows-client
    status:
```

```
filesystems:
  summary=concatfs: mac-client trying to mount
  mirrorfs: mac-client trying to mount
  stripefs: mac-client trying to mount
licenses:
  cxfs_client
nodes:
  summary=aix-client: Inactive
  mac-client: Mounted 0 of 3 filesystems
  solaris-client: Inactive
  windows-client: Inactive
  summary=node(s) not stable, filesystem(s) not stable
switch:
  fcswitch12, fcswitch13
tiebreaker=irix-client
```

This output says that cluster `mycluster` has three nodes, one of which is inactive. All sever-capable administration nodes within the cluster plus the external node `remoteA` can use `cxfs_admin` to make changes to the cluster database. There is one filesystem, which is in stable condition (which means that all nodes that should have the filesystem mounted do in fact have the filesystem mounted; a nodes may correctly not have the filesystem mounted because it is disabled, not in membership, have it explicitly unmounted, or explicitly not allowed to mount it).

Show License Information with `cxfs_admin`

To show the CXFS licenses available for the cluster, use the following command:

```
show licenses
```

For example:

```
cxfs_admin:mycluster> show licenses
status:licenses:
  cxfs_client:
    enterprise:
      allocated=17
      valid=22
    workstation:
      allocated=4
      valid=15
```

CXFS Filesystem Tasks with `cxfs_admin`

The `filesystem` class represents the clustered XVM volumes that can be mounted by CXFS nodes. Before you can create a filesystem definition, you must create the clustered XVM volume and make the filesystem with `mkfs`.

By default, the filesystem:

- Uses the XVM device of the same name
- Enables all nodes to mount the filesystem
- Mounts the filesystem in `/mnt/`
- Is not managed by GRIOv2

To override these defaults, use the optional attributes listed below.

This section discusses the following:

- "Create or Modify a CXFS Filesystem with `cxfs_admin`" on page 287
- "Mount a CXFS Filesystem with `cxfs_admin`" on page 292
- "Unmount a CXFS Filesystem with `cxfs_admin`" on page 293
- "Relocate the Metadata Server for a Filesystem with `cxfs_admin`" on page 293
- "Delete a CXFS Filesystem with `cxfs_admin`" on page 294
- "Show a CXFS Filesystem" on page 294

Create or Modify a CXFS Filesystem with `cxfs_admin`

Use the following commands to define a filesystem and the nodes on which it may be mounted (line breaks shown here for readability, defaults in parentheses):

```
create filesystem name=filesystemname
  [options=mount_options]
  [forced_unmount=true|false]                (false)
  [mountpoint=mountpoint]                  (/mnt/filesystemname)
  [mounted=true|false]                       (true)
Advanced-mode:
  [device=devicename]                       (filesystemname)
  [servers=server_list]                     (all servers are potential MDS)
```

```
[nodes=nodes_that_can_mount]           (all nodes can mount)  
[mount_new_nodes=true|false]           (true)  
[grio_managed=true|false]              (false)  
[grio_qual_bandwidth=qualified_bandwidth]
```

Note: Relocation is disabled by default. Recovery and relocation are supported only when using standby nodes. Therefore, you should only define multiple metadata servers for a given filesystem if you are using the standby node model. See "Relocation" on page 20.

Basic-mode usage notes:

- name specifies the name of the filesystem:

Note: You must create the CXFS filesystem with `xvm` before you set it up using `cxfs_admin`.

- If you also specify a value for `device`, then `name` can be any string that does not begin with a number or an underscore (`_`), or include any whitespace characters, and can be at most 255 characters. For example, if the full XVM volume name is `/dev/cxvm/concat1`:

```
cxfs_admin:mycluster> create filesystem name=filesys1 device=concat1
```

- If you do not specify a value for `device`, then `name` must be the name of the XVM volume following `/dev/cxvm`. For example:

```
cxfs_admin:mycluster> create filesystem name=concat1
```

Note: Within the GUI, the default is to use the last portion of the device name; for example, for a device name of `/dev/cxvm/d761un0s0`, the GUI will automatically supply a logical filesystem name of `d761un0s0`. The GUI will accept other logical names defined with `cxfs_admin` but the GUI will not allow you to modify a logical name; you must use `cxfs_admin` to modify the logical name.

- options specifies the mount options that are passed to the `mount` operating system command. These mount options control access to the specified filesystem. For a list of supported mount options, see the *CXFS MultiOS Client-Only Guide for SGI InfiniteStorage*. By default, this is unassigned.

Specify multiple mount options as a comma-separated list. For example, the following specifies that the `myfs` filesystem uses `inode64` allocation and does not update the access time stamps for files and directories:

```
cxfs_admin:mycluster> create filesystem name=myfs options=inode64,noatime
```

Note: No validation is done on the mount options in `cxfs_admin`, so an invalid option may prevent the filesystem mounting on all nodes.

- `forced_unmount` controls the action that CXFS takes if there are processes that have open files or directories in the filesystem to be unmounted:
 - If set to `true`, the processes will be killed and the unmount will occur
 - If set to `false`, the processes will not be killed and the filesystem will unmount only after all references to the filesystem have been closed (default)
- `mounted` specifies whether a new filesystem is mounted on all nodes in the cluster (`true`) or not mounted on any nodes (`false`). By default, the new filesystem is mounted on all nodes (`true`).
- `mountpoint` specifies a mount point for the filesystem. The mount point is a directory to which the XVM volume is attached. This directory name must begin with a slash (`/`). The default is `/mnt/filesystemname`.

For example, to create a filesystem named `myfs` and use the default mount point of `/mnt/myfs`:

```
cxfs_admin:mycluster> create filesystem name=myfs
```

To create the `myfs` filesystem but use a mount point of `/tmp/myfs`:

```
cxfs_admin:mycluster> create filesystem name=myfs mountpoint=/tmp/myfs
```

Advanced-mode usage notes:

- `device` is the device name for an XVM volume. The default is the filesystem name specified by name.

Note: Specify only the XVM volume name itself. Do not include `/dev/cxvm/`.

For example, to create a device name of `mydev` for the `myfs` filesystem:

```
cxfs_admin:mycluster> create filesystem name=myfs device=mydev
```

- `servers` specifies the potential metadata servers that can serve the filesystem to the cluster. To specify multiple server capable nodes, use a comma-separated list of node names. The default is all server-capable nodes in the cluster.

For example, to specify that either `node2` or `node3` could be the metadata server, with `node2` being the primary server, for the `myfs` filesystem:

```
cxfs_admin:mycluster> create filesystem name=myfs servers=node2,node3
```

- `nodes` specifies the only nodes that can mount the filesystem as a specified comma-separated list. If you do not specify `nodes` on the `create` command, all nodes can mount the filesystem. If you restrict the nodes on the `create` command line, you can later mount all nodes by specifying all of them with the `nodes` attribute.

For example, to restrict mounting the `myfs` filesystem to nodes `node1` and `node2`:

```
create myfs nodes=node1,node2
```

To add `node3`:

```
modify myfs nodes=node1,node2,node3
```

- `mount_new_nodes` specifies whether a newly created node will automatically mount the filesystem when it gets membership (`true`) or will not mount the filesystem (`false`). By default, new nodes mount all defined filesystems.

For example, to create filesystem `myfs` that is not automatically mounted by new nodes, use the following command:

```
cxfs_admin:mycluster> create filesystem name=myfs mount_new_nodes=false
```

To later mount the filesystem on `node3` after it has been created, use the following command:

```
cxfs_admin:mycluster> mount myfs nodes=node3
```

- `grio_managed` specifies whether a filesystem is managed by GRIOv2 (`true`) or not (`false`). The default is `false`. Setting `grio_managed` to `false` disables GRIO management for the specified filesystem, but it does not reset the `grio_qual_bandwidth` value. In this case, `grio_qual_bandwidth` is left unmodified in the cluster database and ignored.

- `grio_qual_bandwidth` specifies a filesystem's qualified bandwidth in bytes (B suffix), kilobytes (KB), megabytes (MB), or gigabytes (GB), where the units are multiples of 1024. The default is MB for 4000 or less, B for 4001 or greater. If the filesystem is GRIO-managed, you must specify a qualified bandwidth with this attribute. You can modify the qualified bandwidth for a mounted filesystem without taking it offline.

For example, the following commands all create the `myfs` filesystem with a GRIOv2 qualified bandwidth of 1.2 GB/s:

```
cxfs_admin:mycluster> create filesystem name=myfs grio_qual_bandwidth=1288500000
cxfs_admin:mycluster> create filesystem name=myfs grio_qual_bandwidth=1258300KB
cxfs_admin:mycluster> create filesystem name=myfs grio_qual_bandwidth=1288.8MB
cxfs_admin:mycluster> create filesystem name=myfs grio_qual_bandwidth=1.2GB
```

For example, using prompting in basic mode:

```
cxfs_admin:mycluster> create filesystem
name? myfs
options? rw
forced_unmount? false
mountpoint? /mnt/myfs
mounted? true
Filesystem "myfs" has been created, waiting for it to be mounted on all
assigned nodes...
Waiting for filesystem myfs, current status: A server is trying to mount
Waiting for filesystem myfs, current status: node1 trying to mount, node2
trying to mount
Waiting for filesystem myfs, current status: node1 trying to mount
Operation completed successfully
```

For example, using prompting in advanced mode:

```
cxfs_admin:mycluster> create filesystem
Specify the attributes for create filesystem:
name? myfs
options? rw
forced_unmount? false
mountpoint? /mnt/myfs
device? myfs
servers? node1,node2,node3
nodes? node1,node2,node3,node4
mounted? true
```

```
mount_new_nodes? true
grio_managed? false
Filesystem "myfs" has been created, waiting for it to be mounted on all
assigned nodes...
Waiting for filesystem myfs, current status: A server is trying to mount
Waiting for filesystem myfs, current status: node1 trying to mount, node2
trying to mount,node3 trying to mount, node4 trying to mount
Waiting for filesystem myfs, current status: node1 trying to mount
Operation completed successfully
```

Note: After a filesystem has been defined in CXFS, running `mkfs` on it will cause errors to appear in the system log file. To avoid these errors, run `mkfs` before defining the filesystem in CXFS, or delete the CXFS filesystem before running `mkfs`. See "Delete a CXFS Filesystem with `cxfs_admin`" on page 294.

Mount a CXFS Filesystem with `cxfs_admin`

The `mount` command operates on the set of nodes that were specified in the `nodes=odelist` attribute when the filesystem was created. By default, this is all nodes in the cluster.

To mount the filesystem on all enabled nodes in the cluster:

```
mount filesystem
```

To mount the filesystem on specific enabled nodes:

```
mount filesystem nodes=odelist
```

For example, to mount the filesystem `myfs` on only nodes `node2` and `node3`:

```
cxfs_admin:mycluster> mount myfs nodes=node2,node3
```

Note: If any nodes that are set to mount the filesystem are enabled and attached are not in membership, mounting a filesystem will not completely succeed. The filesystem will be marked as ready to be mounted when the nodes achieve membership.

Unmount a CXFS Filesystem with `cxfs_admin`

To unmount a filesystem from all nodes in the cluster:

```
umount filesystem
```

To unmount the filesystem from a specific comma-separated list of nodes:

```
umount filesystem nodes=nodelist
```

For example, to unmount filesystem `myfs` from nodes `node1` and `node3`:

```
cxfs_admin:mycluster> umount myfs nodes=node1,node3
```

Note: If any nodes are not in membership, the filesystem will be marked as not to be mounted when the nodes achieve membership.

Relocate the Metadata Server for a Filesystem with `cxfs_admin`

The `relocate` command forcefully moves a filesystem's metadata server to another node in the cluster that has already been defined as a potential metadata server for that filesystem. This action is typically used to free a server so it can be brought down for maintenance or upgrades. Relocation must also be explicitly enabled in the kernel with the `cxfs_relocation_ok` system tunable parameter (see "Relocation" on page 20).

If relocation is explicitly enabled in the kernel, you can relocate a metadata server to another node by using the following command:

```
relocate filesystem server=new_metadata_server
```

For example:

```
cxfs_admin:mycluster> relocate myfs server=node2
```

CXFS kernel membership is not affected by relocation. However, users may experience a degradation in filesystem performance while the metadata server is relocating.

Delete a CXFS Filesystem with `cxfs_admin`

Use the following command to delete a filesystem:

```
delete [filesystem:]filesystem
```

You cannot delete a mounted filesystem; you must first unmount it. For example, if `myfs` is a unique name in the cluster database:

```
cxfs_admin:mycluster> unmount myfs  
cxfs_admin:mycluster> delete myfs
```

Show a CXFS Filesystem

To show information about all filesystems:

```
show filesystem
```

To show information about a specific filesystem:

```
show [filesystem:]filesystemname
```

For example:

```
cxfs_admin:mycluster> show stripefs  
filesystem:stripefs:  
  forced_unmount=false  
  mount=true  
  mountpoint=/mnt/stripefs  
  nodes:  
    aix-client, irix-client, mac-client, linux-client, mds1, mds2, solaris-client, windows-client  
  options=  
  servers:  
    mds1, mds2  
  status:  
    free=1.31TB  
    nodes:  
      aix-client=inactive  
      irix-client=mounted  
      linux-client=mounted  
      mac-client=trying to mount  
      mds1=mounted  
      mds2=mounted
```

```
solaris-client=inactive
windows-client=mounted
server=mds2
size=1.82TB
summary=mac-client trying to mount
utilization=28%
```

Network Failover Tasks with `cxfs_admin`

To allow the cluster to continue operation if the primary private network fails, you can set up private network failover.

To inform the servers of the failover networks, you must create a `failover_net` network.

Each node in the cluster must have all `private_net` values specified to match the subsets defined by the failover networks in the same order as all other nodes in the cluster.

Command syntax:

```
create failover_net network=IPaddress
mask=IPmask
```

To create two private networks, one on the 192.168.0.x and the other on the 10.0.0.x subnets, use the following command:

```
cxfs_admin:mycluster> create failover_net network=192.168.0.0 mask=255.255.255.0
cxfs_admin:mycluster > create failover_net network=10.0.0.0 mask=255.255.255.0
```

To create a node with failover network support:

```
cxfs_admin:mycluster> create node name=mynode private_net=192.168.0.2,10.0.0.2
```

Switch Tasks with `cxfs_admin`

This section discusses the following:

- "Create a Switch with `cxfs_admin`" on page 296
- "Delete a Switch Definition with `cxfs_admin`" on page 297

- "Show Switches with `cxfs_admin`" on page 297

For general information, see "Fence" on page 30.

Note: Nodes without system controllers require I/O fencing to protect data integrity. A switch is mandatory to support I/O fencing; therefore, multiOS CXFS clusters require a switch. See the release notes for supported switches.

To raise or lower a fence, or update switch port information, use the `hafence` command.

Create a Switch with `cxfs_admin`

To define a new switch, use the following command:

```
create switch name=switch_hostname
  [password=username_password]           (password)
  [user=username]                         (admin)
  [vendor=brocade | qlogic | site-specific_vendor] (brocade)
Advanced-mode:
  [mask=ports_that_will_not_be_fenced]
```

Note: You must define all of the switches within your fabric to which a CXFS client or server is connected.

Basic-mode usage notes:

- `name` specifies the hostname of the Fibre Channel switch; this is used to determine the IP address of the switch.
- `password` specifies the password for the specified `username`. The default is `password`.
- `user` specifies the user name to use when sending a telnet message to the switch. The default is `admin`.
- `vendor` specifies the vendor of the Fibre Channel switch. It can be one of the following values:

 `brocade` (default)
 `qlogic`

site-specific-value

For example, if `myswitch` is a QLogic switch:

```
cxfs_admin:mycluster> create switch name=myswitch vendor=qlogic
```

Advanced-mode usage notes:

- `mask` specifies the ports on the switch that will never be fenced. By default, no ports are masked (and therefore all ports are available for fencing). The value for `mask` is a series of comma-separated port ranges. For example, the following states that ports 0, 4 and 12 to 15 for `myswitch` will never be fenced by CXFS:

```
cxfs_admin:mycluster> create switch name=myswitch mask=0,4,12-15
```

CXFS administration nodes automatically discover the available HBAs and, when fencing is triggered, fence off all of the Fibre Channel HBAs when the `Fence` or `FenceReset` fail action is selected. However, masked HBAs will not be fenced. Masking allows you to prevent the fencing of devices that are attached to the SAN but are not shared with the cluster, to ensure that they remain available regardless of CXFS status. You would want to mask HBAs used for access to tape storage, or HBAs that are only ever used to access local (nonclustered) devices.

Delete a Switch Definition with `cxfs_admin`

To delete a switch, use the following command:

```
delete [switch:]switch_hostname
```

For example, if `myswitch` is a unique name in the cluster database:

```
cxfs_admin:mycluster> delete myswitch
```

Show Switches with `cxfs_admin`

To display all of the switches in the system, use the following command:

```
show switch [output=full_pathname]
```

For example, in basic mode:

```
switch:  
  fcswitch12:
```

```
hostname=fcswitch12
num_ports=32
port:
    0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18,
    19, 20, 21, 22, 23, 24, 25, 26, 27, 28, 29, 30, 31
vendor=brocade
fcswitch13:
hostname=fcswitch13
num_ports=32
port:
    0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18,
    19, 20, 21, 22, 23, 24, 25, 26, 27, 28, 29, 30, 31
vendor=brocade
```

To send the output to the /tmp/switchinfo file:

```
cxfs_admin:mycluster> show switch output=/tmp/switchinfo
```

To display a specific switch:

```
show [switch:]switchname [output=full_pathname]
```

To display mask values, use advanced mode (see "Basic and Advanced Mode" on page 260.) For example, if myswitch is a unique name in the cluster database:

```
cxfs_admin:mycluster> show myswitch
switch:myswitch:
    hostname=myswitch
    mask=
    num_ports=16
    port:
        0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15
    switchid=Switch0
    vendor=brocade
```

To display the switches and ports each host is connected to:

```
show wwns
```

For example:

```
cxfs_admin:mycluster> show wwns
node:mdsl:status:wwns:
```

```
100000062b0f5284:
    switch=fcswitch12
    switch_port=11
100000062b0f5285:
    switch=fcswitch13
    switch_port=11
node:mgs2:status:wwns:
100000062b0f568c:
    switch=fcswitch13
    switch_port=20
100000062b0f568d:
    switch=fcswitch12
    switch_port=20
node:mac-client:status:wwns:
100000062b105528:
100000062b105529:
node:windows-client:status:wwns:
210000e08b12767a:
    switch=fcswitch13
    switch_port=3
210100e08b32767a:
    switch=fcswitch12
    switch_port=3
node:linux-client:status:wwns:
210000e08b1cad7e:
    switch=fcswitch12
    switch_port=25
210100e08b3cad7e:
    switch=fcswitch13
    switch_port=25
node:irix-client:status:wwns:
210000e08b087a3f:
    switch=fcswitch12
    switch_port=0
210100e08b287a3f:
    switch=fcswitch13
    switch_port=0
```

To show full status details for each port on the switch, use one of the following commands:

```
show [switch:]switchname all
show switchname:port
```

For example, for the switch named `fcswitch12`:

```
cxfs_admin:mycluster> show fcswitch12:port
```

Saving and Recreating the Current Configuration with `cxfs_admin`

The `config` command displays a series of commands that represent the current configuration of the objects specified. You can use this output to recreate the configuration of the entire cluster or a subset of it.

By default, `config` displays information at the `cxfs_admin` prompt. To write the configuration output to a file, use the `output` attribute and specify the full pathname of the file to contain the information:

```
config node output=full_pathname
```

You can use the generated file with the `-f` command line option to recreate the configuration at a later time.

Note: For a more readable configuration output (without the related commands), use the `show` command rather than the `config` command.

For example, to display all node configuration commands:

```
config node
```

For example (blank lines and line breaks added here for readability):

```
cxfs_admin:mycluster> config node
```

```
create node name=node1 os=Linux type=server_admin private_net=192.168.0.168
enabled=true hostname=node1.example.com failpolicy=Fence,Reset,Shutdown
nodeid=1 reset_method=powerCycle reset_port=bmc reset_status=disabled reset_node=node2
reset_comms=ipmi reset_device=node1-bmc.mycompany.com
```

```
create node name=node2 os=Linux type=server_admin private_net=192.168.0.185
enabled=true hostname=node2.example.com failpolicy=Fence,Shutdown nodeid=2
```

```
create node name=node3 os=IRIX type=client_only private_net=192.168.0.204
enabled=true hostname=node3.example.com failpolicy=Fence,Shutdown nodeid=3
```

```
create node name=node4 os=Linux type=server_admin private_net=128.162.232.79
enabled=true hostname=node4.example.com failpolicy=Fence,Shutdown nodeid=4
```

To display the configuration commands for a specific node:

```
config [node:]nodename
```

For example, if the name node3 is unique in the database:

```
cxfs_admin:mycluster> config node3
create node name=node3 os=IRIX type=client_only private_net=192.168.0.204
enabled=true hostname=node3.example.com failpolicy=Fence,Shutdown nodeid=3
```

To dump the entire cluster configuration to the /tmp/config.txt file (where * denotes all objects):

```
cxfs_admin:mycluster> config * output=/tmp/config.txt
```

Note: You must give the absolute pathname for the output file.

Following is an example of a file /tmp/buildcluster that creates a cluster named mycluster with two server-capable nodes (mds1 and mds2) and a client of each OS type:



Caution: Line breaks and indentations added in this guide for readability. Each cxfs_admin command must actually be on one line.

```
create cluster name=mycluster id=1 heartbeat_monitor=static
create node name=mds1 os=Linux type=server_admin
  private_net=10.11.0.239,128.162.242.4 enabled=true hostname=mds1
  failpolicy=Reset nodeid=2 reset_method=powerCycle reset_port=bmc
  reset_status=enabled reset_node=mds2 reset_comms=ipmi
  reset_device=bmc-mds1
create node name=mds2 os=Linux type=server_admin
```

```
private_net=10.11.0.242,128.162.242.6 enabled=true hostname=mds2
failpolicy=Reset nodeid=1 reset_method=powerCycle reset_port=bmc
reset_status=enabled reset_node=mds1 reset_comms=ipmi
reset_device=bmc-mds2
create node name=mac-client os=MacOSX type=client_only
private_net=10.11.0.150,128.162.242.193 enabled=true hostname=mac-client
failpolicy=Fence,Shutdown nodeid=9
create node name=windows-client os=Windows type=client_only
private_net=10.11.0.166,128.162.242.241 enabled=true hostname=windows-client
failpolicy=Fence,Shutdown nodeid=6
create node name=aix-client os=AIX type=client_only
private_net=10.11.0.48,128.162.242.197 enabled=true hostname=aix-client
failpolicy=Fence,Shutdown nodeid=12
create node name=solaris-client os=Solaris type=client_only
private_net=10.11.0.60,128.162.242.196 enabled=true hostname=solaris-client
failpolicy=Fence,Shutdown nodeid=13
create node name=linux-client os=Linux type=client_only
private_net=10.11.0.253,128.162.242.226 enabled=true hostname=linux-client
failpolicy=Fence,Shutdown nodeid=4
create node name=irix-client os=IRIX type=client_only
private_net=10.11.0.52,128.162.242.200 enabled=true hostname=irix-client
failpolicy=Fence,Shutdown nodeid=8
modify cluster:mycluster tiebreaker=linux-client
create failover_net network=10.11.0.0 mask=255.255.255.0
create failover_net network=128.162.242.0 mask=255.255.255.0
create switch name=fcs12 vendor=brocade
create filesystem name=concatfs device=concatfs mountpoint=/mnt/concatfs
options=rw,dmi servers=mds1,mds2
nodes=mac-client, windows-client,mds1,mds2,aix-client,solaris-client,linux-client,irix-client
forced_unmount=false mounted=true mount_new_nodes=true
create filesystem name=mirrorfs device=mirrorfs mountpoint=/mnt/mirrorfs
options=rw,dmi servers=mds1,mds2
nodes=mac-client, windows-client,mds1,mds2,aix-client,solaris-client,linux-client,irix-client
forced_unmount=false mounted=true mount_new_nodes=true
create filesystem name=stripefs device=stripefs mountpoint=/mnt/stripesfs
options=rw,dmi servers=mds1,mds2
nodes=mac-client, windows-client,mds1,mds2,aix-client,solaris-client,linux-client,irix-client
forced_unmount=false mounted=true mount_new_nodes=true
```

To use this file to recreate the cluster, first clear the existing cluster configuration and then send the file to `cxfs_admin`:

- Clear the cluster database. See "Clearing the Cluster Database" on page 485.
- Recreate the cluster:

```
mds1:/usr/cluster/bin # /usr/cluster/bin/cxfs_admin -s -f /tmp/buildcluster
Connecting to the local CXFS server...
Node "mds1" has been created, waiting for it to join the cluster...
Please restart all cxfs and cluster services on the server "mds1" to make it join the cluster.
Node "mds2" has been created, waiting for it to join the cluster...
Please restart all cxfs and cluster services on the server "mds2" to make it join the cluster.
Node "mac-client" has been created, waiting for it to join the cluster...
Node "windows-client" has been created, waiting for it to join the cluster...
Node "aix-client" has been created, waiting for it to join the cluster...
Node "solaris-client" has been created, waiting for it to join the cluster...
Node "linux-client" has been created, waiting for it to join the cluster...
Node "irix-client" has been created, waiting for it to join the cluster...
Filesystem "concatfs" has been created, waiting for it to be mounted on all assigned nodes...
Filesystem "mirrorfs" has been created, waiting for it to be mounted on all assigned nodes...
Filesystem "stripefs" has been created, waiting for it to be mounted on all assigned nodes...
Waiting for configuration update to complete
Operation completed successfully
```


Administration and Maintenance

When a CXFS filesystem is mounted, you must perform cluster administration from an active metadata server if you are using the GUI. (You cannot connect the GUI to a client-only node.) When filesystems are not mounted, you can perform administration tasks using the `cxfs_admin` command when logged into any CXFS administration node (one that is installed with the `cxfs_cluster` product), or when the CXFS GUI is connected to any CXFS administration node. You can also use `cxfs_admin` from a server-capable node that has permission to access the CXFS cluster database.

Note: You should perform reconfiguration and/or cluster manipulation (such as adding or deleting filesystems or nodes) on a scheduled cluster maintenance shift and not during production hours. You should stop CXFS services on an administration node before performing maintenance on a node.

The following are the same in CXFS and XFS:

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

The rest of this chapter discusses the following topics:

- "Precedence of Configuration Options" on page 307
- "CXFS Release Versions and Rolling Upgrades" on page 307
- "Upgrading From 3.4.1 Or Earlier" on page 314
- "Manual CXFS Start/Stop" on page 315
- "CXFS and Cluster Administration Initialization Commands" on page 316
- "Using `ha fence` to Manipulate a Switch" on page 318
- "CXFS Port Usage" on page 320
- "CXFS `chkconfig` Arguments" on page 320

- "Configuring Real-Time Filesystems For IRIX Nodes" on page 323
- "Granting Task Execution Privileges to Users" on page 323
- "Transforming an Existing Node into a Client-Only Node" on page 324
- "CXFS Mount Scripts" on page 325
- "Unmounting `lofs` File Systems" on page 327
- "Using `telnet` and I/O Fencing" on page 327
- "Using `fsr` and `xfss_fsr`" on page 328
- "Using `cron` in a CXFS Cluster" on page 328
- "Using Hierarchical Storage Management (HSM) Products" on page 328
- "Discovering the Active Metadata Server for a Filesystem" on page 329
- "Metadata Server Recovery" on page 332
- "Shutdown of the Database and CXFS" on page 332
- "Avoiding a CXFS Restart at Reboot" on page 338
- "Log File Management" on page 339
- "Volume Management" on page 341
- "Disk Management" on page 342
- "Filesystem Maintenance" on page 344
- "Dump and Restore" on page 346
- "Site-Changeable System Tunable Parameters" on page 348
- "Restricted System Tunable Parameters" on page 356
- "Hardware Changes and I/O Fencing" on page 367
- "Configuring Private Network Failover" on page 368
- "Removing and Restoring Cluster Members" on page 369
- "Discovering the WWNs" on page 375

- "Mapping XVM Volumes to Storage Targets" on page 375
- "Switching Between SGIRDAC and SGIAVT Mode for SGI RAID" on page 376

See also Chapter 14, "Cluster Database Management" on page 379.

Note: If you have upgraded directly from IRIX 6.5.12f or earlier, you must manually convert you filesystem definitions to the new format. See "IRIX: Converting Filesystem Definitions for Upgrades" on page 135.

Precedence of Configuration Options

CXFS programs take their configuration options as follows:

- Built-in defaults, which can be overridden by
- Configuration files, which can be overridden by
- Environment variables, which can be overridden by
- Command-line options

CXFS Release Versions and Rolling Upgrades

To identify compatible CXFS releases, see the *CXFS MultiOS Software Compatibility Matrix* that is posted on Supportfolio.

Beginning with CXFS 3.2, SGI lets you upgrade of a subset of nodes from *X.anything* to *X.anything* within the same major-release thread (*X*). This policy lets you to keep your cluster running and filesystems available during the temporary upgrade process.



Caution: You must upgrade all server-capable nodes before upgrading any client-only nodes (servers must run the same or later release as client-only nodes.) Operating a cluster with clients running a mixture of older and newer CXFS versions will result in a performance loss.

After the upgrade process is complete, all nodes should be running the same major-level release *X.Y* (such as 4.0), or any minor-level release with the same major level *X.Y.anything* (such as 4.0.3).



Caution: Although clients that are not upgraded might continue to function in the CXFS cluster without problems, new CXFS functionality may not be enabled until all clients are upgraded and SGI does not provide support for any CXFS problems encountered on the clients that are not upgraded.

SGI recommends the following for server-capable nodes in a production cluster:

- Run the latest CXFS release.
- Run a release that is the same or later than the release run by client-only nodes. (The only exception is if the release in question does not apply to the server-capable nodes.)
- Run the same minor-level release (such as 4.0.3) on all server-capable nodes.

Using the 3.4.2 release as an example, a production cluster could contain server-capable nodes running 3.4.2 and client-only nodes running 3.4, 3.4.1, and 3.4.2; it could contain client-only nodes running 3.4.3 only because there was no server platforms included in 3.4.3. It should not contain any nodes running 3.3.

General Upgrade Procedure

Before you start an upgrade, you should save the CXFS configuration as a precaution before making changes and acquire new CXFS server-side licenses (if required). See Chapter 14, "Cluster Database Management" on page 379, and Chapter 4, "CXFS License Keys" on page 75.

To upgrade a CXFS cluster, do the following:

1. Ensure all server-capable nodes are running the same software release.
2. Upgrade the *standby node*, which is a server-capable administration node that is configured as a potential metadata server for a given filesystem, but does not currently run any applications that will use that filesystem.
3. For the next server-capable node, move all CXFS filesystems running on it to the standby node (making the standby node now the active metadata server for those filesystems). Run `chkconfig cluster off` and then reset the node to force recovery.

Note: When performing upgrades, you should not make any other configuration changes to the cluster (such as adding new nodes or filesystems) until the upgrade of all nodes is complete and the cluster is running normally.

4. Upgrade the server-capable node.
 5. Return the upgraded server-capable node to the cluster. Run `chkconfig cluster on` and reset the node.
-

Note: Skip steps 6, 7, and 8 if your cluster has only two server-capable nodes.

6. For the next server-capable node, move all CXFS filesystems running on it to the standby node (making the standby node now the active metadata server for those filesystems). Run `chkconfig cluster off` and then reset the node to force recovery.
7. Upgrade the server-capable node.
8. Return the upgraded server-capable node to the cluster. Run `chkconfig cluster on` and reset the node.

If your cluster has additional server-capable nodes, repeat steps 6 through 8 for each remaining server-capable node.

9. Return the first CXFS filesystem to the server-capable node that you want to be its metadata server (make it the active metadata server).
10. Return the next CXFS filesystem to the server-capable node that you want to be its metadata server (make it the active metadata server).

Repeat this step as needed for each CXFS filesystem.

11. Upgrade the client-only nodes.
-

Note: There are issues when upgrading from releases prior to 3.4.2 if you have more than two server-capable nodes. See "Upgrading From 3.4.1 Or Earlier" on page 314.

Example Upgrade Process

The following figures show an example upgrade procedure for a three-node cluster with two filesystems (fs1 and fs2), in which all nodes are running CXFS 4.0 at the beginning and NodeB is the standby node.

1 Starting configuration, all nodes running 4.0:

NodeA Server-Capable 4.0	NodeB Server-Capable 4.0	NodeC Server-Capable 4.0	NodeD Client-only 4.0
fs1 (MDS)	fs1 (P)	fs1 (C)	fs1 (C)
fs2 (P)	fs2 (C)	fs2 (MDS)	fs2 (C)

2 Upgrade NodeB to 4.1:

NodeA Server-Capable 4.0	NodeB Server-Capable 4.1	NodeC Server-Capable 4.0	NodeD Client-only 4.0
fs1 (MDS)	fs1 (P)	fs1 (C)	fs1 (C)
fs2 (P)	fs2 (C)	fs2 (MDS)	fs2 (C)

3 On NodeA, run `chkconfig cluster off` and then reset NodeA to force recovery of fs1 onto NodeB:

NodeA Server-Capable 4.0	NodeB Server-Capable 4.1	NodeC Server-Capable 4.0	NodeD Client-only 4.0
	fs1 (MDS)	fs1 (C)	fs1 (C)
	fs2 (C)	fs2 (MDS)	fs2 (C)

4 Upgrade NodeA to 4.1:

NodeA Server-Capable 4.1	NodeB Server-Capable 4.1	NodeC Server-Capable 4.0	NodeD Client-only 4.0
	fs1 (MDS)	fs1 (C)	fs1 (C)
	fs2 (C)	fs2 (MDS)	fs2 (C)

Key:
MDS = active metadata server P = potential metadata server C = client

Figure 13-1 Example Rolling Upgrade Procedure (part 1)

5 On NodeA, run `chkconfig cluster on` and then reset NodeA:

Note: Ensure that there will be no I/O that will be restarted from NodeA to fs1 or fs2 after NodeA is reset.

NodeA Server-Capable 4.1	NodeB Server-Capable 4.1	NodeC Server-Capable 4.0	NodeD Client-only 4.0
fs1 (P)	fs1 (MDS)	fs1 (C)	fs1 (C)
fs2 (P)	fs2 (C)	fs2 (MDS)	fs2 (C)

6 On NodeC, run `chkconfig cluster off` and then reset NodeC to force recovery of fs2 onto NodeA:

NodeA Server-Capable 4.1	NodeB Server-Capable 4.1	NodeC Server-Capable 4.0	NodeD Client-only 4.0
fs1 (P)	fs1 (MDS)		fs1 (C)
fs2 (MDS)	fs2 (C)		fs2 (C)

7 Upgrade NodeC to 4.1:

NodeA Server-Capable 4.1	NodeB Server-Capable 4.1	NodeC Server-Capable 4.1	NodeD Client-only 4.0
fs1 (P)	fs1 (MDS)		fs1 (C)
fs2 (MDS)	fs2 (C)		fs2 (C)

8 On NodeC, run `chkconfig cluster on` and then reset NodeC:

Note: Ensure that there will be no I/O that will be restarted from NodeC to fs2 after NodeC is reset.

NodeA Server-Capable 4.1	NodeB Server-Capable 4.1	NodeC Server-Capable 4.1	NodeD Client-only 4.0
fs1 (P)	fs1 (MDS)	fs1 (C)	fs1 (C)
fs2 (MDS)	fs2 (C)	fs2 (P)	fs2 (C)

Key:
MDS = active metadata server P = potential metadata server C = client

Figure 13-2 Example Rolling Upgrade Procedure (part 2)

- 9 To return the active metadata server for fs2 to NodeC, reset NodeA:

Note: Ensure that there will be no I/O that will be restarted from NodeA to fs2 after NodeA is reset.

NodeA Server-Capable 4.1	NodeB Server-Capable 4.1	NodeC Server-Capable 4.0	NodeD Client-only 4.0
fs1 (P) fs2 (P)	fs1 (MDS) fs2 (C)	fs1 (C) fs2 (MDS)	fs1 (C) fs2 (C)

- 10 To return the active metadata server for fs1 to NodeA, reset NodeB:

NodeA Server-Capable 4.1	NodeB Server-Capable 4.1	NodeC Server-Capable 4.1	NodeD Client-only 4.0
fs1 (MDS) fs2 (P)	fs1 (P) fs2 (C)	fs1 (C) fs2 (MDS)	fs1 (C) fs2 (C)

- 11 Upgrade the client-only NodeD to 4.1 (repeat for all other client-only nodes):

NodeA Server-Capable 4.1	NodeB Server-Capable 4.1	NodeC Server-Capable 4.1	NodeD Client-only 4.1
fs1 (MDS) fs2 (P)	fs1 (P) fs2 (C)	fs1 (C) fs2 (MDS)	fs1 (C) fs2 (C)

Key:
MDS = active metadata server P = potential metadata server C = client

Upgrading From 3.4.1 Or Earlier

If you have a cluster with more than two server-capable nodes and you are upgrading from CXFS 3.4.1 or earlier to CXFS 3.4.2 or later, you must install the new CXFS software according to the age of the server-capable node, from youngest to oldest. Use the `clconf_info` command to display the age of the nodes. If the Age fields are the same, use the CellID field, installing from highest CellID number to lowest CellID number.

For example, suppose all of the nodes in the cluster are server-capable nodes and the output from `clconf_info` is as follows:

```
[root@mvcxfs1 ~]# clconf_info

Event at [2006-06-20 11:20:25]

Membership since Tue Jun 20 11:20:25 2006

-----
Node           NodeID Status  Age   CellID
-----
mvcxfs1        1  up      1     0
mvcxfs2        2  inactive -     1
mvcxfs7        7  DOWN    -     4
mvcxfs16       16 up       2     3
mvcxfs17       17 up       1     2
-----

1 CXFS FileSystems
/dev/cxvm/volume_Als0 on /CXFS/volume_A1 enabled server=(mvcxfs1) 2
client(s)=(mvcxfs17,mvcxfs16) status=UP
```

Therefore, the order of installation would be as follows, from first to last:

1. `mvcxfs17` (because it is the newest, with Age=1 and CellID=2)
2. `mvcxfs1`
3. `mvcxfs16` (because it is the oldest, with Age=2)

You should install the new software on `mvcxfs2` and `mvcxfs7` before they rejoin the cluster.

Manual CXFS Start/Stop

On administration nodes, the `/etc/init.d/cluster` (IRIX) or `/etc/init.d/cxfs_cluster` (SGI ProPack) script will be invoked automatically during normal system startup and shutdown procedures; on client-only nodes, the script is `/etc/init.d/cxfs_client`. This script starts and stops the processes required to run CXFS.

To start up CXFS processes manually, enter the following commands:

- On an administration node:
 - IRIX:


```
# /etc/init.d/cluster start
Starting cluster services: fs2d cmond cad crsd
# /etc/init.d/cxfs start
Starting CXFS Cluster services:
Starting clconfd:
```
 - SGI ProPack:


```
# /etc/init.d/cxfs_cluster start
Starting cluster services: fs2d cmond cad crsd           [ OK ]
# /etc/init.d/cxfs start
```
- On an IRIX client-only node:


```
# /etc/init.d/cxfs_client start
cxfs_client daemon started
```
- On an SGI ProPack client-only node:


```
# /etc/init.d/cxfs_client start
Loading cxfs modules:           [ OK ]
Mounting devfs filesystems:     [ OK ]
Starting cxfs client:           [ OK ]
```

To stop CXFS processes manually, enter the following command:

- On an administration node:
 - IRIX:


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

– SGI ProPack:

```
# /etc/init.d/cxfs stop
# /etc/init.d/cxfs_cluster stop
```

- On an IRIX client-only node:

```
# /etc/init.d/cxfs_client stop
Shutting down CXFS client
```

- On an SGI ProPack client-only node:

```
# /etc/init.d/cxfs_client stop
Stopping cxfs client: [ OK ]
```

Note: There is also a `restart` option that performs a stop and then a start.

To see the current status of the CXFS processes, use the `status` argument. For example, the following output shows that `cxfs_client` is running:

```
# /etc/init.d/cxfs_client status
cxfs_client (pid 3226) is running...
```

The output in the following example shows that the CXFS client is stopped on a client-only node:

```
# /etc/init.d/cxfs_client status
cxfs_client is stopped
```

CXFS and Cluster Administration Initialization Commands

Table 13-1 summarizes the `/etc/init.d` initialization commands used for the CXFS control daemon and the cluster administration daemons. Paths may differ between IRIX and SGI ProPack systems.

Table 13-1 CXFS and Cluster Administration Initialization Commands

IRIX	SGI ProPack	Description
<code>/etc/init.d/cluster start</code>	<code>/etc/init.d/cxfs_cluster start</code>	Starts the <code>fs2d</code> , <code>cmond</code> , <code>cad</code> , and <code>crsd</code> (the cluster administration daemons) on the local node
<code>/etc/init.d/cxfs start</code>	<code>/etc/init.d/cxfs start</code>	Starts <code>clconfd</code> (the CXFS control daemon) on the local node
<code>/etc/init.d/cluster stop</code>	<code>/etc/init.d/cxfs_cluster stop</code>	Stops <code>fs2d</code> , <code>cmond</code> , <code>cad</code> , and <code>crsd</code> on the local node
<code>/etc/init.d/cxfs stop</code>	<code>/etc/init.d/cxfs stop</code>	Stops CXFS in the kernel (which withdraws membership) and <code>clconfd</code> on the local node
<code>/etc/init.d/cluster restart</code>	<code>/etc/init.d/cxfs_cluster restart</code>	Restarts the cluster administration daemons on the local node
<code>/etc/init.d/cxfs restart</code>	<code>/etc/init.d/cxfs start</code>	Restarts <code>clconfd</code> on the local node
<code>/etc/init.d/cluster status</code>	<code>/etc/init.d/cxfs_cluster status</code>	Gives status (running or stopped) of <code>fs2d</code> , <code>cmond</code> , <code>cad</code> , and <code>crsd</code> on the local node
<code>/etc/init.d/cxfs status</code>	<code>/etc/init.d/cxfs status</code>	Gives status (running or stopped) of <code>clconfd</code> on the local node

Using `hafence` to Manipulate a Switch

To add or modify a switch:

```
/usr/cluster/bin/hafence -a -s switchname -u username -p password -m mask [-L vendor]
```

To raise the fence for a node:

```
/usr/cluster/bin/hafence -r nodename
```

To lower the fence for a node:

```
/usr/cluster/bin/hafence -l nodename
```

To query switch status:

```
/usr/cluster/bin/hafence -q -s switchname
```

Usage notes:

- `-a` adds or changes a switch in cluster database
- `-l` lowers the fence for the specified node
- `-L` specifies the vendor name, which loads the appropriate plug-in library for the switch. If you do not specify the vendor name, the default is `brocade`
- `-m` specifies one of the following:

- A list of ports in the switch that will never be fenced. The list has the following form, beginning with the `#` symbol, separating each port number with a comma, and enclosed within quotation marks:

```
"#port,port,port . . ."
```

Each *port* is a decimal integer in the range 0 through 1023. For example, the following indicates that port numbers 2, 4, 5, 6, 7, and 23 will never be fenced:

```
-m "#2,4,5,6,7,23"
```

- A hexadecimal string that represents ports in the switch that will never be fenced. Ports are numbered from 0. If a given bit has a binary value of 0, the port that corresponds to that bit is eligible for fencing operations; if 1, then the port that corresponds to that bit will always be excluded from any fencing operations. For an example, see Figure 11-5 on page 235.

CXFS administration nodes automatically discover the available HBAs and, when fencing is triggered, fence off all of the Fibre Channel HBAs when the `Fence` or `FenceReset` fail action is selected. However, masked HBAs will not be fenced. Masking allows you to prevent the fencing of devices that are attached to the SAN but are not shared with the cluster, to ensure that they remain available regardless of CXFS status. You would want to mask HBAs used for access to tape storage, or HBAs that are only ever used to access local (nonclustered) devices.

- `-p` specifies the password for the specified *username*.
- `-q` queries switch status
- `-r` raises the fence for the specified node
- `-s` specifies the hostname of the Fibre Channel switch; this is used to determine the IP address of the switch.
- `-u` specifies the user name to use when sending a telnet message to the switch.

For example, the following defines a QLogic switch named `myqlswitch` and uses no masking:

```
# /usr/cluster/bin/hafence -a -s myqlswitch -u admin -p *** -L qlogic
```

The above command line will attempt to load the `libcrf_vendor.so` library, which must be installed in a directory that is searched by `dlopen(3)`, which is usually `/usr/lib` on SGI ProPack systems and `/usr/lib32` on IRIX systems. However, the shared library search path is platform dependent and site configurable; therefore, it may be somewhere else if the `LD_LIBRARY_PATH` environment variable has been set. See the `dlopen(3)` man page for details.

The following masks port numbers 2 and 3:

```
# /usr/cluster/bin/hafence -a -s myqlswitch -u admin -p *** -m "#2,3" -L qlogic
```

The following lowers the fence for `client1`:

```
# /usr/cluster/bin/hafence -l client1
```

The following raises the fence for `client1`:

```
# /usr/cluster/bin/hafence -r client1
```

The following queries port status for all switches defined in the cluster database:

```
# /usr/cluster/bin/hafence -q
```

For more information, see the `hafence(1M)` man page. See the release notes for supported switches.

CXFS Port Usage

CXFS uses the following ports:

- Fencing requires TCP port 23 for `telnet` access
- The RPC port mapper requires UDP port 111 and TCP port 111
- The `fs2d` daemon is RPC-based and is dynamically assigned on a TCP port in the range of 600-1023. The instance of `fs2d` that determines the cluster database membership also uses TCP port 5449.
- The `crsd` daemon defaults to UDP port 7500 and is set in `/etc/services`:

```
sgi-crsd          7500/tcp
```

- The CXFS kernel uses ports 5450 through 5453 (TCP for ports 5450 and 5451, UDP for ports 5052 and 5053)
- The server-capable node that is the quorum leader uses UDP port 5449.
- The `cad` daemon defaults to TCP port 9000 and is set in `/etc/services`:

```
sgi-cad          9000/tcp
```

For more information, see Appendix C, "IP Filtering for the CXFS Private Network" on page 529.

CXFS `chkconfig` Arguments

Table 13-2 summarizes the CXFS `chkconfig` arguments for IRIX and SGI ProPack nodes. These settings are not normally manipulated by the administrator; they are set or unset by the CXFS GUI. These settings only control the processes, not the cluster. Stopping the processes that control the cluster will not stop the cluster (that is, will not drop the cluster membership or lose access to CXFS filesystems and cluster volumes), and starting the processes will start the cluster **only** if the CXFS services are marked as activated in the database.

Note: `cxfs_cluster` controls different daemons on IRIX than it does on SGI ProPack.

On SGI ProPack nodes, `chkconfig` settings are saved by updating various symbolic links in the `/etc/rc.n` directories.

The following shows the settings of the arguments on IRIX and SGI ProPack administration nodes:

- IRIX:

```
irix# chkconfig | grep cluster
      cluster          on
      cxfs_cluster     on
```

- SGI ProPack:

```
[root@linux root]# chkconfig --list | grep cxfs fam
cxfs_cluster 0:off 1:off 2:on 3:on 4:on 5:on 6:off
cxfs         0:off 1:off 2:on 3:on 4:on 5:on 6:off
fam          0:off 1:off 2:on 3:on 4:on 5:on 6:off
```

Table 13-2 CXFS chkconfig Arguments

IRIX Admin	SGI ProPack Admin	Client-Only (IRIX or SGI ProPack)	Description
cluster	cxfs_cluster	N/A	Controls the cluster administration daemons (<code>fs2d</code> , <code>crsd</code> , <code>cad</code> , and <code>cmond</code>). If this argument 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 database will not be accessible locally and the node will not be configured to join the cluster.
cxfs_cluster	cxfs	N/A	Controls the <code>clconfd</code> daemon and whether or not the <code>cxfs_shutdown</code> command is used during a system shutdown. The <code>cxfs_shutdown</code> 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: <code>clconfd</code> cannot start unless <code>fs2d</code> is already running.
N/A	N/A	cxfs_client	Controls whether or not the <code>cxfs_client</code> daemon should be started
N/A	fam	N/A	Starts the file alteration monitoring (<code>fam</code>) service, which is required to use the CXFS GUI on SGI ProPack nodes

Configuring Real-Time Filesystems For IRIX Nodes

CXFS can write to real-time files in real-time volumes on IRIX nodes. For more details about real-time volumes, see the *XVM Volume Manager Administrator's Guide*.

When creating the CXFS filesystem, be aware of the following:

- To maintain appropriate performance of the real-time filesystem, do not flag unwritten extents. Use the following command:

```
irix# mkfs_xfs -d unwritten=0
```

- Set the real-time extent size to a large value for maximum performance. This parameter should be a multiple of the basic filesystem block size, and can vary between 4 KB to 1 GB. SGI recommends 128 MB. You can set this value with the following command:

```
irix# mkfs_xfs -r extsize=size_of_real-time_extent
```

- Use a large value for block size. Linux systems are not capable of accessing filesystems with block size larger than the system page size. If the filesystem is to be accessible by all nodes in the cluster, its block size must be the lowest common denominator.

You can set this value with the following command:

```
irix# mkfs_xfs -b size=blocksize
```

Granting Task Execution Privileges to Users

The CXFS GUI lets you grant or revoke access to a specific GUI task for one or more specific users. By default, only `root` may execute tasks in the GUI. Access to the task is only allowed on the node to which the GUI is connected; if you want to allow access on another node in the pool, you must connect the GUI to that node and grant access again.

Note: You cannot grant or revoke tasks for users with a user ID of 0.

CXFS GUI tasks operate by executing underlying privileged commands that are normally accessible only to `root`. When granting access to a task, you are in effect granting access to all of its required underlying commands, which results in also granting access to the other GUI tasks that use the same underlying commands. The `cxfs_admin` command provides similar functionality with the `allow|deny` subcommands.

For instructions about granting or revoking CXFS GUI privileges, see "Privileges Tasks with the GUI" on page 249.

To see which tasks a specific user can currently access, select **View: Users**. Select a specific user to see details about the tasks available to that user.

To see which users can currently access a specific task, select **View: Task Privileges**. Select a specific task to see details about the users who can access it and the privileged commands it requires.

Transforming an Existing Node into a Client-Only Node

If you are upgrading to 6.5.19f from 6.5.17f or earlier and you want to change an existing node with weight 1 (which as of 6.5.18f was defined as a *server-capable administration node*) to be a client-only node, you must do the following:

1. Ensure that the node is not listed as a potential metadata server for any filesystem.
2. Stop the CXFS services on the node.
3. Modify the cluster so that it no longer contains the node.
4. Delete the node definition.
5. Install the node with the `cxfs_client` package and remove the `cluster_admin`, `cluster_control`, and `cluster_services` packages.
6. Reboot the node to ensure that all previous node configuration information is removed.
7. Redefine the node and use a node function of client-only.
8. Modify the cluster so that it contains the node.
9. Start the CXFS services on the node.

CXFS Mount Scripts

Scripts are provided for execution prior to and after a CXFS filesystem is mounted or unmounted on the following platforms:

- On server-capable nodes:

```
/var/cluster/clconfd-scripts/cxfs-pre-mount
/var/cluster/clconfd-scripts/cxfs-post-mount
/var/cluster/clconfd-scripts/cxfs-pre-umount
/var/cluster/clconfd-scripts/cxfs-post-umount
```

The `clconfd` daemon executes the above scripts.

- On client-only nodes:

```
/var/cluster/cxfs_client-scripts/cxfs-pre-mount
/var/cluster/cxfs_client-scripts/cxfs-post-mount
/var/cluster/cxfs_client-scripts/cxfs-pre-umount
/var/cluster/cxfs_client-scripts/cxfs-post-umount
```

The `cxfs_client` daemon executes the above scripts.

The scripts are used by CXFS to ensure that LUN path failover works properly after fencing by executing the following:

```
/etc/init.d/failover stop
/etc/init.d/failover start
```

These scripts can be customized to suit a particular environments. For example, an application could be started when a CXFS filesystem is mounted by extending the `cxfs-post-mount` script. The application could be terminated by changing the `cxfs-pre-umount` script.

On IRIX and SGI ProPack nodes, these scripts also allow you to use NFS to export the CXFS filesystems listed in `/etc/exports` if they are successfully mounted.

The appropriate daemon executes these scripts before and after mounting or unmounting CXFS filesystems specified in the `/etc/exports` file. The files must be named **exactly** as above and must have root execute permission.

Note: The `/etc/exports` file describes the filesystems that are being exported to NFS clients. If a CXFS mount point is included in the `exports` file, the empty mount point is exported unless the filesystem is re-exported after the CXFS mount using the `cxfs-post-mount` script.

The `/etc/exports` file cannot contain any filesystems managed by FailSafe.

The following arguments are passed to the files:

- `cxfs-pre-mount`: filesystem device name and CXFS mounting point
- `cxfs-post-mount`: filesystem device name, CXFS mounting point, and exit code
- `cxfs-pre-umount`: filesystem device name and CXFS mounting point
- `cxfs-post-umount`: filesystem device name, CXFS mounting point, and exit code

Because the filesystem name is passed to the scripts, you can write the scripts so that they take different actions for different filesystems; because the exit codes are passed to the `post` files, you can write the scripts to take different actions based on success or failure of the operation.

The `clconfd` or `cxfs_client` daemon checks the exit code for these scripts. In the case of failure (nonzero), the following occurs:

- For `cxfs-pre-mount` and `cxfs-pre-umount`, the corresponding mount or unmount is not performed.
- For `cxfs-post-mount` and `cxfs-post-umount`, `clconfd` will retry the entire operation (including the `-pre-` script) for that operation.

This implies that if you **do not** want a filesystem to be mounted on a host, the `cxfs-pre-mount` script should return a failure for that filesystem while the `cxfs-post-mount` script returns success.

Note: After the filesystem is unmounted, the mount point is removed.

The following script is run when needed to reprobe the Fibre Channel controllers:

- On server-capable nodes:

```
/var/cluster/clconfd-scripts/cxfs-reprobe
```

- On client-only nodes:

```
/var/cluster/cxfs_client-scripts/cxfs-reprobe
```

You may modify any of these scripts if needed.

Unmounting lofs File Systems

You must unmount lofs mounts of a CXFS filesystem before attempting to unmount the CXFS filesystem. You can use a script such as the following to unexport and locally unmount an lofs filesystem:

```
#!/bin/ksh
#/var/cluster/clconfd-scripts/cxfs-pre-umount
echo "$0: Preparing to unmount CXFS file system \"$1\" "
MNTPOINT=`mount | grep "$1 " | cut -f 3 -d" "`
print "MNTPOINT $MNTPOINT"
if [ -n "${MNTPOINT}" ] ; then
    lofslist=`mount | grep 'type lofs' | grep "${MNTPOINT}" | nawk '{print $3}'`
    set -e
    for lofs in ${lofslist}
    do
        echo "$0: unmounting $lofs"
        umount -k $lofs
    done
    if /usr/etc/exportfs | /sbin/grep -q "${MNTPOINT}" ; then
        echo "$0: unexporting $MNTPOINT"
        /usr/etc/exportfs -u ${MNTPOINT}
    fi
fi
```

Using telnet and I/O Fencing

If there are problems with a node, the I/O fencing software sends a message via the telnet protocol to the appropriate Fibre Channel switch. The switch only allows one telnet session at a time; therefore, if you are using I/O fencing, you must keep the

telnet port on the Fibre Channel switch free at all times. **Do not** perform a telnet to the switch and leave the session connected.

Using `fsr` and `xfs_fsr`

The IRIX `fsr` and the Linux `xfs_fsr` commands can **only** be used on the active metadata server for the filesystem; the `bulkstat` system call has been disabled for CXFS clients. You should use `fsr` or `xfs_fsr` manually, and only on the active metadata server for the filesystem.

Using `cron` in a CXFS Cluster

The `cron` daemon can cause severe stress on a CXFS filesystem if multiple nodes in a cluster start the same filesystem-intensive task simultaneously. An example of such a task is one that uses the `find` command to search files in a filesystem.

Any task initiated using `cron` on a CXFS filesystem should be launched from a single node in the cluster, preferably from the active metadata server.

Using Hierarchical Storage Management (HSM) Products

CXFS supports the use of hierarchical storage management (HSM) products through the data management application programming interface (DMAPI), also known as X/Open Data Storage Management Specification (XSDM). An example of an HSM product is the Data Migration Facility (DMF). DMF is the only HSM product currently supported with CXFS.

The HSM application must make all of its DMAPI interface calls through the active metadata server. The CXFS client nodes do not provide a DMAPI interface to CXFS mounted filesystems. A CXFS client routes all of its communication to the HSM application through the metadata server. This generally requires that the HSM application run on the CXFS metadata server.

To use HSM with CXFS, do the following:

- Install `eo.e.sw.dmi` on each CXFS administration node. For client-only nodes, no additional software is required.
- Use the `dmi` option when mounting a filesystem to be managed.

- Start the HSM application on the active metadata server for each filesystem to be managed.

Discovering the Active Metadata Server for a Filesystem

You can discover the active metadata server using the CXFS GUI, `cxfs_admin`, or `clconf_info`:

- "Discovering the Active Metadata Server with the CXFS GUI" on page 329
- "Discovering the Active Metadata Server with `cxfs_admin`" on page 330
- "Discovering the Active Metadata Server with `clconf_info`" on page 331

See and the following sections.

Discovering the Active Metadata Server with the CXFS GUI

Do the following:

1. Select **View: Filesystems**
2. In the view area, click the name of the filesystem you wish to view. The name of the active metadata server is displayed in the details area to the right.

Figure 13-4 shows an example.

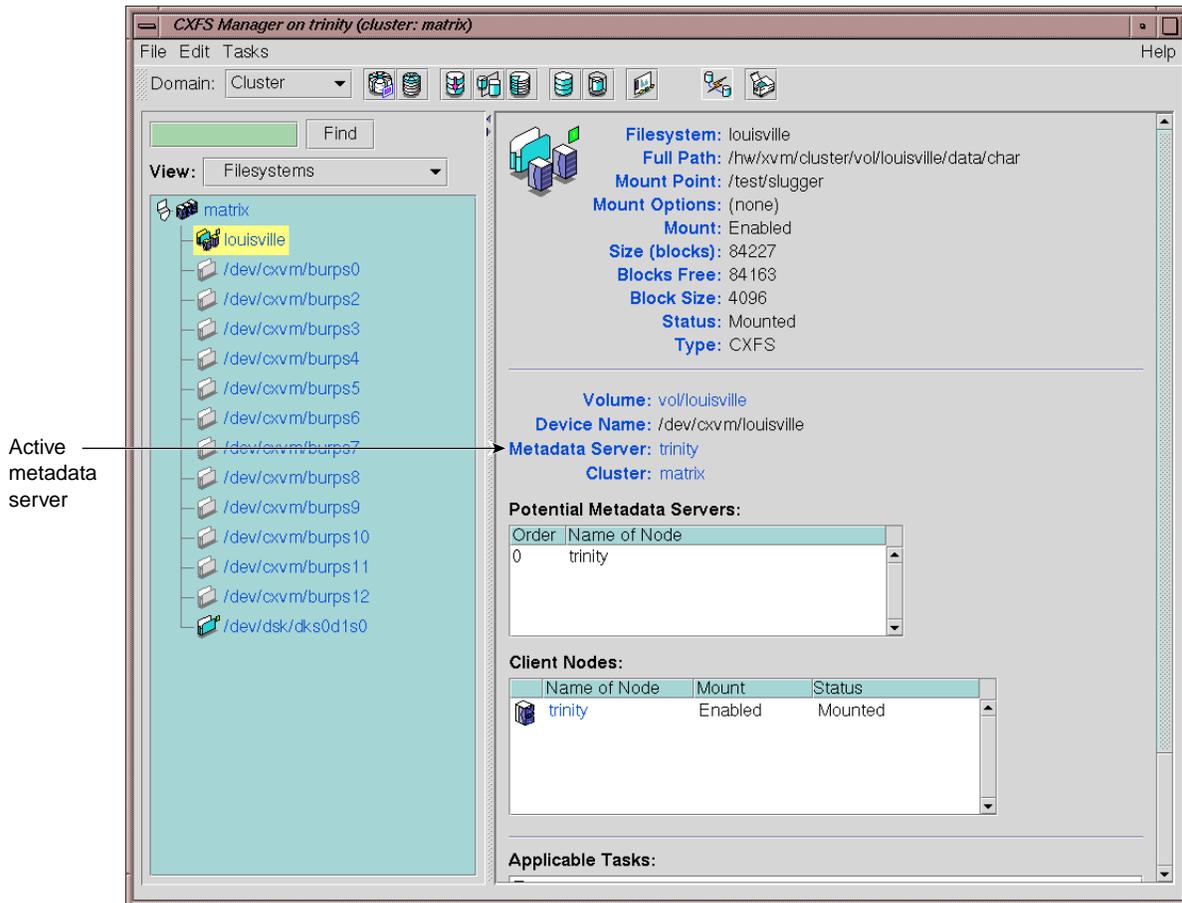


Figure 13-4 Window Showing the Metadata Server

Discovering the Active Metadata Server with `cxfs_admin`

To discover the active metadata server for a filesystem, use the `show` command:

- To show information for all filesystems, including their active metadata servers:

```
show server
```

For example:

```
cxfs_admin:mycluster> show server
filesystem:concatfs:status:server=mds1
filesystem:mirrorfs:status:server=mds1
filesystem:stripefs:status:server=mds2
```

- To show the active metadata server for a specific filesystem:

```
show [filesystem:]filesystem:status:server
```

In the above, you could abbreviate `status` to `*`. For example, if `concatfs` is a unique name in the cluster database:

```
cxfs_admin:mycluster> show concatfs:*:server
filesystem:concatfs:status:server=mds1
```

Discovering the Active Metadata Server with `clconf_info`

You can use the `clconf_info` command to discover the active metadata server for a given filesystem. For example, the following shows that `cxfs7` is the metadata server:

```
cxfs6 # clconf_info
```

```
Event at [2004-04-16 09:20:59]
```

```
Membership since Fri Apr 16 09:20:56 2004
```

Node	NodeID	Status	Age	CellID
cxfs6	6	up	0	2
cxfs7	7	up	0	1
cxfs8	8	up	0	0

```
1 CXFS FileSystems
```

```
/dev/cxvm/concat0 on /concat0 enabled server=(cxfs7) 2 client(s)=(cxfs8,cxfs6)
```

Metadata Server Recovery

Note: Recovery is supported only when using standby nodes.

If the node acting as the metadata server for a filesystem dies, another node in the list of potential metadata servers will be chosen as the new metadata server. This assumes that at least two potential metadata servers are listed when you define a filesystem.

The metadata server that is chosen must be a filesystem client; other filesystem clients will experience a delay during the relocation process. Each filesystem will take time to recover, depending upon the number of active inodes; the total delay is the sum of time required to recover each filesystem. Depending on how active the filesystems are at the time of recovery, the total delay could take up to several minutes per filesystem.

If a CXFS client dies, the metadata server will clean up after the client. Other CXFS clients may experience a delay during this process. A delay depends on what tokens, if any, that the deceased client holds. If the client has no tokens, then there will be no delay; if the client is holding a token that must be revoked in order to allow another client to proceed, then the other client will be held up until recovery returns the failed nodes tokens (for example, in the case where the client has the write token and another client wants to read). The actual length of the delay depends upon the following:

- The total number of exported inodes on the metadata server
- CXFS kernel membership situation
- Whether any servers have died
- Where the servers are in the recovery order relative to recovering this filesystem

The deceased CXFS client is not allowed to rejoin the CXFS kernel membership until all metadata servers have finished cleaning up after the client.

Shutdown of the Database and CXFS

This section tells you how to perform the following:

- "Cluster Database Shutdown" on page 333
- "Normal CXFS Shutdown: Stop CXFS Services" on page 335

- "Forced CXFS Shutdown: Revoke Membership of Local Node" on page 337

If there are problems, see Chapter 19, "Troubleshooting" on page 425. For more information about states, Chapter 17, "Monitoring Status" on page 399.

Cluster Database Shutdown

A *cluster database shutdown* terminates the following user-space daemons that manage the cluster database:

```
cad
clconfd
cmond
crsd
fs2d
```

After shutting down the database on a node, access to the shared filesystems remains available and the node is still a member of the cluster, but the node is not available for database updates. Rebooting of the node results in a restart of all services (restarting the daemons, joining cluster membership, enabling cluster volumes, and mounting CXFS filesystems).

To perform a cluster database shutdown, enter the following:

- IRIX:

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

- SGI ProPack:

```
[root@linux root]# /etc/init.d/cxfs_cluster stop
```

If you also want to disable the daemons from restarting at boot time, enter the following:

- IRIX:

```
irix# chkconfig cluster off
```

- SGI ProPack:

```
[root@linux root]# chkconfig cxfs_cluster off
```

For more information, see "CXFS chkconfig Arguments" on page 320.

Node Status and Cluster Database Shutdown

A cluster database shutdown is appropriate when you want to perform a maintenance operation on the node and then reboot it, returning it to `ACTIVE` status.

If you perform a cluster database shutdown, the node status will be `DOWN`, which has the following impacts:

- The `DOWN` node is still considered part of the cluster, but unavailable.
- The `DOWN` node does not get cluster database updates; however, it will be notified of all updates after it is rebooted.

Missing cluster database updates can cause problems if the kernel portion of CXFS is active. That is, if the node continues to have access to CXFS, the node's kernel level will not see the updates and will not respond to attempts by the remaining nodes to propagate these updates at the kernel level. This in turn will prevent the cluster from acting upon the configuration updates.

Restart the Cluster Database

To restart the cluster database, enter the following:

- IRIX:

```
# /etc/init.d/cluster start
```
- SGI ProPack:

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

Normal CXFS Shutdown: Stop CXFS Services

You should perform a *normal CXFS shutdown* when you want to stop CXFS services on a node and remove it from the CXFS kernel membership quorum. A normal CXFS shutdown does the following:

- Unmounts all the filesystems except those for which it is the active metadata server; those filesystems for which the node is the active metadata server will become inaccessible from the node after it is shut down.
- Terminates the CXFS kernel membership of this node in the cluster.
- Marks the node as `INACTIVE`.

The effect of this is that cluster disks are unavailable and no cluster database updates will be propagated to this node. Rebooting the node leaves it in the shutdown state.

If the node on which you shut down CXFS services is an active metadata server for a filesystem, then that filesystem will be recovered by another node that is listed as one of its potential metadata servers. The server that is chosen must be a filesystem client; other filesystem clients will experience a delay during the recovery process.

If the node on which the CXFS shutdown is performed is the sole potential metadata server (that is, there are no other nodes listed as potential metadata servers for the filesystem), then you should use `umount` to unmount the filesystem from all nodes before performing the shutdown.

To perform a normal CXFS shutdown, you could use the CXFS GUI; see "Stop CXFS Services with the GUI" on page 228. If you are using `cxfs_admin`, you can disable individual nodes with the `disable` command.

Note: This action deactivates CXFS services on **one** node, forming a new CXFS kernel membership after deactivating the node. If you want to stop CXFS services on multiple nodes, you must enter this command multiple times or perform the task using the GUI.

After you stop CXFS services on a node, the node is marked as inactive and is no longer used when calculating the CXFS kernel membership. See "Node Status" on page 406.

Node Status and Stopping CXFS Services

After performing stopping CXFS services on a node, its state will be `INACTIVE`; therefore, it will not impact CXFS kernel membership quorum calculation. See "Normal CXFS Shutdown: Stop CXFS Services" on page 335.

When You Should Not Perform Stop CXFS Services

You should not stop CXFS services under the following circumstances:

- On the *local node*, which is the CXFS administration node on which the cluster manager is running or the node to which the CXFS GUI is connected
- If stopping CXFS services on the node will result in loss of CXFS kernel membership quorum
- If the node is the only available metadata server for one or more active CXFS filesystems

If you want to perform a CXFS shutdown under these conditions, you must perform a forced CXFS shutdown. See "Forced CXFS Shutdown: Revoke Membership of Local Node" on page 337.

Rejoining the Cluster after Stopping CXFS Services

The node will not rejoin the cluster after a reboot. The node will rejoin the cluster only when CXFS services are explicitly reactivated with the CXFS GUI. In `cxfs_admin`, you can disable individual nodes with the `disable` command.

Forced CXFS Shutdown: Revoke Membership of Local Node

A *forced CXFS shutdown* (or *administrative CXFS stop*) is appropriate when you want to shutdown the local node even though it may drop the cluster below its CXFS kernel membership quorum requirement.

CXFS does the following:

- Shuts down all CXFS filesystems on the local node
- Attempts to access the CXFS filesystems result in I/O error (you may need to manually unmount the filesystems)
- Removes this node from the CXFS kernel membership
- Marks the node as `DOWN`
- Disables access from the local node to cluster-owned XVM volumes.
- Treats the stopped node as a failed node and executes the failure action defined for the node in the cluster database (`reset`, `fence`, `fencerreset`, or `shutdown`).



Caution: A forced CXFS shutdown may cause the cluster to fail if the cluster drops below CXFS kernel membership quorum.

If you do a forced CXFS shutdown on an active metadata server, it loses membership immediately. At this point another potential metadata server must take over (and recover the filesystems) or quorum is lost and a forced shutdown follows on all nodes.

If you do a forced CXFS shutdown that forces a loss of quorum, the remaining part of the cluster (which now must also do an administrative stop) will **not** reset the departing node.

To perform an administrative stop, see "Revoke Membership of the Local Node with the GUI" on page 232. If you are using `cxfs_admin`, you can disable individual nodes with the `disable` subcommands.

Node Status and Forced CXFS Shutdown

After a forced CXFS shutdown, the node is still considered part of the configured cluster and is taken into account when propagating the cluster database (these services are still running) and when computing the cluster database (`fs2d`)

membership quorum (this could cause a loss of quorum for the rest of the cluster, causing the other nodes to do a forced CXFS shutdown). The state is `INACTIVE`.

It is important that this node stays accessible and keeps running the cluster infrastructure daemons to ensure database consistency. In particular, if more than half the nodes in the pool are down or not running the infrastructure daemons, cluster database updates will stop being propagated and will result in inconsistencies. To be safe, you should remove those nodes that will remain unavailable from the cluster and pool.

Rejoining the Cluster after a Forced CXFS Shutdown

After a forced CXFS shutdown, the local node will not resume CXFS kernel membership until the node is rebooted or until you explicitly allow CXFS kernel membership for the local node.

With the CXFS GUI, see "Allow Membership of the Local Node with the GUI" on page 232. In `cxfs_admin`, you can disable individual nodes with the `disable|enable` subcommands.

If you perform a forced CXFS shutdown on a CXFS administration node, you must restart CXFS on that node before it can return to the cluster. If you do this while the cluster database still shows that the node is in a cluster and is activated, the node will restart the CXFS kernel membership daemon. Therefore, you may want to do this after resetting the database or after stopping CXFS services.

Reset Capability and a Forced CXFS Shutdown



Caution: If you perform an administrative CXFS stop on an administration node with system reset capability and the stop will not cause loss of cluster quorum, the node will be reset (rebooted) by the appropriate node.

For more information about resets, see "Reset" on page 34.

Avoiding a CXFS Restart at Reboot

If the following `chkconfig` arguments are turned off, the `clconfd` and `cxfs_client` daemons on CXFS administration nodes and client-only nodes,

respectively, will not be started at the next reboot and the kernel will not be configured to join the cluster:

- IRIX administration nodes: `cxfs_cluster`
- SGI ProPack administration nodes: `cxfs`
- Client-only nodes: `cxfs_client`

It is useful to turn these arguments off before rebooting if you want to temporarily remove the nodes from the cluster for system or hardware upgrades or for other maintenance work.

For example, do the following:

- IRIX administration node:

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

- SGI ProPack administration node:

```
[root@linux root]# /sbin/chkconfig cxfs off
[root@linux root]# /sbin/chkconfig cxfs_cluster off
[root@linux root]# reboot
```

For more information, see "CXFS `chkconfig` Arguments" on page 320.

Log File Management

CXFS log files should be rotated at least weekly so that your disk will not become full. This section discusses the following:

- "Rotating Log Files on SGI ProPack" on page 340
- "Rotating All Log Files on IRIX Server-Capable Nodes" on page 340
- "Rotating Large Log Files on IRIX" on page 341

For information about log levels, see "Configure Log Groups with the GUI" on page 231.

Rotating Log Files on SGI ProPack

On SGI ProPack, a package that provides CXFS daemons also supplies scripts to rotate the log files for those daemons.

SGI places the following scripts in `/etc/logrotate.d` for server-capable nodes:

```
/etc/logrotate.d/cluster_admin  
/etc/logrotate.d/cluster_control  
/etc/logrotate.d/cxfs_cluster
```

SGI places the following script in `/etc/logrotate.d` for client-only nodes:

```
/etc/logrotate.d/cxfs_client
```

To customize log rotation, edit the supplied CXFS log rotation scripts mentioned above.

Rotating All Log Files on IRIX Server-Capable Nodes

On IRIX server-capable nodes, you can run the `/var/cluster/cmgr-scripts/rotatelogs` script to copy all files to a new location (this script works regardless of the configuration tool used). This script saves log files with the day and the month name as a suffix. If you run the script twice in one day, it will append the current log file to the previous saved copy. The `root` crontab file has an entry to run this script weekly.

The script syntax is as follows:

```
/var/cluster/cmgr-scripts/rotatelogs [-h] [-d|-u]
```

If no option is specified, the log files will be rotated. Options are as follows:

- h Prints the help message. The log files are not rotated and other options are ignored.
- d Deletes saved log files that are older than one week before rotating the current log files. You cannot specify this option and `-u`.
- u Unconditionally deletes all saved log files before rotating the current log files. You cannot specify this option and `-d`.

By default, the `rotatelogs` script will be run by `crontab` once a week, which is sufficient if you use the default log levels. If you plan to run with a high debug level

for several weeks, you should reset the crontab entry so that the `rotatelogs` script is run more often.

On heavily loaded machines, or for very large log files, you may want to move resource groups and stop CXFS services before running `rotatelogs`.

Rotating Large Log Files on IRIX

On IRIX, you can use a script such as the following to copy large files to a new location. The files in the new location will be overwritten each time this script is run.

```
#!/bin/sh
# Argument is maximum size of a log file (in characters) - default: 500000

size=${1:-500000}
find /var/cluster/ha/log -type f ! -name '*.OLD' -size +${size}c -print | while read log_file; do
    cp ${log_file} ${log_file}.OLD
    echo '*** LOG FILE ROTATION ' `date` '***' > ${log_file}
done
```

Also see "`cad.options` on CXFS Administration Nodes" on page 129, and "`fs2d.options` on CXFS Administration Nodes" on page 130

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: The `xvm` command must be run on a CXFS administration node. If you try to run an XVM command before starting the CXFS daemons, you will get a warning message and be put into XVM's *local domain*.

When you are in XVM's local domain, you could define your filesystems, but then when you later start up CXFS you will not see the filesystems. When you start up CXFS, XVM will switch to *cluster domain* and the filesystems will not be recognized because you defined them in local domain; to use them in the cluster domain, you would have to use the `give` command. Therefore, it is better to define the volumes directly in the cluster domain.

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

Disk Management

This section describes the CXFS differences for backups, NFS, quotas, and Samba.

Disk 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.

Note: Do not run backups on a client node because it causes heavy use of non-swappable kernel memory on the metadata server. During a backup, every inode on the filesystem is visited, and if done from a client, it imposes a huge load on the metadata server. The metadata server may experience typical out-of-memory symptoms, and in the worst case can even become unresponsive or crash.

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 filesystems. You should run the NFS server on the CXFS active metadata server for optimal performance.

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 any IRIX or Linux node in the cluster.

For more information about setting quotas, see *XFS for Linux Administration* and *IRIX Admin: Disks and Filesystems*.

Samba

You can run Samba on top of CXFS, allowing Windows machines to support CXFS and have access to the filesystem. Samba should run on the active metadata server for optimal performance. You should not serve the same CXFS filesystem from multiple nodes in a cluster.

The architecture of Samba assumes that each share is exported by a single server. Because all Samba client accesses to files and directories in that share are directed through a single Samba server, the Samba server is able to maintain private metadata state to implement the required concurrent access controls (in particular, share modes, write caching and oplock states). This metadata is not necessarily promulgated to the filesystem and there is no protocol for multiple Samba servers exporting the same share to communicate this information between them.

Running multiple Samba servers on one or more CXFS clients exporting a single share that maps to a common underlying filesystem has the following risks:

- File data corruption from writer-writer concurrency
- Application failure due to inconsistent file data from writer-reader concurrency

These problems do not occur when a single Samba server is deployed, because that server maintains a consistent view of the metadata used to control concurrent access across all Samba clients.

It may be possible to deploy multiple Samba servers under one of the following circumstances:

- There are no writers, so a read-only share is exported
- Application-level protocols and/or work-flow guarantee that only one application is ever writing a file, and concurrent file writing and reading does not take place



Caution: The onus is on the customer to ensure these conditions are met, as there is nothing in the Samba architecture to verify it. Therefore, SGI recommends that you do not use multiple Samba servers.

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 database instead.

As the administrator, you will supply the CXFS filesystem configuration by using the CXFS GUI or `cxfs_admin`.

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.



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 CXFS GUI or `cxfs_admin`.



Caution: Do not attempt to use the `mount` 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 nodes.

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.

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 CXFS GUI or `cxfs_admin`.

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 active metadata servers, the unmount cannot succeed before all of the CXFS 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 CXFS GUI or `cxfs_admin` 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 CXFS 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 using the CXFS GUI or `cxfs_admin`.
2. Change the domain of the XVM volume from a cluster volume to a local volume using the XVM `give` command. See the *XVM Volume Manager Administrator's Guide*.
3. Mount the filesystem as an XFS filesystem. See *IRIX Admin: Disks and Filesystems*.
4. Use the `xfs_growfs` command or the CXFS GUI task; see "Grow a Filesystem with the GUI" on page 240.
5. Unmount the XFS filesystem. See *IRIX Admin: Disks and Filesystems*.
6. Change the domain of the XVM volume back to a cluster volume using the `give` command. See the *XVM Volume Manager Administrator's Guide*.
7. Mount the filesystem as a CXFS filesystem by using the GUI or `cxfs_admin`

Dump and Restore

You must perform the backup of a CXFS filesystem from the metadata server of that filesystem. The `xfsdump` and `xfsrestore` commands make use of special system calls that will only function on the metadata server.

If there are multiple potential metadata servers for a filesystem and the primary server goes down because of this problem, the backup metadata server will gather information on all open files in the cluster. Unless the backup server has much larger memory than the primary server, the result is that it too will go down with exactly the same symptoms that caused the primary server to crash.

You must perform dump and restore procedures from the active metadata server.

The filesystem can have active clients during a dump process.

In a clustered environment, a CXFS filesystem may be directly accessed simultaneously by many CXFS clients and the active metadata server. With failover or metadata server reassignment, a filesystem may, over time, have a number of metadata servers. Therefore, in order for `xfsdump` to maintain a consistent inventory, it must access the inventory for past dumps, even if this information is located on another node.

SGI recommends that the inventory be made accessible by potential metadata server nodes in the cluster using one of the following methods:

- Relocate the inventory to a shared filesystem.

For example, where *shared_filesystem* is replaced with the actual name of the filesystem to be shared:

- On the node currently containing the inventory, enter the following:

```
# cd /var
# cp -r xfsdump /shared_filesystem
# mv xfsdump xfsdump.bak
# ln -s /shared_filesystem/xfsdump xfsdump
```

- On all other administration nodes in the cluster, enter the following:

```
# cd /var
# mv xfsdump xfsdump.bak
# ln -s /shared_filesystem/xfsdump xfsdump
```

- Export the directory using an NFS shared filesystem.

For example:

- On the IRIX node currently containing the inventory, add `/var/xfsdump` to `/etc/exports` and then enter the following:

```
irix# exportfs -a
```

(On SGI ProPack, the path is `/var/lib/xfsdump`.)

- On all other IRIX administration nodes in the cluster, enter the following:

```
# cd /var
# mv xfsdump xfsdump.bak
# ln -s /hosts/hostname/var/xfsdump xfsdump
```

Note: It is the IRIX `/var/xfsdump` directory (SGI ProPack `/var/lib/xfsdump`) that should be shared, rather than the IRIX `/var/xfsdump/inventory` directory (SGI ProPack `/var/lib/xfsdump/inventory`). If there are inventories stored on various nodes, you can use `xfsinvutil` to merge them into a single common inventory, prior to sharing the inventory among the cluster.

Site-Changeable System Tunable Parameters

This section lists the CXFS system tunable parameters that you can change. SGI recommends that you use the same settings on all applicable nodes in the cluster.

Note: Before changing any parameter, you should understand the ramifications of doing so on your system. Contact your SGI support person for guidance.

To manipulate these parameters on a running system, you can use the Linux `sysctl` command or the IRIX `systemtune` command. For more information, see the `sysctl(1M)`, `systemtune(1M)`, and `modules.conf(5)` man pages.

Linux organizes the tunables in a hierarchy, therefore you must specify the entire "path" to the tunable. The first part of that path is given under the "Location" entry in the following sections. For example, the full path to the tunable `cxfsd_sync_force` is `fs.cxfs.cxfsd_sync_force`.

Example of a query using `sysctl`:

```
[root@linux root]# sysctl fs.cxfs.cxfsd_sync_force
fs.cxfs.cxfsd_sync_force = 8372224
```

Example of setting a value using `sysctl`:

```
[root@linux root]# sysctl fs.cxfs.cxfsd_sync_force=0
fs.cxfs.cxfsd_sync_force = 0
```

Note: There cannot be spaces around the = character when setting a value.

There are also several configuration files that you can use to set a tunable automatically. SGI recommends that you use the `/etc/modprobe.conf.local` file. This file specifies options for modules and can be used to set options that cannot be

set with `sysctl`. To set an option, add a line of the following format to `/etc/modprobe.conf.local`:

```
options modulename tunablename=value
```

In this guide, the `modulename` value to be used is given in the "Location" entry in the following sections. For example, `sgi-cxfs` is the module name for the tunable `cxfsd_sync_force`. Therefore to set the value of `cxfsd_sync_force` to `0x7fc000`, you would add the following line to `/etc/modprobe.conf.local`:

```
options sgi-cxfs cxfsd_sync_force=0x7fc000
```

Note: SGI does not recommend using `/etc/sysctl.conf` because it is a global configuration file that might be affected by upgrades of non-related software.

Site-Changeable Static Parameters

Static parameters require a reboot to take affect. On IRIX, you must build and boot new kernels, which happens automatically during a normal boot process. On any of the Linux flavors supported in this CXFS release, you must specify the parameter in `/etc/modprobe.conf.local`.

`mtcp_hb_period`

Specifies (in hundredths of a second) the length of time that CXFS waits for heartbeat from other nodes before declaring node failure. SGI recommends a value of 500 (5 seconds). You should only change this value at the recommendation of SGI support. The same value must be used on all nodes in the cluster.

Range of values:

- Default: 500
- Minimum: 100
- Maximum: 12000

Location:

- IRIX: `/var/sysgen/mtune/cell`
- Linux: `kernel.cell` (`sgi-cell` module)

mtcp_hb_watchdog

Controls the behavior of the heartbeat monitor watchdog. This facility monitors the generation of CXFS heartbeats in the kernel.

Range of values:

- 0 specifies that there is no use of watchdog (default)
- 1 specifies that watchdog expiration causes CXFS shutdown
- 2 specifies that watchdog expiration causes panic

Location:

- IRIX: `/var/sysgen/mtune/cell`
- Linux: `kernel.cell` (`sgi-cell` module)

mtcp_nodelay

Specifies whether to enable or disable `TCP_NODELAY` on CXFS message channels.

Range of values:

- 0 disables
- 1 enables (default)

Location:

- IRIX: `/var/sysgen/mtune/cell`
- Linux: `kernel.cell` (`sgi-cell` module)

mtcp_rpc_thread

Specifies whether metadata messages are sent from a separate thread in order to save stack space.

Range of values:

- 0 disables (default for most nodes)
- 1 enables (default for RHEL 4 32-bit nodes)

Location:

- IRIX: `/var/sysgen/mtune/cell`
- Linux: `kernel.cell` (`sgi-cell` module)

`rhelpld_max`

Specifies the maximum number of `rhelpld` threads to run. The `rhelpld` threads help out recovery and relocation tasks. They are used for asynchronous inode reconstruction, parallel recoveries, and so on. The `rhelpld` thread pool is global in nature and gets created during module load time.

Range of values:

- Default: 0, which specifies an automatically calculated value that will be 4 times the number of CPUs, as long as it is in the range 0 through 128. To disable automatic `rhelpld_max` calculation, set `rhelpld_max` to a non-zero value.
- Minimum: 0
- Maximum: 128

Location:

- IRIX: `/var/sysgen/mtune/cxfs`
- Linux: `fs.cxfs` (`sgi-cxfs` module)

`rhelpld_min`

Specifies the minimum number of `rhelpld` threads to run.

Range of values:

- Default: 0, which specifies an automatically calculated value that is 4 times the number of CPUs or 128, whichever is smaller. To disable automatic `rhelpld_min` calculation, set `rhelpld_min` to a non-zero value. When the value is set explicitly, the maximum is 8.
- Minimum: 0
- Maximum: 8

Location:

- IRIX: `/var/sysgen/mtune/cxfs`

- Linux: `fs.cxfs` (`sgi-cxfs` module)

Site-Changeable Dynamic Parameters

Dynamic parameters take affect as soon as they are changed.

`cms_local_fail_action`

Specifies the action to take when a local node detects that it has failed:

Range of values:

- 0 withdraws from the cluster (default)
- 1 halts
- 2 reboots

Location:

- IRIX: `/var/sysgen/mtune/cell`
- Linux: `kernel.cell` (`sgi-cell` module)

`cxfs_client_push_period`

Specifies (in hundredths of a second) how long that a client may delay telling the metadata server that it has updated the `atime` timestamp of a file. The default for both `cxfs_client_push_period` and `cxfs_server_push_period` is 1/4 of a second, so `atime` updates are delayed by up to 1/2 second by default. See also "`cxfs_server_push_period`" on page 354.

Range of values:

- Default: 25
- Minimum: 0
- Maximum: 1000

Location:

- IRIX: `/var/sysgen/mtune/cxfs`
- Linux: `fs.cxfs` (`sgi-cxfs` module)

cxfs_dcvn_timeout

Specifies the timeout (in seconds) of the dcvn idle period before returning tokens to the server.

Range of values:

- Default: 60
- Minimum: 5
- Maximum: 3600

Location:

- IRIX: `/var/sysgen/mtune/cxfs`
- Linux: `fs.cxfs` (`sgi-cxfs` module)

cxfs_extents_delta

Specifies whether or not to optimize the way extent lists are sent across the private network by sending a delta when possible. This parameter is dynamically tunable.

Range of values:

- 0 does not optimize
- 1 optimizes (default)

Location:

- IRIX: `/var/sysgen/mtune/cxfs`
- Linux: `fs.cxfs` (`sgi-cxfs` module)

cxfs_punch_hole_restrict

Specifies whether or not to allow exported files to have their extents freed by DMAPi via `dm_punch_hole()`.

Range of values:

- 0 allows extents to be freed (default)
- 1 does not allow extents to be freed

Location:

- IRIX: `/var/sysgen/mtune/cxfs`
- Linux: `fs.cxfs` (`sgi-cxfs` module)

`cxfs_relocation_ok`

Specifies whether relocation is disabled or enabled (must be specified on the active metadata server):

Range of values:

- 0 disables relocation (default)
- 1 enables relocation

Note: Relocation is disabled by default and is only supported on standby nodes.

Location:

- IRIX: `/var/sysgen/mtune/cxfs`
- Linux: `fs.cxfs` (`sgi-cxfs` module)

`cxfs_server_push_period`

Specifies (in hundredths of a second) how long that a metadata server may delay broadcasting to the clients that it has updated the `atime` timestamp. The default for both `cxfs_client_push_period` and `cxfs_server_push_period` is 1/4 of a second, so `atime` updates are delayed by up to 1/2 second by default. See also "`cxfs_client_push_period`" on page 352.

Range of values:

- Default: 25
- Minimum: 0
- Maximum: 1000

Location:

- IRIX: `/var/sysgen/mtune/cxfs`

- Linux: `fs.cxfs` (`sgi-cxfs` module)

`cxfsd_max`

Specifies the maximum number of `cxfsd` threads to run per CXFS filesystem. (The `cxfsd` threads do the disk block allocation for delayed allocation buffers in CXFS and the flushing of buffered data for files that are being removed from the local cache by the metadata server.) The threads are allocated at filesystem mount time. The value of the `cxfsd_max` parameter at mount time remains in effect for a filesystem until it is unmounted.

Legal values for `cxfsd_max`:

- Default: 0, which specifies the value of `cxfsd_min` + 2. (The value for `cxfsd_max` is always at least `cxfsd_min` + 2, even if that forces the kernel to increase the value beyond 2048.) To disable automatic `cxfsd_max` calculation, set `cxfsd_max` to a non-zero value.
- Minimum: 16
- Maximum: 2048

Note: The value for `cxfsd_max` cannot be less than the value specified for `cxfsd_min`.

Location:

- IRIX: `/var/sysgen/mtune/cxfs`
- Linux: `fs.cxfs` (`sgi-cxfs` module)

`cxfsd_min`

Specifies the minimum number of `cxfsd` threads to run per CXFS filesystem. The value of the `cxfsd_min` parameter at mount time remains in effect for a filesystem until it is unmounted.

Legal values:

- Default: 0, which specifies an automatically calculated value that will be 2 times the number of CPUs (the number of actual running `cxfsd` threads is dynamic), as long as it is in the range 16 through 2048. To disable automatic `cxfsd_min` calculation, set `cxfsd_min` to a non-zero value.

- Minimum: 16
- Maximum: 2048

Location:

- IRIX: `/var/sysgen/mtune/cxfs`
- Linux: `fs.cxfs` (`sgi-cxfs` module)

mtcp_mesg_validate

Enables checksumming. Normally, this is not needed and is only used if TCP data corruption is suspected.

Legal values:

- 0 performs no validation (default)
- 1 generates checksums, but does not perform validation
- 2 generates and validates checksums, warns (via a SYSLOG message) on validation failure
- 3 generates and validates checksums, warns and returns an error message on validation failure
- 4 generates and validates checksums, warns and panics on validation error

Location:

- IRIX: `/var/sysgen/mtune/cell`
- Linux: `kernel.cell` (`sgi-cell` module)

Restricted System Tunable Parameters

This section lists the CXFS system tunable parameters that are provided for debugging purposes.



Caution: You must not modify any of these parameters unless directed to do so by SGI support.

Restricted Static Parameters

Static parameters require a reboot to take affect. On IRIX, you must build and boot new kernels, which happens automatically during a normal boot process. On Linux, you must specify the parameter in `/etc/modprobe.conf.local`. For more information, see "Site-Changeable System Tunable Parameters" on page 348.

`cxfs_extents_block_size`

Specifies the size in kilobytes of the units to use for memory allocations for extent lists on CXFS filesystems. You should only change this parameter for debugging purposes.

Range of values:

- Default: 0 (page size of the platform)
- Minimum: 0
- Maximum: 256

Location:

- IRIX: `/var/sysgen/mtune/cxfs`
- Linux: `fs.cxfs` (`sgi-cxfs` module)

`cxfs_extents_delta_depth`

Specifies the number of changes to the extent list kept by the CXFS metadata server for generating extents deltas. You should only change it for for debugging purposes.

Range of values:

- Default: 5
- Minimum: 0
- Maximum: 32

Location:

- IRIX: `/var/sysgen/mtune/cxfs`
- Linux: `fs.cxfs` (`sgi-cxfs` module)

cxfs_shutdown_time

Specifies the time (in seconds) that other nodes will wait for the node to take media offline after they have recognized that it has lost quorum, if the node has neither fencing nor reset configured. SGI recommends a value of 50 (0.5 seconds).

Range of values:

- Default: 50
- Minimum: 0
- Maximum: 6000

Location:

- IRIX: `/var/sysgen/mtune/cell`
- Linux: `kernel.cell` (`sgi-cell` module)

mesg_delay_time

Specifies the amount of time to delay messages, in nsecs.

Range of values:

- Default: 0
- Minimum: 0
- Maximum: 1000000

Location:

- IRIX: `/var/sysgen/mtune/cell`
- Linux: `kernel.cell` (`sgi-cell` module)

mtcp_reserve_size

Sets the size of the TCP window in bytes.

Range of values:

- Default: 61440
- Minimum: 2048

- Maximum: 1073741824

Location:

- IRIX: `/var/sysgen/mtune/cell`
- Linux: `kernel.cell` (`sgi-cell` module)

Restricted Dynamic Parameters

Dynamic parameters take affect as soon as they are changed.

`cell_tkm_feature_disable`

Disables selected features of the token module by setting a flag bit to one of the following:

- 0x1 disables speculative token acquisition
- 0x2 disables token prefetching
- 0x4 uses multiple RPCs to obtain a token set if rank and class conflict

Range of values:

- Default: 0
- Minimum: 0
- Maximum: 0x7fff

Location:

- IRIX: `/var/sysgen/mtune/cell`
- Linux: `kernel.cell` (`sgi-cell` module)

Note: This parameter supersedes the following parameters:

- `cxfs_prefetch`, which enabled/disabled token obtain optimization
 - `cxfs_speculative_token`, which enabled/disabled speculative vnode token fetching
-

`cms_fence_timeout`

Specifies the number of seconds to wait for `clconfd` to acknowledge a fence request. If a non-zero value is set and the time-out expires, CXFS takes the action specified by the `cms_fence_timeout_action` parameter.

Range of values:

- Default: 0 (infinite wait)
- Minimum: 0
- Maximum: 10000

Location:

- IRIX: `/var/sysgen/mtune/cell`
- Linux: `kernel.cell` (`sgi-cell` module)

`cms_fence_timeout_action`

Specifies the action to be taken when `clconfd` does not acknowledge a fence request (determined by `cms_fence_timeout`). `cms_fence_timeout_action` may be set to one of the following.

Legal values:

- 0 proceeds as if the fence return an error. This causes the node waiting for the fence acknowledgement to forcibly withdraw from the cluster, equivalent to a forced CXFS shutdown that occurs when a node loses quorum (default). If `clconfd` is still present and functioning properly, it will then restart the kernel `cms` daemon and the node will attempt to rejoin the cluster.
- 1 proceeds as if the fence succeeded. This clears all pending fence requests and continues (that is, fakes acknowledgment). **CAUTION:** Setting this value is potentially dangerous.
- 2 panics the local node

Location:

- IRIX: `/var/sysgen/mtune/cell`
- Linux: `kernel.cell` (`sgi-cell` module)

cms_reset_error_override

Specifies whether or not to ignore reset errors. You should only set this value to 1 for testing purposes, and **never** on a production system.

Legal values:

- 0 does not ignore reset errors (default)
- 1 ignores reset errors

Location:

- IRIX: `/var/sysgen/mtune/cell`
- Linux: `kernel.cell` (`sgi-cell` module)

cms_reset_timeout

Specifies the number of seconds to wait for `clconfd` to acknowledge a reset request. If you specify a non-zero value and the time-out expires, CXFS takes the action specified by the `cms_reset_timeout_action` parameter.

Range of values:

- Default: 0 (infinite wait)
- Minimum: 0
- Maximum: 10000

Location:

- IRIX: `/var/sysgen/mtune/cell`
- Linux: `kernel.cell` (`sgi-cell` module)

cms_reset_timeout_action

Specifies the action to be taken when `clconfd` does not acknowledge a reset request (determined by `cms_reset_timeout`).

Legal values:

- 0 proceeds as if the reset returned an error. This causes the node waiting for the reset acknowledgement to forcibly withdraw from the cluster, equivalent to a

forced CXFS shutdown that occurs when a node loses quorum (default). If `clconfd` is still present and functioning properly, it will then restart the kernel `cms` daemon and the node will attempt to rejoin the cluster.

- 1 proceeds as if the reset succeeded. This clears all pending resets and continues (that is, fakes acknowledgment). **CAUTION:** Setting this value is potentially dangerous.
- 2 panics the local node

Location:

- IRIX: `/var/sysgen/mtune/cell`
- Linux: `kernel.cell` (`sgi-cell` module)

`cms_trace_enable`

Enables or disables `cms` tracing on a non-DEBUG kernel. and determines the number of trace entries allocated.

Range of values:

- Default: 0 (disables)
- Minimum: 0
- Maximum: 1048576

Location:

- IRIX: `/var/sysgen/mtune/cell`
- Linux: `kernel.cell` (`sgi-cell` module)

`cxfs_recovery_slowdown`

Slows down recovery by inserting delays (measured in ms).

Range of values:

- Default: 0
- Minimum: 0
- Maximum: 60000

Location:

- IRIX: `/var/sysgen/mtune/cell`
- Linux: `kernel.cell` (`sgi-cell` module)

`cxfs_recovery_timeout_panic`

Specifies the action taken when a node with stalled recovery is discovered.

Legal values:

- 0 shuts down a node with stalled recovery (default)
- 1 panics a node with stalled recovery

Location:

- IRIX: `/var/sysgen/mtune/cell`
- Linux: `kernel.cell` (`sgi-cell` module)

`cxfs_recovery_timeout_period`

Specifies the time in seconds between recovery timeout polls.

Range of values:

- Default: 60
- Minimum: 0 (disables recovery polls)
- Maximum: 3600

Location:

- IRIX: `/var/sysgen/mtune/cell`
- Linux: `kernel.cell` (`sgi-cell` module)

`cxfs_recovery_timeout_stalled`

Specifies the time in seconds after which a node whose status is not changing is considered to have a stalled recovery.

Range of values:

- Default: 600
- Minimum: 0 (disables timeout)
- Maximum: 3600

Location:

- IRIX: `/var/sysgen/mtune/cell`
- Linux: `kernel.cell` (`sgi-cell` module)

cxfs_recovery_timeout_start

Specifies the time in seconds following a recovery before the recovery timeout monitoring begins.

Range of values:

- Default: 60
- Minimum: 0
- Maximum: 3600

Location:

- IRIX: `/var/sysgen/mtune/cell`
- Linux: `kernel.cell` (`sgi-cell` module)

cxfs_token_fault_tolerant

Specifies whether to tolerate certain recoverable errors in the token subsystem. The least significant 4 bits are used in non-DEBUG kernels, the next 4 bits are used in DEBUG kernels (SGI internal only). In each group of 4 bits, the most significant bit determines whether the system will panic if an error condition is detected. The next bit determines whether part of the code path doing error detection and/or handling is enabled or disabled. The last 2 bits are interpreted as a debug level:

- 0 = No messages are printed
- 1 = Debug level 1
- 2 = Debug level 2

- 3 = Debug level 3

Figure 13-5 displays the interpretation of the bits.

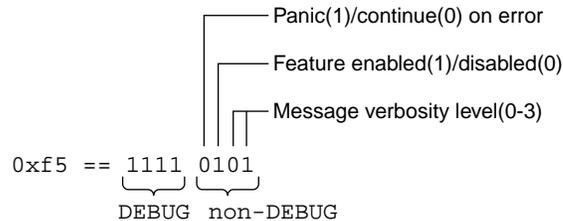


Figure 13-5 Value Bit Interpretation

Range of values:

- Default: 0xf5 (prints only some messages in the non-DEBUG case, and prints all messages and panics in the DEBUG case)
- Minimum: 0 (disables all messages and diagnostics, DEBUG and non-DEBUG)
- Maximum: 0xff (enables panics on error detection, maximum verbosity for diagnostic messages, DEBUG and non-DEBUG)

Location:

- IRIX: /var/sysgen/mtune/cxfs
- Linux: fs.cxfs (sgi-cxfs module)

cxfs_token_track

Specifies whether to verify that a client complies with the token locking hierarchy.

Range of values (for more information, see "cxfs_token_fault_tolerant" on page 364):

- Default: 0x0 (prints only some messages in the non-DEBUG case, and prints all messages and panics in the DEBUG case)
- Minimum: 0 (disables all messages and diagnostics, DEBUG and non-DEBUG)

- Maximum: 0xff (enables panics on error detection, maximum verbosity for diagnostic messages, DEBUG and non-DEBUG)

Location:

- IRIX: /var/sysgen/mtune/cxfs
- Linux: fs.cxfs (sgi-cxfs module)

cxfs_validate_objid

Specifies the following:

- Server: specifies whether to check that an objid received from a client corresponds to an object of the expected type
- Client: specifies the level of reporting upon receipt of an EBADOBJID error from the server.

Range of values (for more information, see "cxfs_token_fault_tolerant" on page 364):

- Default: 0xf5 (prints only some messages in the non-DEBUG case, and prints all messages and panics in the DEBUG case)
- Minimum: 0 (disables all messages and diagnostics, DEBUG and non-DEBUG)
- Maximum: 0xff (enables panics on error detection, maximum verbosity for diagnostic messages, DEBUG and non-DEBUG)

Location:

- IRIX: /var/sysgen/mtune/cxfs
- Linux: fs.cxfs (sgi-cxfs module)

cxfs_verify_existence_token

Specifies whether or not to verify that a client has the existence token before trying to obtain additional tokens.

Range of values (for more information, see "cxfs_token_fault_tolerant" on page 364):

- Default: `0xf5` (prints only some messages in the non-DEBUG case, and prints all messages and panics in the DEBUG case)
- Maximum: `0` (disables all messages and diagnostics, DEBUG and non-DEBUG)
- Minimum: `0xff` (enables panics on error detection, maximum verbosity for diagnostic messages, DEBUG and non-DEBUG)

Location:

- IRIX: `/var/sysgen/mtune/cxfs`
- Linux: `fs.cxfs` (`sgi-cxfs` module)

`cxfsd_sync_force`

Specifies a bitmask indicating `cxfsd` tasks that must be run synchronously, as opposed to the normal case where they are run asynchronously by threads from a `cxfsd` thread pool. The bits correspond to the opcodes in `cxfsd.h` (SGI internal source file).

Range of values:

- Default: `0`
- Minimum: `0`
- Maximum: `0x7fffffff`

Location:

- IRIX: `/var/sysgen/mtune/cxfs`
- Linux: `fs.cxfs` (`sgi-cxfs` module)

Hardware Changes and I/O Fencing

If you use I/O fencing and then make changes to your hardware configuration, you must verify that switch ports are properly enabled so that they can discover the WWPN of the HBA for I/O fencing purposes.

You must check the status of the switch ports involved whenever any of the following occur:

- An HBA is replaced on a node
- A new node is plugged into the switch for the first time
- A Fibre Channel cable rearrangement occurs

Note: The affected nodes should be shutdown before rearranging cables.

To check the status, use the following command on a CXFS administration node:

```
hafence -v
```

If any of the affected ports are found to be disabled, you must manually enable them before starting CXFS on the affected nodes:

1. Connect to the switch using `telnet`.
2. Use the `portenable` command to enable the port.
3. Close the `telnet` session.

After the port is enabled, the metadata server will be able to discover the new (or changed) WWPN of the HBA connected to that port and thus correctly update the switch configuration entries in the cluster database.

Configuring Private Network Failover

This section provides an example of modifying a cluster to provide private network failover by using the `cxfs_admin` command.

Do the following:

1. Create the failover network subnets. For example:

```
cxfs_admin:mycluster> create failover_net network=192.168.0.0 mask=255.255.255.0
cxfs_admin:mycluster> create failover_net network=192.168.1.0 mask=255.255.255.0
```

2. Disable all nodes (which shuts down the cluster):

```
cxfs_admin:mycluster> disable *
```

3. Update each node to include a private network. For example:

```
cxfs_admin:mycluster> modify red private_net=192.168.0.1,192.168.1.1
cxfs_admin:mycluster> modify yellow private_net=192.168.0.2,192.168.1.2
```

4. Enable all nodes:

```
cxfs_admin:mycluster> enable *
```

For more information, see Chapter 12, "Reference to `cxfs_admin` Tasks" on page 255.

Removing and Restoring Cluster Members

This section discusses removing and restoring cluster members for maintenance:

- "Removing a Metadata Server from the Cluster" on page 369
- "Restoring a Metadata Server to the Cluster" on page 371
- "Removing a Single Client-Only Node from the Cluster" on page 371
- "Restoring a Single Client-Only Node to the Cluster" on page 373
- "Stopping CXFS for the Entire Cluster" on page 374
- "Restarting the Entire Cluster" on page 374

These procedures are the absolute safest way to perform these tasks but in some cases are not the most efficient. They should be followed if you have having problems using standard operating procedures (performing a stop/start of CXFS services or a simple host shutdown or reboot).

Removing a Metadata Server from the Cluster

If you have a cluster with multiple metadata servers and you must perform maintenance on one of them, you must stop CXFS services on it.

To remove a metadata server (`exMDS` for example) from the cluster, do the following:

1. Enable relocation by using the `cxfs_relocation_ok` system tunable parameter. See "Relocation" on page 20.

2. For each filesystem for which exMDS is the active metadata server, manually relocate the metadata services from exMDS to a backup metadata server by using the CXFS GUI or `cxfs_admin`. For example:

```
cxfs_admin:mycluster> relocate fs1 server=node2
```

3. Disable relocation. See "Relocation" on page 20.

Note: If you do not perform steps 1–3 in a system reset configuration, exMDS will be reset shortly after losing its membership. The machine will also be configured to reboot automatically instead of stopping in the PROM. This means that you must watch the console and intervene manually to prevent a full reboot.

In a fencing configuration, exMDS will lose access to the SAN when it is removed from the cluster membership

4. Stop the CXFS services for the exMDS node to be shut down using the CXFS GUI or `cxfs_admin` running on another metadata server. For example:

```
cxfs_admin:mycluster> disable exMDS
```

5. Shut down exMDS.

If you do not want the cluster administration daemons and the CXFS control daemon to run during maintenance, execute the following commands:

- IRIX:

```
irix-exMDS# chkconfig cxfs_cluster off  
irix-exMDS# chkconfig cluster off
```

- SGI ProPack:

```
[root@linux-exMDS root]# chkconfig cxfs off  
[root@linux-exMDS root]# chkconfig cxfs_cluster off
```

If you do an upgrade of the cluster software, these arguments will be automatically reset to on and the cluster administration daemons and the CXFS control daemon will be started.

For more information, see "CXFS `chkconfig` Arguments" on page 320.

Restoring a Metadata Server to the Cluster

To restore a metadata server to the cluster, do the following:

1. Allow the cluster administration daemons and CXFS control daemon to be started upon reboot:

- IRIX:

```
irix-exMDS# chkconfig cxfs_cluster on
irix-exMDS# chkconfig cluster on
```

- SGI ProPack:

```
[root@linux-exMDS root]# chkconfig cxfs on
[root@linux-exMDS root]# chkconfig cxfs_cluster on
```

2. Immediately start cluster administration daemons on the node:

- IRIX:

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

- SGI ProPack:

```
exMDS# /etc/init.d/cxfs_cluster start
```

3. Immediately start the CXFS control daemon on the node:

```
exMDS# /etc/init.d/cxfs start
```

4. Start CXFS services on this node from another CXFS administration node:

```
anotherAdmin# start cx_services on node exMDS for cluster clustername force
```

Removing a Single Client-Only Node from the Cluster

To remove a single client-only node from the cluster, do the following:

1. Verify that the configuration is consistent by running the following on each active metadata server and comparing the output:

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

If the client is not consistent with the metadata servers, or if the metadata servers are not consistent, then you should abort this procedure and address the health of the cluster. If a client is removed while the cluster is unstable, attempts to get the client to rejoin the cluster are likely to fail.

2. Flush the system buffers on the client you want to remove in order to minimize the amount of buffered information that may be lost:

```
client# sync
```

3. Stop CXFS services on the client:

```
client# /etc/init.d/cxfs_client stop
client# chkconfig cxfs_client off
```

Note: The path to `cxfs_client` varies across the operating systems supported by CXFS. For more information, see *CXFS MultiOS Client-Only Guide for SGI InfiniteStorage*.

4. Verify that CXFS services have stopped:

- Verify that the CXFS client daemon is not running on the client (success means no output):

```
client# ps -ef | grep cxfs_client
client#
```

- Monitor the `cxfs_client` log on the client you wish to remove and look for filesystems that are unmounting successfully. For example:

```
Apr 18 13:00:06 cxfs_client: cis_setup_fsес Unmounted green0: green0 from /cxfs/green0
```

- Monitor the `SYSLOG` on the active metadata server and look for membership delivery messages that do not contain the removed client. For example, the following message indicates that cell 2 (`client`), the node being shut down, is not included in the membership:

```
Apr 18 13:01:03 5A:o200a unix: NOTICE: Cell 2 (client) left the membership
Apr 18 13:01:03 5A:o200a unix: NOTICE: Membership delivered for cells 0x3
Apr 18 13:01:03 5A:o200a unix: Cell(age): 0(7) 1(5)
```

- Use the following command to show that filesystems are not mounted:

```
client# df -hl
```

5. Verify that the configuration is consistent and does not contain the removed client by running the following on each active metadata server and comparing the output:

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

Restoring a Single Client-Only Node to the Cluster

To restore a single client-only node to the cluster, do the following:

1. Verify that the configuration is consistent by running the following on each active metadata server and comparing the output:

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

2. Start CXFS on the client-only node:

```
client# chkconfig cxfsc_client on
client# /etc/init.d/cxfsc_client start
```

Note: The path to `cxfsc_client` varies across the operating systems supported by CXFS. For more information, see *CXFS MultiOS Client-Only Guide for SGI InfiniteStorage*.

3. Verify that CXFS has started:

- Verify that the CXFS client daemon is running on the client-only node:

```
client# ps -ef | grep cxfsc_client
root      716      1  0 12:59:14 ?        0:05 /usr/cluster/bin/cxfsc_client
```

- Monitor the `SYSLOG` on the active metadata server and look for a cell discovery message for the client and a membership delivered message containing the client cell. For example (line breaks added for readability):

```
Apr 18 13:07:21 4A:o200a unix: WARNING: Discovered cell 2 (woody)
[priority 1 at 128.162.240.41 via 128.162.240.34]
Apr 18 13:07:31 5A:o200a unix: NOTICE: Cell 2 (client) joined the membership
Apr 18 13:07:31 5A:o200a unix: NOTICE: Membership delivered for cells 0x7
Apr 18 13:07:31 5A:o200a unix: Cell(age): 0(9) 1(7) 2(1)
```

- Monitor the `cxfsc_client` log on the client you and look for filesystem mounts that are processing successfully. For example:

```
Apr 18 13:06:56 cxfsc_client: cis_setup_fsesc Mounted green0: green0 on /cxfsc/green0
```

- Use the following command to show that filesystems are mounted:

```
client# df -hl
```

4. Verify that the configuration is consistent and contains the client by running the following on each active metadata server and comparing the output:

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

Stopping CXFS for the Entire Cluster

To stop CXFS for the entire cluster, do the following:

1. Stop CXFS services on a client-only node:

```
client# /etc/init.d/cxfs_client stop
```

Repeat this step on each client-only node.

2. Stop CXFS services on a metadata server:

```
MDS# /etc/init.d/cxfs stop
```

Repeat this step on each potential metadata server.

3. Stop the cluster daemons on a metadata server:

- IRIX:

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

- SGI ProPack:

```
[root@linuxMDS root]# /etc/init.d/cxfs_cluster stop
```

Repeat this step on each potential metadata server.

Restarting the Entire Cluster

To restart the entire cluster, do the following:

1. Start the cluster daemons on a potential metadata server:

- IRIX:

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

- SGI ProPack:

```
/etc/init.d/cxfs_cluster start
```

Repeat this step on each potential metadata server.

2. Start CXFS services on a metadata server:

```
MDS# /etc/init.d/cxfs start
```

Repeat this step on each potential metadata server.

3. Start CXFS services on a client-only node:

```
client# /etc/init.d/cxfs_client start
```

Repeat this step on each client-only node.

Discovering the WWNs

The `cxfs-enumerate-wwns` script enumerates the worldwide names (WWNs) on the host that are known to CXFS. For example, from an IRIX metadata server with two single-port HBAs:

```
irix# /var/cluster/clconfd-scripts/cxfs-enumerate-wwns
# cxfs-enumerate-wwns
# scsi @ /hw/scsi_ctlr/0/bus
# scsi @ /hw/scsi_ctlr/1/bus
# scsi @ /hw/scsi_ctlr/2/bus
# scsi @ /hw/scsi_ctlr/3/bus
210000e08b12ba14
# scsi @ /hw/scsi_ctlr/4/bus
210100e08b32ba14
```

Mapping XVM Volumes to Storage Targets

This section discusses mapping XVM volumes to storage targets on the IRIX and SGI ProPack platforms.

Mapping XVM Volumes to Storage Targets on IRIX

Do the following:

1. Get visible controller port WWNs.

2. Display the desired fields:

```
ls -d -l /dev/dsk/* | egrep -v "dks|root|swap" | cut -f4 -d"/" | sort -u
```

Mapping XVM Volumes to Storage Targets on SGI ProPack

You can use the `cxfs-enumerate-wwns` script to map XVM volumes to storage targets (assuming that CXFS software is installed.)

Do the following:

- SGI ProPack server-capable nodes:

```
# /var/cluster/clconfd-scripts/cxfs-enumerate-wwns | grep -v "#" | sort -u
```

- SGI ProPack client-only nodes:

```
# /var/cluster/cxfs_client-scripts/cxfs-enumerate-wwns | grep -v "#" | sort -u
```

Switching Between SGIRDAC and SGI AVT Mode for SGI RAID

To convert from SGIRDAC to SGI AVT, do the following:

1. Install the latest supported firmware on the RAID.
2. Determine the IP address for one of the controllers on each RAID box.
3. Make a script `settype.scr` that contains the following line:

```
set storageArray defaultHostType="modename";
```

Note: The capitalization and punctuation in the above line are required.

To switch to SGI AVT mode, use the following line:

```
set storageArray defaultHostType="SGIAVT";
```

To switch to SGIRDAC mode, use the following line:

```
set storageArray defaultHostType="SGIRDAC";
```

For the InfiniteStorage 220, use the CLI client to set the host type to SGI AVT:

```
smicli -w SA_WWID -c 'set storageArray defaultHostType="SGIAVT";'
```

To determine the value for *SA_WWID*, invoke the following:

```
smicli -d -w
```

For example:

```
# smicli -d -w
unit1 600a0b80002459d40000000045003fbc      localhost
      |---> SA_WWID
```

4. Run the following for one of the controllers per RAID box:

```
/opt/tpssm/client/tpssmcli RAID_IPaddress -f settype.scr
```

For example:

```
# /opt/tpssm/client/tpssmcli 192.168.0.1 -f settype.scr
Performing syntax check...

Syntax check complete.

Executing script...

Script execution complete.

tpssmcli completed successfully.
```


Cluster Database Management

This chapter contains the following:

- "Performing Cluster Database Backup and Restoration" on page 379
- "Checking the Cluster Configuration with `cxfs-config`" on page 382

Performing Cluster Database Backup and Restoration

You should perform a database backup whenever you want to save the database and be able to restore it to the current state at a later point.

You can use the following methods to restore the database:

- If the database is accidentally deleted from a node, use the `fs2d` daemon to replicate the database from another node in the pool.
- If you want to be able to recreate the current configuration, use the `config` command in `cxfs_admin`. You can then recreate this configuration by using the output file and the `cxfs_admin -f` option or running the script generated.
- If you want to retain a copy of the database and all node-specific information such as local logging, use the `cdbBackup` and `cdbRestore` commands. You should periodically backup the cluster database on all administration nodes using the `cdbBackup` command either manually or by adding an entry to the `root crontab` file.

Restoring a Deleted Database from Another Node

If the database has been accidentally deleted from an individual administration node, you can restore it by synchronizing with the database on another administration node.

Note: Do not use this method if the cluster database has been corrupted, because the database on another node will also be corrupted. In the case of corruption, you must reinstate a backup copy of the database. See "Using `cxfs_admin` to Recreate the Cluster Database" on page 381.

Do the following:

1. Stop the CXFS service on the node with the deleted database by running the following command on one node:



Caution: If you omit this step, the target node might be reset by another node.

```
cmgr> stop cx_services on node nodename for cluster clustername
```

2. Stop the CXFS control daemon (`clconfd`) by running the following command on each administration node:



Caution: Running this command will completely shut down all CXFS filesystem access on the current node.

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

3. Stop the CXFS cluster administration daemons (`cad`, `cmond`, `crsd`, and `fs2d`) by running the following command on each administration node:

- IRIX:

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

- SGI ProPack:

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

4. Run `cdbreinit` on administration nodes that are missing the cluster database.
5. Wait for the following message to be logged to the `syslog`:

```
fs2d[PID]: Finished receiving CDB sync series from machine nodename
```

6. Restart the CXFS control daemon by running the following command on each administration node:

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

The `cdbreinit` command will restart cluster daemons automatically. The `fs2d` daemon will then replicate the cluster database to those nodes from which it is missing

Using `cxfs_admin` to Recreate the Cluster Database

See "Saving and Recreating the Current Configuration with `cxfs_admin`" on page 300

Using `cdbBackup` and `cdbRestore` for the Cluster Database and Logging Information

The `cdbBackup` and `cdbRestore` commands backup and restore the cluster database and node-specific information, such as local logging information. You must run these commands individually for each node.

To perform a backup of the cluster, use the `cdbBackup` command on each node.



Caution: Do not make configuration changes while you are using the `cdbBackup` command.

To perform a restore, run the `cdbRestore` command on each node. You can use this method for either a missing or a corrupted cluster database. Do the following:

1. Stop CXFS services on all nodes in the cluster.
2. Stop the cluster administration daemons on each administration node.
3. Remove the old database by using the `cdbreinit` command on each node.
4. Stop the cluster administration daemons again (these were restarted automatically by `cdbreinit` in the previous step) on each node.
5. Use the `cdbRestore` command on each node.
6. Start the cluster administration daemons on each node.

For example, to backup the current database, clear the database, and then restore the database to all administration nodes, do the following as directed on administration nodes in the cluster:

On each node:

```
# /usr/cluster/bin/cdbBackup
```

On one node:

```
cmgr> stop cx_services for cluster clusterA
```

On each node:

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

On each node:

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

```
SGI ProPack:  
# /etc/init.d/cxfs_cluster stop
```

On each:

```
# /usr/cluster/bin/cdbreinit
```

On each node (again):

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

```
SGI ProPack:  
# /etc/init.d/cxfs_cluster stop
```

On each node:

```
# /usr/cluster/bin/cdbRestore
```

On each node:

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

```
SGI ProPack:  
# /etc/init.d/cxfs_cluster start
```

For more information, see the `cdbBackup` and `cdbRestore` man page.

Checking the Cluster Configuration with `cxfs-config`

The `cxfs-config` command displays and checks configuration information in the cluster database. You can run it on any administration node in the cluster.

By default, `cxfs-config` displays the following:

- Cluster name and cluster ID
- Tiebreaker node

- Networks for CXFS failover networks
- Nodes in the pool:
 - Node ID
 - Cell ID (as assigned by the kernel when added to the cluster and stored in the cluster database)
 - Status of CXFS services (configured to be enabled or disabled)
 - Operating system
 - Node function
- CXFS filesystems:
 - Name, mount point (enabled means that the filesystem is configured to be mounted; if it is not mounted, there is an error)
 - Device name
 - Mount options
 - Potential metadata servers
 - Nodes that should have the filesystem mounted (if there are no errors)
 - Switches:
 - Switch name, user name to use when sending a telnet message, mask (a hexadecimal string representing a 64-bit port bitmap that indicates the list of ports in the switch that will not be fenced)
 - Ports on the switch that have a client configured for fencing at the other end
- Warnings or errors

For example:

```
thump# /usr/cluster/bin/cxfs-config
Global:
  cluster: topiary (id 1)
  tiebreaker: <none>

Networks:
  net 0: type tcpip 192.168.0.0      255.255.255.0
```

14: Cluster Database Management

```
net 1: type tcpip 134.14.54.0      255.255.255.0
```

Machines:

```
node leesa: node 6      cell 2  enabled  Linux32  client_only
fail policy: Fence
nic 0: address: 192.168.0.164 priority: 1 network: 0
nic 1: address: 134.14.54.164 priority: 2 network: 1
```

```
node thud: node 8      cell 1  enabled  IRIX      client_admin
fail policy: Fence
nic 0: address: 192.168.0.204 priority: 1 network: 0
nic 1: address: 134.14.54.204 priority: 2 network: 1
```

```
node thump: node 1     cell 0  enabled  IRIX      server_admin
fail policy: Fence
nic 0: address: 192.168.0.186 priority: 1 network: 0
nic 1: address: 134.14.54.186 priority: 2 network: 1
```

Filesystems:

```
fs dxm: /mnt/dxm          enabled
device = /dev/cxvm/tp9500a4s0
options = []
servers = thump (1)
clients = leesa, thud, thump
```

Switches:

```
switch 0: admin@asg-fcsw1      mask 0000000000000000
port 8: 210000e08b0ead8c thump
port 12: 210000e08b081f23 thud
```

```
switch 1: admin@asg-fcsw0      mask 0000000000000000
```

Warnings/errors:

```
enabled machine leesa has fencing enabled but is not present in switch database
```

The following options are of particular interest:

- -all lists all available information

- `-ping` contacts each NIC in the machine list and displays if the packets is transmitted and received. For example:

```
node leesa: node 6      cell 2  enabled  Linux32  client_only
fail policy: Fence
nic 0: address: 192.168.0.164 priority: 1
      ping: 5 packets transmitted, 5 packets received, 0.0% packet loss
      ping: round-trip min/avg/max = 0.477/0.666/1.375 ms
nic 1: address: 134.14.54.164 priority: 2
      ping: 5 packets transmitted, 5 packets received, 0.0% packet loss
      ping: round-trip min/avg/max = 0.469/0.645/1.313 ms
```

- `-xfs` lists XFS information for each CXFS filesystem, such as size. For example:

```
Filesystems:
fs dxm: /mnt/dxm          enabled
      device = /dev/cxvm/tp9500a4s0
      options = []
      servers = thump (1)
      clients = leesa, thud, thump
      xfs:
        magic: 0x58465342
        blocksize: 4096
        uuid: 3459ee2e-76c9-1027-8068-0800690dac3c
        data size 17.00 Gb
```

- `-xvm` lists XVM information for each CXFS filesystem, such as volume size and topology. For example:

```
Filesystems:
fs dxm: /mnt/dxm          enabled
      device = /dev/cxvm/tp9500a4s0
      options = []
      servers = thump (1)
      clients = leesa, thud, thump
      xvm:
        vol/tp9500a4s0          0 online,open
          subvol/tp9500a4s0/data 35650048 online,open
            slice/tp9500a4s0    35650048 online,open

      data size: 17.00 Gb
```

- -check performs extra verification, such as XFS filesystem size with XVM volume size for each CXFS filesystem. This option may take a few moments to execute.

The following example shows errors reported by `cxfs-config`:

```
aiden # /usr/cluster/bin/cxfs-config -check -all
Global:
  cluster: BP (id 555)
  cluster state: enabled
  tiebreaker:
Networks:
  net 0: type tcpip 10.11.0.0          255.255.255.0
  net 1: type tcpip 128.162.242.0     255.255.255.0

Machines:
  node aiden: node 27560 cell 0  enabled  IRIX    server_admin
    hostname: aiden.example.com
    fail policy: Fence, Shutdown
    nic 0: address: 10.11.0.241 priority: 1 network: 0
      ping: 5 packets transmitted, 5 packets received, 0.0% packet loss
      ping: round-trip min/avg/max = 0.136/0.171/0.299 ms
    nic 1: address: 128.162.242.12 priority: 2 network: 1
      ping: 5 packets transmitted, 5 packets received, 0.0% packet loss
      ping: round-trip min/avg/max = 0.130/0.171/0.303 ms

  node brigid: node 31867 cell 2  enabled  IRIX    server_admin
    hostname: brigid.example.com
    fail policy: Fence, Shutdown
    nic 0: address: 10.11.0.240 priority: 1 network: 0
      ping: 5 packets transmitted, 5 packets received, 0.0% packet loss
      ping: round-trip min/avg/max = 0.303/0.339/0.446 ms
    nic 1: address: 128.162.242.11 priority: 2 network: 1
      ping: 5 packets transmitted, 5 packets received, 0.0% packet loss
      ping: round-trip min/avg/max = 0.336/0.430/0.799 ms

  node flynn: node 1      cell 1  enabled  linux64 client_only
    hostname: flynn.example.com
    fail policy: Fence, Shutdown
    nic 0: address: 10.11.0.234 priority: 1 network: 0
      ping: 5 packets transmitted, 5 packets received, 0.0% packet loss
      ping: round-trip min/avg/max = 0.323/0.370/0.539 ms
    nic 1: address: 128.162.242.189 priority: 2 network: 1
```

```
ping: 5 packets transmitted, 5 packets received, 0.0% packet loss
ping: round-trip min/avg/max = 0.283/0.312/0.424 ms
```

Filesystems:

```
fs concatfs: /concatfs          enabled
  device = /dev/cxvm/concatfs
  force = true
  options = [rw,quota]
  servers = aiden (1), brigid (2)
  clients = aiden, brigid, flynn
  xvm:
    vol/concatfs                0 online,open
      subvol/concatfs/data      2836134016 online,open
      subvol/concatfs/data      2836134016 online,open
        concat/concat0         2836134016 online,tempname,open
          slice/lun2s0          1418067008 online,open
          slice/lun3s0          1418067008 online,open
```

```
data size: 1.32 TB
```

```
xfs:
  magic: 0x58465342
  blocksize: 4096
  uuid: 9616ae39-3a50-1029-8896-080069056bf5
  data size 1.32 TB
```

```
fs stripefs: /stripefs          enabled
  device = /dev/cxvm/stripefs
  force = true
  options = [rw,quota]
  servers = aiden (1), brigid (2)
  clients = aiden, brigid, flynn
  xvm:
    vol/stripefs                0 online,open
      subvol/stripefs/data      2836133888 online,open
      stripe/stripe0            2836133888 online,tempname,open
        slice/lun0s0            1418067008 online,open
        slice/lun1s0            1418067008 online,open
```

```
data size: 1.32 TB
```

```
xfs:
  magic: 0x58465342
```

14: Cluster Database Management

```
blocksize: 4096
uuid: 9616ae38-3a50-1029-8896-080069056bf5
data size 1.32 TB
```

Switches:

```
switch 0: 32 port brocade admin@fcswitch12      port 12: 210000e08b00e6eb brigid
           port 28: 210000e08b041a3a aiden

switch 1: 32 port brocade admin@fcswitch13      port 7: 210000e08b08793f flynn
           port 12: 210100e08b28793f flynn
```

cxfs-config warnings/errors:

```
server aiden fail policy must not contain "Shutdown" for cluster with
even number of enabled servers and no tiebreaker
server brigid fail policy must not contain "Shutdown" for cluster with
even number of enabled servers and no tiebreaker
```

For a complete list of options, see the `cxfs-config` man page.

Coexecution with FailSafe

This chapter discusses the following:

- "Why Run CXFS and FailSafe Together?" on page 389
- "Coexecution Release Levels" on page 390
- "Size of the Coexecution Cluster" on page 390
- "Cluster Type" on page 390
- "Metadata Server Node Types" on page 392
- "Separate GUIs" on page 392
- "Conversion" on page 392
- "Network Interfaces" on page 393
- "Metadata Servers and Failover Domain" on page 393
- "CXFS Resource Type for FailSafe" on page 393
- "Core-Dump Completion" on page 395

Also see "Communication Paths in a Coexecution Cluster" on page 503.

Why Run CXFS and FailSafe Together?

CXFS allows groups of computers to coherently share large amounts of data while maintaining high performance.

The SGI FailSafe product provides a general facility for providing highly available (HA) services. If one of the administration nodes in the cluster or one of the node's components fails, a different administration node in the cluster restarts the HA services of the failed node. To CXFS clients, the services on the replacement node are indistinguishable from the original services before failure occurred. It appears as if the original node has crashed and rebooted quickly. The CXFS clients notice only a brief interruption in the HA service.

FailSafe assumes that CXFS filesystems are highly available and will recover from CXFS failures (including loss of CXFS membership). FailSafe will wait for CXFS to

recover CXFS filesystems before the resource group containing the CXFS resource is started on another FailSafe node in the cluster.

You can therefore use FailSafe in a CXFS cluster (known as *coexecution*) to provide HA services (such as NFS or web) running on a CXFS filesystem. This combination provides high-performance shared data access for highly available applications in a clustered system.

Coexecution Release Levels

CXFS 6.5.10 or later and IRIS FailSafe 2.1 or later (plus relevant patches) may be installed and run on the same system.

Size of the Coexecution Cluster

A subset of administration nodes in a coexecution cluster can be configured to be used as FailSafe nodes; a coexecution cluster can have up to eight nodes that run FailSafe.

All nodes in a CXFS cluster will run CXFS, and up to eight of those administration nodes can also run FailSafe. All administration nodes must run IRIX (FailSafe is not supported on SGI ProPack). Even when you are running CXFS and FailSafe, there is still only one pool, one cluster, and one cluster configuration.

It is recommended that a production cluster be configured with an odd number of server-capable nodes. (A cluster with reset cables and only two server-capable nodes is supported, but there are inherent issues with this configuration; see "CXFS Recovery Issues in a Cluster with Only Two Server-Capable Nodes " on page 526.)

Cluster Type

The cluster can be one of three types:

- **FailSafe.** In this case, all nodes will also be of type `FailSafe`. The nodes must all be administration nodes.
- **CXFS.** In this case, all nodes will be of type `CXFS`. The nodes can be either administration nodes or client-only nodes.

- CXFS and FailSafe (coexecution). In this case, all nodes will be a mix of type CXFS (any nodes running other operating systems) and type CXFS and FailSafe (administration nodes), using FailSafe for application-level high availability and CXFS.

Note: Although it is possible to configure a coexecution cluster with type FailSafe only nodes, SGI does not support this configuration.

Figure 15-1 describes some of the various legal and illegal combinations.

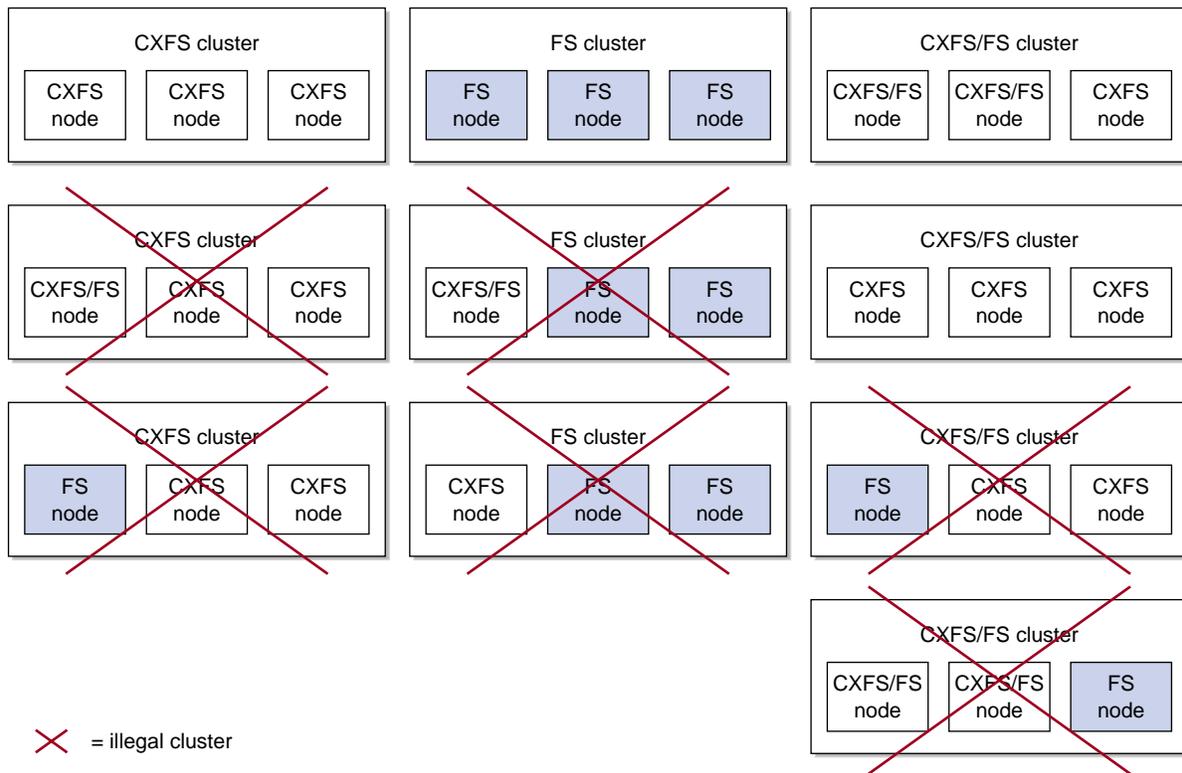


Figure 15-1 Cluster and Node Type Combinations

Metadata Server Node Types

All potential metadata server nodes must be of one of the following types:

- CXFS
- CXFS and FailSafe

Separate GUIs

There is one `cmgr` (`cluster_mgr`) command but separate graphical user interfaces (GUIs) for CXFS and for FailSafe. You must manage CXFS configuration with the CXFS GUI and FailSafe configuration with the FailSafe GUI; you can manage both with `cmgr`.

Conversion

Using the CXFS GUI or `cmgr`, you can convert an existing FailSafe cluster and nodes to type CXFS or to type CXFS and FailSafe. You can perform a parallel action using the FailSafe GUI. A converted node can be used by FailSafe to provide application-level high-availability and by CXFS to provide clustered filesystems. See "Set Up an Existing FailSafe Cluster for CXFS with the GUI" on page 205.

However:

- You cannot change the type of a node if the respective HA or CXFS services are active. You must first stop the services for the node.
- The cluster must support all of the functionalities (FailSafe and/or CXFS) that are turned on for its nodes; that is, if your cluster is of type CXFS, then you **cannot** modify a node that is already part of the cluster so that it is of type FailSafe. However, the nodes do not have to support all the functionalities of the cluster; that is, you can have a CXFS node in a CXFS and FailSafe cluster.

See "Convert a Node to CXFS or FailSafe with `cmgr`" on page 592, and "Convert a Cluster to CXFS or FailSafe with `cmgr`" on page 602.

Network Interfaces

For FailSafe, you must have at least two network interfaces. However, CXFS uses only one interface for **both** heartbeat and control messages. (The CXFS GUI appears to let you select only heartbeat or only control for a network, but you must not choose these selections.)

Metadata Servers and Failover Domain

The metadata server list must exactly match the failover domain list (the names and the order of names).

CXFS Resource Type for FailSafe

FailSafe provides a CXFS resource type that can be used to fail over applications that use CXFS filesystems. CXFS resources must be added to the resource group that contain the resources that depend on a CXFS filesystem. The CXFS resource type name is the CXFS filesystem mount point.

The CXFS resource type has the following characteristics:

- It does not start all resources that depend on CXFS filesystem until the CXFS filesystem is mounted on the local node.
- The `start` and `stop` action scripts for the CXFS resource type do not mount and unmount CXFS filesystems, respectively. (The `start` script waits for the CXFS filesystem to become available; the `stop` script does nothing but its existence is required by FailSafe.) Users should use the CXFS GUI or `cmgr` command to mount and unmount CXFS filesystems.
- It monitors CXFS filesystem for failures.
- Optionally, for applications that must run on a CXFS metadata server, the CXFS resource type relocates the CXFS metadata server when there is an application failover. In this case, the application failover domain (AFD) for the resource group should consist of the CXFS metadata server and the meta-data server backup nodes.

The CXFS filesystems that an NFS server exports should be mounted on all nodes in the failover domain using the CXFS GUI or the `cmgr` command.

For example, following are the commands used to create resources NFS, CXFS, and statd_unlimited based on a CXFS filesystem mounted on /FC/lun0_s6. (This example assumes that you have defined a cluster named test-cluster and have already created a failover policy named cxfs-fp and a resource group named cxfs-group based on this policy. Line breaks added for readability.)

```
cmgr> define resource /FC/lun0_s6 of resource_type CXFS in cluster test-cluster
Enter commands, when finished enter either "done" or "cancel"
```

Type specific attributes to create with set command:

Type Specific Attributes - 1: relocate-mds

No resource type dependencies to add

```
resource /FC/lun0_s6 ? set relocate-mds to false
resource /FC/lun0_s6 ? done
```

=====

```
cmgr> define resource /FC/lun0_s6 of resource_type NFS in cluster test-cluster
Enter commands, when finished enter either "done" or "cancel"
```

Type specific attributes to create with set command:

Type Specific Attributes - 1: export-info
Type Specific Attributes - 2: filesystem

No resource type dependencies to add

```
resource /FC/lun0_s6 ? set export-info to rw
resource /FC/lun0_s6 ? set filesystem to /FC/lun0_s6
resource /FC/lun0_s6 ? done
```

=====

```
cmgr> define resource /FC/lun0_s6/statmon of resource_type statd_unlimited in cluster
test-cluster
```

Enter commands, when finished enter either "done" or "cancel"

Type specific attributes to create with set command:

Type Specific Attributes - 1: ExportPoint

Resource type dependencies to add:

Resource Dependency Type - 1: NFS

```
resource /FC/lun0_s6/statmon ? set ExportPoint to /FC/lun0_s6
resource /FC/lun0_s6/statmon ? add dependency /FC/lun0_s6 of type NFS
resource /FC/lun0_s6/statmon ? done
```

```
=====
cmgr> define resource_group cxfs-group in cluster test-cluster
Enter commands, when finished enter either "done" or "cancel"

resource_group cxfs-group ? set failover_policy to cxfs-fp
resource_group cxfs-group ? add resource /FC/lun0_s6 of resource_type NFS
resource_group cxfs-group ? add resource /FC/lun0_s6 of resource_type CXFS
resource_group cxfs-group ? add resource /FC/lun0_s6/statmon of resource_type statd_unlimited
resource_group cxfs-group ? done
```

For more information about resource groups and failover domains, see the *FailSafe Administrator's Guide for SGI InfiniteStorage*.

Core-Dump Completion

When used in a FailSafe co-execution cluster, the system kernel core-dump may not complete before FailSafe timers expire and issue a reset of the system. This will result in an incomplete core-dump. To avoid this, set the IRIX `dump_level` system tunable parameter to the value 2 to shorten the time required to generate a dump.

Trusted IRIX and CXFS

CXFS has been qualified in an SGI Trusted IRIX cluster with the Data Migration Facility (DMF) and Tape Management Facility (TMF).

If you want to run CXFS and Trusted IRIX, all server-capable nodes in the cluster must run Trusted IRIX. The client-only nodes can run IRIX. Other platforms are not supported in a cluster with Trusted IRIX.

Installation Tips for CXFS and Trusted IRIX

SGI recommends that you install all of the software products you intend to run (Trusted IRIX, CXFS, DMF, TMF, and so on) at the same time.

After installing these products, you must do the following:

1. From the system console, go to the system maintenance menu. For example:

```
# init 0
```

(If your system is set to automatically reboot to multiuser mode, you will need to press Esc to reach the menu.)

2. Choose 5 from the menu in order to enter the command monitor:

```
System Maintenance Menu
```

```
1) Start System
2) Install System Software
3) Run Diagnostics
4) Recover System
5) Enter Command Monitor
```

```
Option? 5
```

3. Enter single user mode by using the single command:

```
>> single
```

4. Enter the root password when prompted.

5. Ensure that you are in the root directory:

```
# cd /
```

6. Set the following attributes for Trusted IRIX and CXFS:

```
# suattr -C all+eip
```

7. Execute the Trusted IRIX configuration command, which sets the appropriate extended attributes on files:

```
# /etc/trix.config
```

For more information, see:

- *Trusted IRIX Read Me First Notice*
- *Trusted IRIX/CMW Security Features User's Guide*

Mandatory Access Controls

In a mixed Trusted IRIX and IRIX cluster, an IRIX CXFS client will require but not have a mandatory access control (MAC) label associated with its credentials when it attempts to access a Trusted IRIX server. In order to address this, a MAC label is provided in one of the following ways:

- The filesystem can be mounted with the `eag:mac-ip=label` option to specify the label used for IRIX CXFS clients.
- If the mount option is not used, the default label in the `rhost` database entry for the IRIX original node is used.
- If the `rhost` database entry is unavailable or invalid, the following label is used: `msenlow, minthigh`.

Monitoring Status

You can view the system status in the following ways:

Note: Administrative tasks must be performed using one of the following tools:

- The CXFS GUI when it is connected to a CXFS administration node (a node that has the `cluster_admin` software package installed)
- `cxfs_admin` (you must logged in as `root` on a host that has permission to access the CXFS cluster database)

You must run administration commands on a CXFS administration node; you run the `cxfs_info` status command on a client-only node.

- Monitor log files. See "Status in Log Files" on page 400
- Use the GUI or the `tail` command to view the end of the system log file:
 - IRIX: `/var/adm/SYSLOG`
 - SGI ProPack: `/var/log/messages`
- Keep continuous watch on the state of a cluster using the GUI view area, the `clconf_info -e` command, or the following `cxfs_admin` command:

```
cxfs_admin -i clustername -r -c "status interval=seconds"
```
- Query the status of an individual node or cluster using the CXFS GUI or `cxfs_admin`.
- Manually test the filesystems with the `ls` command.
- Monitor the system with Performance Co-Pilot. You can use Performance Co-Pilot to monitor the read/write throughput and I/O load distribution across all disks and for all nodes in the cluster. The activity can be visualized, used to generate alarms, or archived for later analysis. You can also monitor XVM statistics. See the *Performance Co-Pilot for IA-64 Linux User's and Administrator's Guide*, *Performance Co-Pilot for IRIX Advanced User's and Administrator's Guide*, the *Performance Co-Pilot Programmer's Guide*, and the `dkvis`, `pmie`, `pmieconf`, and `pmlogger` man pages.

Note: You must manually install the XVM statistics for the Performance Co-Pilot package; it is not installed by default. See Chapter 6, "IRIX CXFS Installation" on page 101.

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

- "Status in Log Files"
- "Cluster Status" on page 402
- "Node Status" on page 406
- "XVM Statistics" on page 408
- "I/O Fencing Status" on page 409
- "Heartbeat Timeout Status" on page 411

Status in Log Files

You should monitor the following log files listed for problems:

- Administration node logs:
 - System log:
 - IRIX: `/var/adm/SYSLOG`
 - SGI ProPack: `/var/log/messages`
 - Events from the GUI and `clconfd`: `/var/cluster/ha/log/cad_log`
 - Kernel status: `/var/cluster/ha/log/clconfd_hostname`
 - Command line interface log: `/var/cluster/ha/log/cli_hostname`
 - Monitoring of other daemons: `/var/cluster/ha/log/cmond_log`
 - Reset daemon log: `/var/cluster/ha/log/crsd_hostname`

- Output of the diagnostic tools such as the serial and network connectivity tests:
/var/cluster/ha/log/diags_hostname
- Cluster database membership status: */var/cluster/ha/log/fs2d_log*
- System administration log, which contains a list of the commands run by the GUI:
 - IRIX: */var/sysadm/salog*
 - SGI ProPack: */var/lib/sysadm/salog*
- Client-only node log files:
 - *cxfs_client* log file:
 - IRIX: */var/adm/cxfs_client*
 - SGI ProPack: */var/log/cxfs_client*
 - System log:
 - IRIX: */var/adm/SYSLOG*
 - SGI ProPack: */var/log/messages*

Look for a Membership delivered message to indicate that a cluster was formed.

- Output of the diagnostic tools such as the serial and network connectivity tests:
/var/cluster/ha/log/diags_hostname
- The SGI ProPack platform uses the logrotate system utility to rotate the *cxfs_client* logs:
 - The */etc/logrotate.conf* file specifies how often system logs are rotated
 - The */etc/logrotate.d/cxfs_client* file specifies the manner in which *cxfs_client* logs are rotated

For information about client-only nodes running other operating systems, see *CXFS MultiOS Client-Only Guide for SGI InfiniteStorage*.

If the disk is filling with log messages, see "Log File Management" on page 339.



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

Cluster Status

You can monitor system status with the following tools:

- The `cxfs_admin` command on any host that has `monitor` access the CXFS cluster database (see "Setting `cxfs_admin` Access Permissions" on page 268)
- The CXFS GUI connected to a CXFS administration node
- The `clconf_info` command on a CXFS administration node
- The `cxfs_info` command on a client-only node

Also see "Key to Icons and States" on page 201.

Check Cluster Status with the GUI

The easiest way to keep a continuous watch on the state of a cluster is to use the view area and choose the following:

 Edit
 > Expand All

The cluster status can be one of the following:

- **ACTIVE**, which means the cluster is up and running.
- **INACTIVE**, which means the start CXFS 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 services are not running on the node performing the query. For more information, see in "Node Status" on page 406.

Check Cluster Status with `clconf_info`

If the cluster is up, you can see detailed information by using `/usr/cluster/bin/clconf_info` on a CXFS administration node.

The `clconf_info` command has the following options:

<code>-e</code>	Waits for events from <code>clconfd</code> and displays the new information
<code>-n <i>nodename</i></code>	Displays information for the specified logical node name
<code>-p</code>	Persists until the membership is formed
<code>-q</code>	(Quiet mode) Decreases verbosity of output. You can repeat this option to increase the level of quiet; that is, <code>-qq</code> specifies more quiet (less output) than <code>-q</code> .
<code>-s</code>	Sorts the output alphabetically by name for nodes and by device for filesystems. By default, the output is not sorted.
<code>-v</code>	(Verbose mode) Specifies the verbosity of output (<code>-vv</code> specifies more verbosity than <code>-v</code>). Deferred implementation.

For example:

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

Event at [2004-04-16 09:20:59]

Membership since Fri Apr 16 09:20:56 2004

Node           NodeID Status  Age    CellID
-----
leesa          0  inactive  -      0
whack          2  up        16     3
lustre         8  up        5      5
thud           88  up        16     1
cxfs2         102  DOWN      -      2

2 CXFS FileSystems
/dev/cxvm/tp9500_0 on /mnt/cxfs0  enabled  server=(whack)  2 client(s)=(thud,lustre)  status=UP
/dev/cxvm/tp9500a4s0 on /mnt/tp9500a4s0  disabled  server=()  0 client(s)=()  status=DOWN
```

This command displays the following fields:

- Node is the node name.
- NodeID is the node ID.
- Status is the status of the node, which may be up, DOWN, or inactive.
- Age indicates how many membership transitions in which the node has participated. The age is 1 the first time a node joins the membership and will increment for each time the membership changes. This number is dynamically allocated by the CXFS software (the user does not define the age).
- CellID is the cell ID, which is allocated when a node is added into the cluster definition with the GUI or `cxfs_admin`. It persists until the node is removed from the cluster. The kernel also reports the cell ID in console messages.

Check Cluster/Node/Filesystem/Switch Status with `cxfs_admin`

To query node and cluster status, use the following `cxfs_admin` command on any host that has monitor access (see "Setting `cxfs_admin` Access Permissions" on page 268) the CXFS cluster database:

```
status
```

To continuously redisplay an updated status, enter an interval in seconds:

```
status interval=seconds
```

To stop the updates, send an interrupt signal (usually Ctrl+C).

The most common states for nodes include:

- Disabled: The node is not allowed to join the cluster
- Inactive: The node is not in cluster membership
- Stable: The node is in membership and has mounted all of its filesystems

The most common states for filesystems include:

- Mounted: All enabled nodes have mounted the filesystem
- Unmounted: All nodes have unmounted the filesystem

Any other state (not mentioned above) requires attention by the administrator.

For example (a * character indicates a server-capable node):

```

cxfs_admin:mycluster> status
Cluster      : mycluster
Tiebreaker   : irix-client
Licenses     : enterprise   allocated 12 of 278
              workstation  allocated 4 of 15
-----
Node          Cell ID  Status
-----
mds1 *       6       Stable
mds2 *       0       Stable
aix-client   4       Stable
irix-client  1       Stable
mac-client   3       Inactive
solaris-client 2       Stable
windows-client 5       Stable
-----
Filesystem    Mount Point      Status
-----
concatfs     /mnt/concatfs    Mounted (mds1)
mirrorfs     /mnt/mirrorfs    Mounted (mds1)
stripefs     /mnt/stripefs    Mounted (mds1)
-----
Switch        Port Count  Known Fenced Ports
-----
fcswitch12   32         None
fcswitch13   32         None

```

Check Cluster/Node/Filesystem Status with `cxfs_info`

The `cxfs_info` command provides information about the cluster status, node status, and filesystem status. `cxfs_info` is run from a client-only node. The path to `cxfs_info` varies by platform.

You can use the `-e` option to display information continuously, updating the screen when new information is available; use the `-c` option to clear the screen between updates. For less verbose output, use the `-q` (quiet) option.

For example, on a Solaris node named `cxfssun4`:

```
cxfssun4 # /usr/cxfs_cluster/bin/cxfs_info
cxfs_client status [timestamp Sep 03 12:16:06 / generation 18879]

Cluster:
  sun4 (4) - enabled
Local:
  cxfssun4 (2) - enabled, state: stable, cms: up, xvm: up, fs: up
Nodes:
  cxfs27      enabled up    1
  cxfs28      enabled up    0
  cxfsnt4     enabled up    3
  cxfssun4    enabled up    2
  mesabi      enabled DOWN  4
Filesystems:
  lun1s0      enabled mounted   lun1s0      /lun1s0
  mirror0     disabled unmounted mirror0     /mirror0
```

Node Status

The following sections provide different methods to monitor node status. Also see "Check Cluster/Node/Filesystem Status with `cxfs_info`" on page 405 and "Check Cluster/Node/Filesystem/Switch Status with `cxfs_admin`" on page 404.

Monitoring Node Status with the GUI

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 services are started and the node is part of the CXFS kernel membership. For more information, see "CXFS Kernel Membership, Quorum, and Tiebreaker" on page 512.

- **DOWN**, which means that although CXFS services are started and the node is defined as part of the cluster, the node is not in the current CXFS kernel membership.
- **INACTIVE**, which means that the start CXFS services task has not been run.
- **UNKNOWN**, which means that the state cannot be determined because CXFS services are not running on the node performing the query.

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

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

- **INACTIVE** for its own state, because it can determine that the start CXFS 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

You can use the view area to monitor the status of the nodes. Select **View: Nodes and Cluster**.

To determine whether a node applies to CXFS, to FailSafe, or both, double-click the node name in the display.

Monitoring Node Status with `cxfs_admin`

See "Check Cluster/Node/Filesystem/Switch Status with `cxfs_admin`" on page 404.

Monitoring Node Status with `clconf_info`

You can use the `clconf_info` command to monitor the status of the nodes in the cluster. It uses the same node states as the CXFS GUI. See "Monitoring Node Status with the GUI" on page 406.

For example:

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

```
Event at [2004-04-16 09:20:59]
```

```
Membership since Fri Apr 16 09:20:56 2004
```

Node	NodeID	Status	Age	CellID
leesa	0	inactive	-	0
whack	2	up	16	3
lustre	8	up	5	5
thud	88	up	16	1
cxfs2	102	DOWN	-	2

```
2 CXFS FileSystems
```

```
/dev/cxvm/tp9500_0 on /mnt/cxfs0 enabled server=(whack) 2 client(s)=(thud,lustre) status=UP
```

```
/dev/cxvm/tp9500a4s0 on /mnt/tp9500a4s0 disabled server=() 0 client(s)=() status=DOWN
```

XVM Statistics

Note: This feature assumes that you have installed the `pcp_eoe` and `pcp_eoe.sw.xvm` packages; see Chapter 6, "IRIX CXFS Installation" on page 101.

You can use Performance Co-Pilot to monitor XVM statistics. To do this, you must enable the collection of statistics:

- To enable the collection of statistics for the local host, enter the following:

```
$ pmstore xvm.control.stats_on 1
```

- To disable the collection of statistics for the local host, enter the following:

```
$ pmstore xvm.control.stats_on 0
```

You can gather XVM statistics in the following ways:

- By using the `pmval` command from the IRIX `pcp_eoe.sw.monitor` package and the SGI ProPack bit `pcp` RPM. It can be used to produce an ASCII report of

selected metrics from the `xvm` group in the Performance Co-Pilot namespace of available metrics.

- By using the optional `pmgxvm` command provided with the Performance Co-Pilot `pcp.sw.monitor` package (an optional product available for purchase).

If you have the `pcp.sw.monitor` package, you can also use the `pmchart` command to view time-series data in the form of a moving graph. Figure 17-1 shows an example.

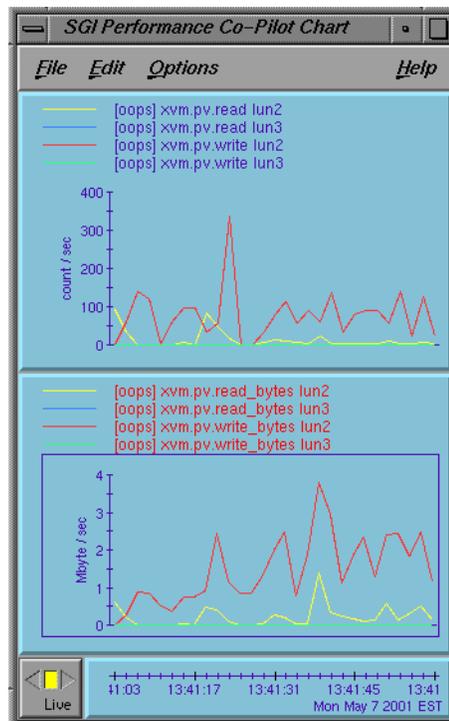


Figure 17-1 pmgxvm chart

I/O Fencing Status

To check the current fencing status, do one of the following:

- Select **View: Switches** in the GUI view area
- Use the `show switch` command within `cxfs_admin`
- Use the `hafence` command as follows:

```
/usr/cluster/bin/hafence -q
```

For example, the following output shows that all nodes are enabled.

```
# /usr/cluster/bin/hafence -q
Switch[0] "ptg-brocade" has 8 ports
  Port 1 type=FABRIC status=enabled hba=210000e08b0102c6 on host thunderbox
  Port 2 type=FABRIC status=enabled hba=210000e08b01fec5 on host whack
  Port 5 type=FABRIC status=enabled hba=210000e08b027795 on host thump
  Port 6 type=FABRIC status=enabled hba=210000e08b019ef0 on host thud
```

A fenced port shows `status=disabled`. For example:

```
# /usr/cluster/bin/hafence -q
Switch[0] "brocade04" has 16 ports
  Port 4 type=FABRIC status=enabled hba=210000e08b0042d8 on host o200c
  Port 5 type=FABRIC status=enabled hba=210000e08b00908e on host cxfs30
  Port 9 type=FABRIC status=enabled hba=2000000173002d3e on host cxfssun3
```

Verbose (`-v`) output would be as follows:

```
# /usr/cluster/bin/hafence -v
Switch[0] "brocade04" has 16 ports
  Port 0 type=FABRIC status=enabled hba=2000000173003b5f on host UNKNOWN
  Port 1 type=FABRIC status=enabled hba=2000000173003adf on host UNKNOWN
  Port 2 type=FABRIC status=enabled hba=210000e08b023649 on host UNKNOWN
  Port 3 type=FABRIC status=enabled hba=210000e08b021249 on host UNKNOWN
  Port 4 type=FABRIC status=enabled hba=210000e08b0042d8 on host o200c
  Port 5 type=FABRIC status=enabled hba=210000e08b00908e on host cxfs30
  Port 6 type=FABRIC status=enabled hba=2000000173002d2a on host UNKNOWN
  Port 7 type=FABRIC status=enabled hba=2000000173003376 on host UNKNOWN
  Port 8 type=FABRIC status=enabled hba=2000000173002c0b on host UNKNOWN
  Port 9 type=FABRIC status=enabled hba=2000000173002d3e on host cxfssun3
  Port 10 type=FABRIC status=enabled hba=2000000173003430 on host UNKNOWN
  Port 11 type=FABRIC status=enabled hba=200900a0b80c13c9 on host UNKNOWN
  Port 12 type=FABRIC status=disabled hba=0000000000000000 on host UNKNOWN
  Port 13 type=FABRIC status=enabled hba=200d00a0b80c2476 on host UNKNOWN
  Port 14 type=FABRIC status=enabled hba=1000006069201e5b on host UNKNOWN
```

```
Port 15 type=FABRIC status=enabled hba=1000006069201e5b on host UNKNOWN
```

A status of `enabled` for an `UNKNOWN` host indicates that the port is connected to a system that is not a node in the cluster. A status of `disabled` for an `UNKNOWN` host indicates that the node has been fenced (disabled), and the port may or may not be connected to a node in the cluster. A status of `enabled` with a specific name host indicates that the port is not fenced and is connected to the specified node in the cluster.

To check current failure action settings, use the `show failpolicy` command in `cxfs_admin` or the `cms_failconf` command as follows:

```
/usr/cluster/bin/cms_failconf -q
```

For example, the following output shows that all nodes except `thud` have the system default failure action configuration. The node `thud` has been configured for fencing and resetting.

```
# cms_failconf -q
CMS failure configuration:
    cell[0] whack      Reset Shutdown
    cell[1] thunderbox  Reset Shutdown
    cell[2] thud       Fence Reset
    cell[3] thump      Reset Shutdown
    cell[4] terry      Reset Shutdown
    cell[5] leesa      Reset Shutdown
```

Heartbeat Timeout Status

You can use Performance Co-Pilot or the IRIX `icrash` command to monitor heartbeat timeouts. For example, the following command prints the CXFS kernel messaging statistics:

```
# icrash -e "load cxfs; mtcp_stats"
corefile = /dev/mem, namelist = /unix, outfile = stdout
```

```
Please wait.....
Loading default Sial macros.....
```

```
>> load cxfs
```

```
>> mtcp_stats
STATS @ 0xc000000001beebb8
Max delays: discovery 500767 multicast 7486 hb monitor 0
hb generation histogram:(0:0)(1:0)(2:0)(3:0)(4:0)(5:0)
Improperly sized alive mesgs 0 small 0 big 0
Alive mesgs with: invalid cell 0 invalid cluster 0 wrong ipaddr 2
Alive mesgs from: unconfigured cells 100 cells that haven't discovered us 6000
mtcp_config_cell_set 0x0000000000000007
cell 0:starting sequence # 77 skipped 0
hb stats init @ 15919:(0:1)(1:478301)(2:29733)(3:0)(4:0)
cell 1:starting sequence # 0 skipped 0
hb stats init @ 360049:(0:1)(1:483337)(2:21340)(3:0)(4:0)
cell 2:starting sequence # 0 skipped 0
```

The following fields contain information that is helpful to analyzing heartbeat timing:

- **discovery:** The maximum time in HZ that the discovery thread (that is, the thread that processes incoming heartbeats) has slept. Because nodes generate heartbeats once per second, this thread should never sleep substantially longer than 100 HZ.

A value much larger than 100 suggests either that it was not receiving heartbeats or that something on the node prevented this thread from processing the heartbeats.
- **multicast:** The thread that generates heartbeats sleeps for 100 HZ after sending the last heartbeat and before starting on the next. This field contains the maximum time in HZ between the start and end of that sleep. A value substantially larger than 100 indicates a problem getting the thread scheduled; for example, when something else on the node is taking all CPU resources.
- **monitor:** The maximum time in HZ for the heartbeat thread to do its sleep and send its heartbeat. That is, it contains the value for `multicast` plus the time it takes to send the heartbeat. If this value is substantially higher than 100 but `multicast` is not, it suggests a problem in acquiring resources to send a heartbeat, such as a memory shortage.
- **gen_hist:** A histogram showing the number of heartbeats generated within each interval. There are 6 buckets tracking each of the first 5 seconds (anything over 5 seconds goes into the 6th bucket).
- **hb_stats:** Histograms for heartbeats received. There is one histogram for each node in the cluster.

- `seq_stats`: Number of consecutive incoming heartbeats that do not have consecutive sequence numbers. There is one field for each node. A nonzero value indicates a lost heartbeat message.
- `overdue`: Time when an overdue heartbeat is noticed. There is one field per node.
- `rescues`: Number of heartbeats from a node that are overdue but CXFS message traffic has been received within the timeout period.
- `alive_small`: Number of times a heartbeat message arrived that was too small, (that is, contained too few bytes).
- `alive_big`: Number of times a heartbeat arrived that was too large.
- `invalid_cell`: Number of heartbeats received from nodes that are not defined in the cluster
- `invalid_cluster`: Number of heartbeats received with the wrong cluster ID
- `wrong_ipaddr`: Number of heartbeats received with an IP address that does not match the IP address configured for the node ID
- `not_configured`: Number of heartbeats received from nodes that are not defined in the cluster
- `unknown`: Number of heartbeats from nodes that have not received the local node's heartbeat

Migration from an IRIX Cluster to an SGI ProPack Cluster

CXFS supports a running cluster with a single type of operating system for server-capable nodes: either all IRIX or all SGI ProPack. The information in this chapter will help you migrate from an IRIX cluster to an SGI ProPack cluster. It discusses the following:

- "An Overview of Differences Between IRIX and Linux System Administration" on page 415
- "Migration Procedure" on page 416
- "Migration Troubleshooting" on page 423

For assistance, contact SGI Managed Services.

An Overview of Differences Between IRIX and Linux System Administration

If you are migrating from a cluster with IRIX metadata servers to a cluster with SGI ProPack metadata servers, you should understand the differences between IRIX and Linux system administration. The details of these differences are beyond the scope of this guide, but a brief overview includes:

- Installation tools
- Mount options
- Paths
- Location of kernel system-tunable parameters

For more information, see the operating system documentation.

See also:

- "SGI ProPack Limitations and Considerations" on page 114
- Chapter 13, "Administration and Maintenance" on page 305
- Appendix D, "Operating System Path Differences" on page 533

Migration Procedure

Note: The following procedure assumes that the filesystems in the cluster you want to migrate do not have block sizes greater than the system page size and that they are not real-time filesystems. These types of filesystems are supported on IRIX but not on SGI ProPack.

The example in this chapter uses `cmgr`, but you could perform a similar procedure using `cxfs_admin` or the GUI. The example begins with a cluster named `performance` having a two IRIX server-capable nodes named `rum` and `snake` and a Solaris client-only node named `ray`:

```
rum # clconf_info
```

```
Event at [2004-02-13 07:57:17]
```

```
Membership since Thu Feb 12 15:15:26 2004
```

Node	NodeID	Status	Age	CellID
snake	1	up	2	1
rum	2	up	2	2
ray	3	up	1	0

```
1 CXFS FileSystems
```

```
/dev/cxvm/V9500 on /cxfs/V9500 enabled server=(snake) 2 client(s)=(ray,rum) status=UP
```

Do the following:

1. Unmount the CXFS filesystems cluster-wide within CXFS. For example:

```
cmgr> admin cxfs_unmount cxfs_filesystem V9500
```

```
cxfs_unmount operation successful
```

2. Mount and unmount the filesystems locally, which will ensure that the XFS log plays back cleanly. For example:

```
# mount /dev/cxvm/V9500 /mnt
```

```
# umount /mnt
```

3. Stop CXFS services on all nodes. For example on the IRIX node run:

```
cmgr> stop cx_services for cluster performance

CXFS services have been deactivated in cluster performance
```

Note: If you use `cxfs_admin`, you must issue a command for each node.

4. Define the administration node with the SGI ProPack operating system type. For example on the IRIX node run:

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

Hostname[optional] ? bang
Is this a FailSafe node <true|false> ? false
Is this a CXFS node <true|false> ? true
Operating System <IRIX|Linux32|Linux64|AIX|HPUX|Solaris|MacOSX|Windows> ? Linux64
Node Function <server_admin|client_admin|client_only> ? server_admin
Node ID[optional] ? 64
Partition ID[optional] ? (0)
Do you wish to define failure hierarchy[y/n]:y
Hierarchy option 0 <System|FenceReset|Fence|Reset|Shutdown>[optional] ? Fence
Hierarchy option 1 <System|FenceReset|Fence|Reset|Shutdown>[optional] ?
Reset type <powerCycle|reset|nmi> ? (powerCycle)
Do you wish to define system controller info[y/n]:n
Number of Network Interfaces ? (1)
NIC 1 - IP Address ? bang-p
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
```

Successfully defined node bang

5. Add the SGI ProPack administration node to the cluster. For example on the IRIX node run:

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

Is this a FailSafe cluster <true|false> ? (false)
Is this a CXFS cluster <true|false> ? (true)
```

```
Cluster Notify Cmd [optional] ?
Cluster Notify Address [optional] ?
Cluster CXFS mode <normal|experimental>[optional] ? (normal)
Cluster ID ? (1)
```

Current nodes in cluster performance:

```
Node - 1: ray
Node - 2: snake
Node - 3: rum
```

No networks in cluster performance

Add nodes to or remove nodes/networks from cluster performance
Enter "done" when completed or "cancel" to abort

```
performance ? add node bang
performance ? done
Added node <bang> to cluster <performance>
Successfully modified cluster performance
```

6. Modify the CXFS filesystems to remove the IRIX administration nodes as metadata servers and add the new SGI ProPack administration node as metadata server. For example, on the IRIX node rum:

```
cmgr> modify cxfs_filesystem V9500
```

(Enter "cancel" at any time to abort)

```
Device ? (/dev/cxvm/V9500)
Mount Point ? (/cxfs/V9500)
Mount Options[optional] ?
Use Forced Unmount ? <true|false> ? (false)
Grio Qualified Bandwidth[optional] ?
Grio managed filesystem ? <true|false>[optional] ?
Default Local Status ? (enabled)
```

MODIFY CXFS FILESYSTEM OPTIONS

- 0) Modify Server.
- 1) Add Server.
- 2) Remove Server.

- 3) Add Enabled Node.
- 4) Remove Enabled Node.
- 5) Add Disabled Node.
- 6) Remove Disabled Node.
- 7) Show Current Information.
- 8) Cancel. (Aborts command)
- 9) Done. (Exits and runs command)

Enter option:**2**

Current servers:

CXFS Server 1 - Rank: 0	Node: rum
CXFS Server 2 - Rank: 1	Node: snake

Server Node ? **rum**

- 0) Modify Server.
- 1) Add Server.
- 2) Remove Server.
- 3) Add Enabled Node.
- 4) Remove Enabled Node.
- 5) Add Disabled Node.
- 6) Remove Disabled Node.
- 7) Show Current Information.
- 8) Cancel. (Aborts command)
- 9) Done. (Exits and runs command)

Enter option:**2**

Current servers:

CXFS Server 1 - Rank: 1	Node: snake
-------------------------	-------------

Server Node ? **snake**

- 0) Modify Server.
- 1) Add Server.
- 2) Remove Server.
- 3) Add Enabled Node.
- 4) Remove Enabled Node.
- 5) Add Disabled Node.
- 6) Remove Disabled Node.

- 7) Show Current Information.
- 8) Cancel. (Aborts command)
- 9) Done. (Exits and runs command)

Enter option:1

No current servers

Server Node ? **bang**

Server Rank ? **1**

- 0) Modify Server.
- 1) Add Server.
- 2) Remove Server.
- 3) Add Enabled Node.
- 4) Remove Enabled Node.
- 5) Add Disabled Node.
- 6) Remove Disabled Node.
- 7) Show Current Information.
- 8) Cancel. (Aborts command)
- 9) Done. (Exits and runs command)

Enter option:9

Successfully modified cxfs_filesystem V9500

After you complete this step, the filesystems would show the following information:

```
cmgr> show cxfs_filesystem V9500
```

```
Name: V9500
Device: /dev/cxvm/V9500
Mount Point: /cxfs/V9500
Forced Unmount: false
Global Status: disabled
Default Local Status: enabled
```

```
Server Name: bang
```

```
Rank: 1
```

7. Remove the IRIX administration nodes from the cluster. For example, switching to the SGI ProPack node bang:

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

```
Is this a FailSafe cluster <true|false> ? (false)  
Is this a CXFS cluster <true|false> ? (true)  
Cluster Notify Cmd [optional] ?  
Cluster Notify Address [optional] ?  
Cluster CXFS mode <normal|experimental>[optional] ? (normal)  
Cluster ID ? (1)
```

Current nodes in cluster performance:

```
Node - 1: ray  
Node - 2: snake  
Node - 3: rum  
Node - 4: bang
```

Add nodes to or remove nodes/networks from cluster performance
Enter "done" when completed or "cancel" to abort

```
performance ? remove node rum  
performance ? remove node snake  
performance ? done  
Successfully modified cluster performance
```

8. Delete the IRIX administration nodes from the pool. For example, from the SGI ProPack node bang:

```
cmgr> delete node rum  
Deleted node (rum).
```

```
cmgr> delete node snake  
Deleted node (snake).
```

9. Start CXFS services for all nodes in the cluster. For example, from the SGI ProPack node bang:

```
cmgr> start cx_services for cluster performance
```

CXFS services have been activated in cluster performance

10. Mount the CXFS filesystems. For example, from the SGI ProPack node bang:

```
cmgr> admin cxfs_mount cxfs_filesystem V9500

cxfs_mount operation successful
```

11. If you are running other storage software products on server-capable nodes, confirm that you have installed either TPSSM or SMI, as appropriate for your hardware.

After completing this procedure, the cluster information is as follows:

```
[root@bang root]# clconf_info

Event at [2004-02-13 08:44:18]

Membership since Fri Feb 13 08:44:13 2004

-----
Node          NodeID Status  Age    CellID
-----
ray           3  up      1      0
bang         64  up      1      3
-----

1 CXFS FileSystems
/dev/cxvm/V9500 on /cxfs/V9500  enabled  server=(bang) 1 client(s)=(ray)  status=UP
```

For more information about using the `cmgr` command to perform this procedure, see the following:

- "Unmount a CXFS Filesystem with `cmgr`" on page 619
- "Stop CXFS Services with `cmgr`" on page 605
- "Define a Node with `cmgr`" on page 577
- "Modify a Cluster with `cmgr`" on page 601
- "Modify a CXFS Filesystem with `cmgr`" on page 619
- "Modify a Cluster with `cmgr`" on page 601
- "Delete a Node with `cmgr`" on page 593
- "Start CXFS Services with `cmgr`" on page 605
- "Mount a CXFS Filesystem with `cmgr`" on page 618

For more information about using the GUI, see the following:

- "Unmount CXFS Filesystems with the GUI" on page 246
- "Stop CXFS Services with the GUI" on page 228
- "Define a Node with the GUI" on page 208
- "Add or Remove Nodes in the Cluster with the GUI" on page 216
- "Modify a CXFS Filesystem with the GUI" on page 244
- "Add or Remove Nodes in the Cluster with the GUI" on page 216
- "Delete a Node with the GUI" on page 222
- "Start CXFS Services with the GUI" on page 227
- "Mount CXFS Filesystems with the GUI" on page 245

Migration Troubleshooting

The following sections discuss possible problems you may encounter after migrating from an IRIX cluster to an SGI ProPack cluster:

- "Filesystems Will Not Mount" on page 423
- "DMF Filesystems Will Not Mount" on page 424
- "DMF Filesystems Will Not Mount" on page 424

Filesystems Will Not Mount

Messages such as the following indicate that the filesystem was not cleanly unmounted from the IRIX metadata server:

```
Jan 29 22:06:07 4A:cxfs2 kernel: XFS: nil uuid in log - IRIX style log
Jan 29 22:06:07 5A:cxfs2 kernel: Starting XFS recovery on filesystem:
xvm-0 (dev: xvm-0)
Jan 29 22:06:07 4A:cxfs2 kernel: XFS: dirty log written in incompatible
format - can't recover
```

To resolve this problem, you must return to the IRIX node and then mount and unmount the filesystem locally on the IRIX node in order to replay the dirty log messages (as in step 2 above in "Migration Procedure" on page 416).



Caution: Do not steal the XVM volumes to the local host. Mounting `/dev/cxvm/volname` locally on `/mnt` is sufficient.

DMF Filesystems Will Not Mount

If you have DMF filesystems and have `dmi` as a mount option, you must edit the `/etc/sysconfig/sysctlfile` to turn on DMAPi probing in order to mount CXFS filesystems. Change the bottom line from:

```
DMAPI_PROBE="no"
```

to:

```
DMAPI_PROBE="yes"
```

Do Not Use `extlog` or `rtfs` Filesystems

If you have SGI ProPack server-capable nodes, you cannot use `extlog` or `rtfs` filesystems.

Troubleshooting

Configuring and administering a CXFS cluster can be a complex task. In general, most problems can be solved by rebooting a node. However, the topics in this chapter may help you avoid rebooting:

- "Troubleshooting Strategy" on page 425
- "Common Problems" on page 442
- "Understanding Error Messages" on page 458
- "Corrective Actions" on page 484
- "Reporting Problems to SGI" on page 492

You must connect the GUI to a node that has the `cluster_admin` software package installed. You can perform administrative tasks with `cxfs_admin` from any host with the appropriate access and network connection. See the *CXFS MultiOS Client-Only Guide for SGI InfiniteStorage* for additional troubleshooting information.

Troubleshooting Strategy

To troubleshoot CXFS problems, do the following:

- "Know the Troubleshooting Tools" on page 426
- "Identify the Cluster Status" on page 438
- "Eliminate a Residual Cluster" on page 440
- "Determine If a Node Is Fenced" on page 440
- "Locate the Problem" on page 441
- "Redirect Switch Logs" on page 442

To avoid problems in the first place, follow the recommendations in Chapter 9, "Best Practices" on page 139.

Know the Troubleshooting Tools

This section provides an **overview** of the tools required to troubleshoot CXFS:



Caution: Many of the commands listed are beyond the scope of this book and are provided here for quick reference only. See the other guides and man pages referenced for complete information before using these commands.

- "Physical Storage Tools" on page 426
- "Cluster Configuration Tools" on page 429
- "Cluster Control Tools" on page 430
- "Networking Tools" on page 431
- "Cluster/Node Status Tools" on page 431
- "Performance Monitoring Tools" on page 433
- "Kernel Status Tools" on page 434
- "Log Files" on page 437
- "Gather Cluster Configuration with `cxfsdump`" on page 437

Physical Storage Tools

Understand the following physical storage tools:

- To display the hardware inventory:

- IRIX:

```
irix# /sbin/hinv
```

- SGI ProPack:

```
[root@linux root]# hwinfo --short
```

If the output is not what you expected, do a probe for devices and perform a SCSI bus reset, using the following commands:

- IRIX:

```
irix# /usr/sbin/scsiha -pr bus_number
```

– SGI ProPack:

- QLogic SCSI or Fibre Channel: use the following to probe the LUN on the specified *hostname*:

```
[root@linux root]# echo "- - -" > /sys/class/scsi_host/hostname/scan
```

Each "-" character is a wildcard for bus, target, and LUN, respectively. Newer SCSI and all FC controllers have a single bus per function, but two functions in the dual-port controllers. For example, if you added a new LUN to a RAID (and the RAID is target 3) for a host named *host3*:

```
[root@linux root]# echo "0 3 -" > /sys/class/scsi_host/host3/scan
```

QLogic Fibre Channel: use the following to discover and build a new table for the LUN, where 3 is the host number:

```
[root@linux root]# echo "scsi-qlascan" >/proc/scsi/qla2xxx/3
```

- LSI: use the `lsiutil` tool to scan the HBA, selecting option 8 to scan for devices:

```
[root@linux root]# lsiutil
```

LSI Logic MPT Configuration Utility, Version 1.41, November 23, 2005

4 MPT Ports found

	Port Name	Chip Vendor/Type/Rev	MPT Rev	Firmware Rev
1.	/proc/mpt/ioc0	LSI Logic 53C1030 B2	102	01032710
2.	/proc/mpt/ioc1	LSI Logic 53C1030 B2	102	01032710
3.	/proc/mpt/ioc2	LSI Logic FC949X A1	105	01030300
4.	/proc/mpt/ioc3	LSI Logic FC949X A1	105	01030300

Select a device: [1-4 or 0 to quit] 3

1. Identify firmware, BIOS, and/or FCode
2. Download firmware (update the FLASH)
4. Download/erase BIOS and/or FCode (update the FLASH)
8. Scan for devices
10. Change IOC settings (interrupt coalescing)
13. Change FC Port settings
16. Display logged-in devices
20. Diagnostics

19: Troubleshooting

```
21. RAID actions
22. Reset bus
23. Reset target
30. Beacon on
31. Beacon off
60. Show non-default settings
61. Restore default settings
98. Reset FC link
99. Reset port
```

Main menu, select an option: [1-99 or e for expert or 0 to quit] 8

FC949X's link is online, type is fabric direct attach, speed is 2 Gbaud

B__T__L	Type	Vendor	Product	Rev	WWPN	PortId
0 127 0	Disk	SGI	TP9300	0612	200d00a0b8131841	021500
0 127 1	Disk	SGI	TP9300	0612		
0 127 2	Disk	SGI	TP9300	0612		
0 127 31	Disk	SGI	Universal Xport	0612		
0 128 0	Disk	SGI	TP9300	0612	200c00a0b8131841	021400
0 128 1	Disk	SGI	TP9300	0612		
0 128 2	Disk	SGI	TP9300	0612		
0 128 31	Disk	SGI	Universal Xport	0612		
0 129 0	Disk	SGI	TP9100 F PSEUDO	5903	23000050cc007d2c	021300
0 130 0	Disk	SGI	TP9100 F PSEUDO	5903	22000050cc007d2c	021200
	FC949X Port				100000062b0e4248	021700
	FCP Initiator				210000e08b1058d4	021000
	FCP Initiator				210100e08b3058d4	021100
	FCP Initiator				100000062b0e4249	021600
	Non-FCP				20fc006069c021b6	fffffc
	Non-FCP				2007006069c021b6	fffffe

You can run the `cxfs-reprobe` script look for devices and perform a SCSI bus reset if necessary. `cxfs-reprobe` will also issue an XVM probe to tell XVM that there may be new devices available:

– On server-capable nodes:

```
server# /var/cluster/clconfd-scripts/cxfs-reprobe
```

- On client-only nodes:

```
client# /var/cluster/cxfs_client-scripts/cxfs-reprobe
```

- To configure I/O devices on an IRIX node, use the following command:

```
irix# /sbin/ioconfig -f /hw
```

- To show the physical volumes, use the xvm command:

```
# /sbin/xvm show -v phys/
```

See the *XVM Volume Manager Administrator's Guide*.

Cluster Configuration Tools

Understand the following cluster configuration tools:

- To configure XVM volumes, use the xvm command:

```
# /sbin/xvm
```

See the *XVM Volume Manager Administrator's Guide*.

- To configure CXFS nodes and cluster, use the CXFS GUI or cxfs_admin:

- The GUI:

```
# /usr/sbin/cxfsmgr
```

See "GUI Features" on page 191 and Chapter 11, "Reference to GUI Tasks" on page 183.

- The cxfs_admin command:

See "Initial Setup with the cxfs_admin Command" on page 172 and Chapter 12, "Reference to cxfs_admin Tasks" on page 255.

- To reinitialize the database, use the cdbreinit command:

```
# /usr/cluster/bin/cdbreinit
```

See "Recreating the Cluster Database" on page 489.

- To check the cluster configuration, use the following command from a server-capable administration node in the cluster:

```
# /usr/cluster/bin/cxfs-config -all -check
```

SGI recommends that you run this command after any significant configuration change or whenever problems occur. For more information, see "Checking the Cluster Configuration with `cxfs-config`" on page 382.

Cluster Control Tools

Understand the cluster control tools:

- "Cluster Administration Daemons" on page 24
- "CXFS Client Daemon" on page 26

These commands are useful if you know that filesystems are available but are not indicated as such by the cluster status, or if cluster quorum is lost. However, note that `/etc/init.d/cxfs stop` and `/etc/init.d/cxfs stop` will cause CXFS to completely shut down on the local node.

See the following:

- "Ensure Cluster Database Membership Quorum Stability" on page 142
- "Restarting CXFS Services" on page 484
- "Clearing the Cluster Database" on page 485
- "Stopping and Restarting Cluster Administration Daemons" on page 489
- "CXFS Services" on page 25

Running this command on the metadata server will cause its filesystems to be recovered by another potential metadata server. See "Cluster Services Tasks with the GUI" on page 227.

Note: Relocation and recovery are supported only when using standby nodes. Relocation is disabled by default.

- To revoke and allow CXFS kernel membership on the local node, forcing recovery on the metadata server for the local node, use the GUI or the following `cxfs_admin` command:

```
cxfs_admin:clustername> disable node:nodename
```

Wait until recovery is complete before issuing a subsequent:

```
cxfs_admin:clustername> enable node:nodename
```

The local node cannot rejoin the CXFS kernel membership until recovery is complete.

Also see the following:

- "Revoke Membership of the Local Node with the GUI" on page 232
- "Allow Membership of the Local Node with the GUI" on page 232
- "Disable a Node with cxfs_admin" on page 279
- "Enable a Node with cxfs_admin" on page 279

Networking Tools

Understand the following networking tools:

- To send packets to network hosts:

- IRIX:

```
irix# /usr/etc/ping
```

- SGI ProPack:

```
[root@linux root]# /bin/ping
```

- To show network status:

- IRIX:

```
irix# /usr/etc/netstat
```

- SGI ProPack:

```
[root@linux root]# /bin/netstat
```

Cluster/Node Status Tools

Understand the following cluster/node status tools:

- To show which cluster daemons are running:

```
# ps -ef | grep cluster
```

See "Verify that the Cluster Daemons are Running" on page 164.

- To see cluster and filesystem status, use one of the following:
 - GUI:

```
# /usr/sbin/cxfsmgr
```

See "Display a Cluster with the GUI" on page 227.
 - `cxfsmgr` command:

```
# /usr/cluster/bin/cxfsmgr -c status
```

See "Display a Cluster with `cxfsmgr`" on page 285.
 - `clconf_info` command:

```
# /usr/cluster/bin/clconf_info
```
 - `cxfsmgr` command on an IRIX or SGI ProPack client-only node:

```
# /usr/cluster/bin/cxfsmgr
```

- To see the mounted filesystems:

- IRIX:

```
irix# /sbin/mount  
irix# /usr/sbin/df
```

- SGI ProPack:

```
[root@linux root]# /bin/mount  
[root@linux root]# /bin/df
```

You can also use the `df` command to report the number of free disk blocks

- To show volumes:

```
# /sbin/xvm show vol/
```

See the *XVM Volume Manager Administrator's Guide*.

Performance Monitoring Tools

Understand the following performance monitoring tools:

- To monitor system activity:

```
# /usr/bin/sar
```

- To monitor file system buffer cache activity on IRIX nodes:

```
irix# /usr/sbin/bufview
```

Note: Do not use bufview interactively on a busy IRIX node; run it in batch mode.

- To monitor operating system activity data on an IRIX node:

```
irix# /usr/sbin/osview
```

- To monitor system input/output device loading on an SGI ProPack node, use the iostat(1) command. For example, to monitor at 2-second intervals for 1000000 times:

```
[root@linux root]# iostat 2 1000000
```

- To monitor process status, memory consumption, paging activity, block I/O operations, interrupts, context switches, and processor usage on an SGI ProPack node, use the vmstat(8) command. For example, to monitor at 1-second intervals for 1000 times:

```
[root@linux root]# vmstat -a -n 1 1000
```

- To monitor the statistics for an XVM volume, use the xvm command:

```
# /sbin/xvm change stat on {concatname|stripename|physname}
```

See the *XVM Volume Manager Administrator's Guide*.

- To monitor system performance, use Performance Co-Pilot. See the *Performance Co-Pilot for IA-64 Linux User's and Administrator's Guide*, *Performance Co-Pilot for IRIX Advanced User's and Administrator's Guide*, the *Performance Co-Pilot Programmer's Guide*, and the pmie and pmieconf man pages.

Kernel Status Tools

Understand the following kernel status tools (this may require help from SGI service personnel):

- To determine IRIX kernel status, use the `icrash` command:

```
# /usr/bin/icrash
>> load -F cdfs
```

Note: Add the `-v` option to these commands for more verbose output.

- `cdfs` to list CXFS commands
- `dcvn` to obtain information on a single client vnode
- `dcvnlist` to obtain a list of active client vnodes
- `dsvn` to obtain information on a single server vnode
- `dsvnlist` to obtain a list of active server vnodes
- `mesglist` to trace messages to the receiver (you can pass the displayed object address to the `dsvn` command to get more information about the server vnodes and pass the thread address to the `mesgargs` command to get more information about the stuck message). For example (line breaks shown here for readability):

```
>> mesglist
```

```
Cell:2
```

TASK ADDR	MSG ID	TYPE	CELL	MESSAGE	Time(Secs)	Object
0xe0000030e5ba8000	14	Snt	0	I_dsvn_fcntl	0	N/A
0xe0000030e5ba8000	14	Cbk	0	I_ucopy_copyin	0	N/A
0xa80000000bb77400	1210	Rcv	0	I_dsxvn_allocate_1	1:06	(dsvn_t*)0xa80000000a7f8900

```
>> mesgargs 0xa80000000bb77400
```

```
(dsvn_t*)0xa80000000a7f8900
(dsxvn_allocate_1_in_t*)0xa800000001245060
  objid=0xa80000000a7f8910 (dsvn=0xa80000000a7f8900)
  offset=116655
```

```

length=0x1
total=1
mode=2
bmapi_flags=0x7
wr_ext_count=0
&state=0xa8000000012450b0 credid=NULLID
lent_tokens=0xa800000 (DVN_TIMES_NUM(SWR) | DVN_SIZE_NUM(WR) | DVN_EXTENT_NUM(RD))
reason_lent=0x24800000 (DVN_TIMES_NUM(CLIENT_INITIATED) | DVN_SIZE_NUM(CLIENT_INITIATED) |
DVN_EXTENT_NUM(CLIENT_INITIATED))
lender_cell_id=0
(dsxvn_allocate_1_inout_t*)0xa800000001245110
cxfs_flags=0x200
cxfs_gen=4661

```

```

>> dsvn 0xa80000000a7f8900
(dsvn_t*)0xa80000000a7f8900:
  flags 0x10
  kq.next 0xc000000001764508 kq.prev 0xc000000001764508
  &tsclient 0xa80000000a7f8a30 &tserver 0xa80000000a7f8a80
  bhv 0xa80000000a7f8910 dsvfs 0xa800000026342b80
  (cfs_frlock_info_t*)0xa80000000bfee280:
    wait: none
    held: none
  vp 0xa8000000224de500 v_count 2 vrgen_flags 0x0
  dmvn 0x0000000000000000
  objid 0xa80000000a7f8910 gen 4 obj_state 0xa80000000a7f8940
  (dsxvn_t*)0xa80000000a7f8900:
    dsvn 0xa80000000a7f8900 bdp 0xa800000010b52d30
    tkclient 0xa80000000a7f8a30 tserver 0xa80000000a7f8a80
    ext gen 4661 io_users 2 exclusive_io_cell -1
    oplock 0 oplock_client -1 &dsx_oplock_lock 0xa80000000a7f8b9

```

- `sinfo` to show clients/servers and filesystems
- `sthrad | grep cmsd` to determine the CXFS kernel membership state. You may see the following in the output:
 - `cms_dead()` indicates that the node is dead

- `cms_follower()` indicates that the node is waiting for another node to create the CXFS kernel membership (the leader)
 - `cms_leader()` indicates that the node is leading the CXFS kernel membership creation
 - `cms_declare_membership()` indicates that the node is ready to declare the CXFS kernel membership but is waiting on resets
 - `cms_nascent()` indicates that the node has not joined the cluster since starting
 - `cms_shutdown()` indicates that the node is shutting down and is not in the CXFS kernel membership
 - `cms_stable()` indicates that the CXFS kernel membership is formed and stable
- `tcp_channels` to determine the status of the connection with other nodes
 - `t -a -w filename` to trace for CXFS
 - `t cms_thread` to trace one of the above threads
- To determine SGI ProPack kernel status, use the KDB built-in kernel debugger.

When `kdb` is enabled, a system panic will cause the debugger to be invoked and the keyboard LEDs will blink. The `kdb` prompt will display basic information. To obtain a stack trace, enter the `bt` command at the `kdb` prompt:

```
kdb> bt
```

To get a list of current processes, enter the following:

```
kdb> ps
```

To backtrace a particular process, enter the following, where `PID` is the process ID:

```
kdb> btP PID
```

To get a dump, enter the following:

```
kdb> sr d
```

To exit the debugger, enter the following:

```
kdb> go
```

If the system will be run in graphical mode with `kdb` enabled, SGI highly recommends that you use `kdb` on a serial console so that the `kdb` prompt can be seen.

- To invoke internal kernel routines that provide useful debugging information, use the `idbg` command:

```
# /usr/sbin/idbg
```

- Use the appropriate version of `lcrash` and load the CXFS kerntypes:

```
# lcrash -x /boot/sgi-cxfs-kerntypes-kernelversion-architecturetype
```

Note: Do not use the version of `lcrash` that is shipped with SLES 9. Use the version of `lcrash` that is available from Supportfolio.

Log Files

Understand the log files discussed in "Status in Log Files" on page 400.

Gather Cluster Configuration with `cxfsdump`

Before reporting a problem to SGI, you should use the `cxfsdump` command to gather configuration information about the CXFS cluster, such as network interfaces, CXFS registry information, I/O, and cluster database contents. This will allow SGI support to solve the problem more quickly.

Note: In cluster mode (the default), the `cxfsdump` command requires `rsh/ssh` and `rcp/scp` access across all nodes in the cluster. You can use the `-secure` option to use secure remote connections.

You should run `cxfsdump` from a CXFS administration node in the cluster:

```
# /usr/cluster/bin/cxfsdump
```

The output will be placed in a file in the directory `/var/cluster/cxfsdump-data` directory on the CXFS administration node on which the `cxfsdump` command was run. The `cxfsdump` command will report the name and location of the file when it is finished.

To gather information about just the local node, use the `cxfsdump -local` option.

On Windows nodes, use the following menu selection to access the `\Program Files\CXFS\cxfsdump.exe` command:

```
Start
  > Programs
    > CXFS
      > CXFS Dump
```

You can configure the location of the dump by selecting the directory from a **browse for folder** dialog or type in the path in the edit field.

On Windows nodes, the `cxfsdump /?` command displays a help message. The `cxfsdump -help` command displays a help message on other nodes.

For more information about client-only nodes, see *CXFS MultiOS Client-Only Guide for SGI InfiniteStorage*.

Identify the Cluster Status

When you encounter a problem, identify the cluster status by answering the following questions:

- Are the cluster daemons running? See "Verify that the Cluster Daemons are Running" on page 164.
- Is the cluster state consistent on each node? Run the `clconf_info` command on each CXFS administration node and compare.
- Which nodes are in the CXFS kernel membership? Check the cluster status and the following files:
 - IRIX: `/var/adm/SYSLOG`
 - SGI ProPack: `/var/log/messages`
- Which nodes are in the cluster database (`fs2d`) membership? See the `/var/cluster/ha/log/fs2d_log` files on each CXFS administration node.
- Is the database consistent on all CXFS administration nodes? Determine this logging in to each administration node and examining the `/var/cluster/ha/log/fs2d_log` file and database checksum.
- Log onto the various CXFS client nodes or use the GUI view area display with details showing to answer the following:

- Are the devices available on all nodes? Use the following:
 - The `xvm` command to show the physical volumes:

```
xvm:cluster> show -v phys/
```
 - Is the client-only node in the cluster? Use the `cxfs_info` command.
 - List the contents of the `/dev/cxvm` directory with the `ls` command:

```
# ls /dev/cxvm
```
 - Use the `hinv` command to display the hardware inventory. See "Physical Storage Tools" on page 426.
- Are the filesystems mounted on all nodes? Use `mount` and `clconf_info` commands.
- Which node is the metadata server for each filesystem? Use the `clconf_info` command.

On the metadata server, use the `clconf_info` command.

- Is the metadata server in the process of recovery? Use the IRIX `icrash` command to search for messages and look at the following files:
 - IRIX: `/var/adm/SYSLOG`
 - SGI ProPack: `/var/log/messages`

Note: You must run the `sial` scripts version of `icrash` commands. See "Kernel Status Tools" on page 434.

Messages such as the following indicate that recovery status:

- In process:

```
Mar 13 11:31:02 1A:p2 unix: ALERT: CXFS Recovery: Cell 1: Client Cell 0 Died, Recovering </scratch/p9/local>
```

- Completed:

```
Mar 13 11:31:04 5A:p2 unix: NOTICE: Signaling end of recovery cell 1
```

- Are there any long running (>20 seconds) kernel messages? Use the `icrash mesglist` command to examine the situation. For an example, see "Kernel Status Tools" on page 434.

- If filesystems are not mounting, do they appear online in XVM? You can use the following `xvm` command:

```
xvm:cluster> show vol/*
```

Eliminate a Residual Cluster

Before you start configuring another new cluster, make sure no nodes are still in a CXFS membership from a previous cluster. Enter the following to check for a `cmsd` kernel thread:

- IRIX:

```
irix# icrash -e 'stthread | grep cmsd'
```

Note: You must run the `sial` scripts version of `icrash` commands. See "Kernel Status Tools" on page 434.

- SGI ProPack:

```
[root@linux root]# ps -ef | grep cmsd
```

If the output shows a `cmsd` kernel thread, perform a forced CXFS shutdown by entering the following:

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

Then check for a `cmsd` kernel thread again.

After waiting a few moments, if the `cmsd` kernel thread still exists, you must reboot the machine or leave it out of the new cluster definition. It will not be able to join a new cluster in this state and it may prevent the rest of the cluster from forming a new CXFS membership.

Determine If a Node Is Fenced

To determine if a node is fenced, log in to a CXFS administration node and use the `cxfs_admin status` command or the `hafence(1M)` command.

The following messages are logged when fencing changes:

```
Raising fence on cell cellID (nodename)
```

```
Lowering fence on cell cellID (nodename)
```

Locate the Problem

To locate the problem, do the following:

- Examine the log files (see "Log Files" on page 437):
 - Search for errors in all log files. See "Status in Log Files" on page 400. Examine all messages within the timeframe in question.
 - Trace errors to the source. Try to find an event that triggered the error.
- Use the `sial` scripts version of the IRIX `icrash` commands. See "Kernel Status Tools" on page 434 and the `icrash` man page.
- Use detailed information from the view area in the GUI to drill down to specific configuration information.
- Run the **Test Connectivity** task in the GUI. See "Test Node Connectivity with the GUI" on page 223.
- Determine how the nodes of the cluster see the current CXFS kernel membership by entering the following command on each CXFS administration node:

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

For more information, see "Check Cluster Status with `clconf_info`" on page 403.

- Check the following file on each CXFS administration node to make sure the CXFS filesystems have been successfully mounted or unmounted:
 - IRIX: `/var/adm/SYSLOG`
 - SGI ProPack: `/var/log/messages`

If a mount/unmount fails, the error will be logged and the operation will be retried after a short delay.

- Use the `sar` system activity reporter to show the disks that are active. For example, the following example for IRIX will show the disks that are active, put the disk name at the end of the line, and poll every second for 10 seconds:

```
irix# sar -DF 1 10
```

For more information, see the `sar` man page.

- Use the IRIX `bufview` filesystem 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` man page.

- Get a dump of the cluster database. You can extract such a dump with the following command:

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

Redirect Switch Logs

Brocade switch problems can cause CXFS to behave abnormally. For easier troubleshooting, use the `syslogdipadd` function on the switch to redirect its `syslogd` information to up to six potential metadata servers in the cluster. SGI recommends logging to at least two potential metadata servers on which you troubleshoot issues and look for error messages. The `syslogd` information is the same as that given by `errshow` command on the switch.

For example, on each switch, define the metadata server nodes MDS1 and MDS2 to which the switch can redirect its `syslogd` output:

```
switch:admin > syslogdipadd ipaddress_MDS1
switch:admin > syslogdipadd ipaddress_MDS2
```

The entries from the switch can be sorted because they are prefixed by the switch name, which is standard `syslogd` behavior.

Common Problems

The following are common problems and solutions:

- "Client Membership Loss" on page 443
- "Node is Permanently Fenced" on page 446
- "Cannot Access Filesystem" on page 446
- "Log Files Consume Too Much Disk Space" on page 446

- "Unable to Define a Node" on page 446
- "System is Hung" on page 447
- "Node is Detected but Never Joins Membership" on page 447
- "Cell ID Count and Membership delivered Messages" on page 447
- "You Cannot Log In" on page 448
- "I/O Error in Filesystem" on page 448
- "Cannot Mount Filesystems" on page 449
- "GUI Displays Invalid Filesystems" on page 449
- "Multiple `client_timeout` Values" on page 449
- "No HBA WWPNs are Detected" on page 450
- "XFS Internal Errors in System Log File" on page 452
- "Multiple Ethernet Interfaces on Altix Systems" on page 452
- "Clients Unable to Remount Filesystems" on page 453
- "Forced Filesystem Shutdown Messages and XFS File Corruption" on page 453
- "GUI Will Not Run" on page 454
- "IPMI Issues" on page 455
- "`cxfs_admin` Output is Not Current" on page 457
- "`clconfd` Is Not Running" on page 457

Client Membership Loss

The following messages indicate that a client has lost membership (line breaks added here for readability):

```
Mar 15 10:55:35 5A:mvxcfs2 kernel: Error -1 reading mesg header channel 0 cell 4 (mvxcfs17)
[priority 1 at 192.168.17.173 via 192.168.17.48]
Mar 15 10:55:35 4A:mvxcfs2 kernel: Error receiving messages from cell 4 (mvxcfs17) tcpchannel 0
[priority 1 at 192.168.17.173 via 192.168.17.48]
Mar 15 10:55:36 5A:mvxcfs2 kernel: Error -1 reading mesg header channel 1 cell 4 (mvxcfs17)
```

```
[priority 1 at 192.168.17.173 via 192.168.17.48]
Mar 15 10:55:36 4A:mvxfs2 kernel: Error receiving messages from cell 4 (mvxfs17) tcpchannel 1
[priority 1 at 192.168.17.173 via 192.168.17.48]
Mar 15 10:55:36 5A:mvxfs2 kernel: Error -1 reading mesg header channel 1 cell 4 (mvxfs17)
[priority 2 at 163.154.17.173 via 163.154.17.48]
Mar 15 10:55:36 4A:mvxfs2 kernel: Error receiving messages from cell 4 (mvxfs17) tcpchannel 1
[priority 2 at 163.154.17.173 via 163.154.17.48]
Mar 15 10:55:36 4A:mvxfs2 kernel: Transport failure cell 4
[priority 2 at 163.154.17.173 via 163.154.17.48] 0 of 2 interfaces up
Mar 15 10:55:36 6A:mvxfs2 kernel: Heartbeat Monitor:Failure time-stamp 295789 ticks:Last heartbeat
time-stamp 289940 ticks:Time-stamp delta 5849 ticks (5 seconds):Heartbeat timeout 5120 ticks (5 seconds)
```

The Error receiving and Error reading messages indicate that the message channel went down. The last message, which includes the Heartbeat Monitor string, contains other strings that give a clues as to why the channel was disconnected. Table 19-1 on page 445 lists all of the possible strings that may be included.

Table 19-1 Error Strings

String	Description
Aggregate Recover Transport	Failover has forced the transport down because the remote node has detected an error on the transport.
Aggregate Send	An error has occurred while attempting to send a message on the underlying socket. The most likely reason is that the message channel has been disconnected by the remote end.
Cell Up	An error occurred while attempting to establish a connection with the remote node.
disable heartbeat	A configuration change has eliminated the node from the cluster or the local node is shutting down CXFS.
Failure time-stamp	The time-stamp in ticks of when the error was detected.
Heartbeat Processing	A heartbeat has been received from the node that indicates it has dropped the local node from its set of known nodes.
Heartbeat Monitor	A heartbeat timeout has been detected.
Heartbeat timeout	The configured timeout in ticks and in seconds.
Last heartbeat time-stamp	The time-stamp in ticks when the last heartbeat from the remote node was received.
Message Failure	One of the following: <ul style="list-style-type: none"> • An internal messaging error (for example, a corrupt header has been received) . This brings down all transports connected to the remote node. This is a serious error that indicates a problem in the local node, the remote node, or the network that is causing corruption. • A socket error has occurred while attempting to send a message. The most likely reason is that the message channel has been disconnected by the remote end.
Receive Thread	A socket error has occurred when attempting to receive a message. The most likely reason is that the message channel has been disconnected by the remote end.
Time-stamp delta	The difference in ticks and in seconds. If this delta is greater than the configured heartbeat timeout, then it is definitively a heartbeat timeout.

In the above example, the last message indicates that there is a heartbeat timeout because the string `Heartbeat Monitor` is included. The message also indicates that the error was detected at 295789 ticks (`Failure time-stamp string`) and that the configured timeout is 5120 ticks or 5 seconds (the `Heartbeat timeout string`). The delta is 5849 ticks or 5 seconds (the `Time-stamp delta string`), therefore it is a heartbeat timeout because the delta is greater than the configured heartbeat timeout.

Node is Permanently Fenced

If you are unable to raise the fence on a node, it may be that the switch ports are unable to determine the WWPN. See "Hardware Changes and I/O Fencing" on page 367.

Cannot Access Filesystem

If you cannot access a filesystem, check the following:

- Is the filesystem enabled? Check the GUI and `clconf_info` command.
- Were there mount errors?

Log Files Consume Too Much Disk Space

If the log files are consuming too much disk space, you should rotate them; see "Log File Management" on page 339. You may also want to consider choosing a less-verbose log level; see the following:

- "`cad.options` on CXFS Administration Nodes" on page 129
- "`fs2d.options` on CXFS Administration Nodes" on page 130
- "Configure Log Groups with the GUI" on page 231

Unable to Define a Node

If you are unable to define a node, it may be that there are hostname resolution problems. See "Hostname Resolution and Network Configuration Rules" on page 93.

System is Hung

The following may cause the system to hang:

- Overrun disk drives.
- Heartbeat was lost. In this case, you will see a message that mentions `withdrawl` of node.
- As a last resort, do a non-maskable interrupt (NMI) of the system and contact SGI. (The NMI tells the kernel to panic the node so that an image of memory is saved and can be analyzed later.) For more information, see the owner's guide for the node.

Make the following files available:

- System log file:
 - IRIX: `/var/adm/SYSLOG`
 - SGI ProPack: `/var/log/messages`
- IRIX `vmcore.#.comp`
- IRIX `unix.#`

Node is Detected but Never Joins Membership

If a node is detected in the system log file but it never receives a `Membership delivered` message, it is likely that there is a network problem.

See "Configuring System Files" on page 127.

Cell ID Count and `Membership delivered` Messages

The `Membership delivered` messages in the system log file include a bitmask with a bit set for the cell IDs of nodes that are members of the new CXFS membership. The `Membership delivered` messages are followed by one or more messages starting with `Cell(age):` that print the individual cell IDs and the ages of their membership. `0xXXX` is a binary bitmask of cells included in the membership. In the following example, cell 0 has been in the last 21 CXFS memberships:

```
NOTICE: Membership delivered for cells 0x3.  
Cell(age): 0(21) 1(12)
```

If the Membership delivered messages are appearing frequently in the system log file, it may indicate a network problem:

- Nodes that are stable and remain in the membership will have a large membership version number.
- Nodes that are having problems will be missing from the messages or have a small membership version number.

See "Configuring System Files" on page 127.

You Cannot Log In

If you cannot log in to a CXFS administration node, you can use one of the following commands, assuming the node you are on is listed in the other nodes' `.rhosts` files:

```
# rsh hostname ksh -i
# rsh hostname csh -i
```

I/O Error in Filesystem

The following message indicates a problem (output lines wrapped here for readability):

```
ALERT: I/O error in filesystem ("/mnt") metadata dev 0xbd block 0x41df03 ("xlog_iodone")
ALERT:      b_error 0 b_bcount 32768 b_resid 0
NOTICE: xfs_force_shutdown(/mnt,0x2) called from line 966 of file ../fs/xfs/xfs_log.c.
Return address = 0xc000000008626e8
ALERT: I/O Error Detected.  Shutting down filesystem: /mnt
ALERT: Please umount the filesystem, and rectify the problem(s)
```

You can fix this problem using `xfs_repair` only if there is no metadata in the XFS log. See "Forced Filesystem Shutdown Messages and XFS File Corruption" on page 453, for the appropriate procedure.

I/O errors can also appear if the node is unable to access the storage. This can happen for several reasons:

- The node has been physically disconnected from the SAN
- A filesystem shutdown due to loss of membership
- A filesystem shutdown due to lost of the metadata server

- The node has been fenced out of the SAN

Cannot Mount Filesystems

If you are unable to raise the fence on a node, it may be that the switch ports are unable to determine the WWPN. See "Hardware Changes and I/O Fencing" on page 367.

If you have defined filesystems and then rename your cluster (by deleting the old cluster and defining a new cluster), CXFS will not be able to mount the existing filesystems. This happens because the clustered XVM volume on which your CXFS filesystem resides is not accessible to the new cluster, and the volumes are therefore considered as foreign.

In order to mount the filesystem on the new cluster, you must use the XVM `steal` command to bring the clustered XVM volume into the domain of the new cluster. For more information, see the *XVM Volume Manager Administrator's Guide*.

GUI Displays Invalid Filesystems

If you create new slices on a previously sliced disk that have the same starting blocks as slices already existing on the disk, and if the old slices had filesystems, then the GUI will display those old filesystems even though they may not be valid.

Multiple `client_timeout` Values

A `client_timeout` value is set by the `clconfd` and `cxfs_client` daemons. The value depends on the order in which filesystems are mounted on the various nodes. The value adapts to help ensure that all filesystems get mounted in a timely manner. The value has no effect on the filesystem operation after it is mounted.

The value for `client_timeout` may differ among nodes, and therefore having multiple values is not really a problem.

The `retry` value is forced to be 0 and you cannot change it.



Caution: You should not attempt to change the `client_timeout` value. Improperly setting the values for `client_timeout` and `retry` could cause the mount command to keep waiting for a server and could delay the availability of the CXFS filesystems.

No HBA WWPNs are Detected

On most platforms, the `cxfs_client` software automatically detects the world wide port names (WWPNs) of any supported host bus adapters (HBAs) in the system that are connected to a switch that is configured in the cluster database. These HBAs will then be available for fencing.

However, if no WWPNs are detected, there will be messages logged to the following file:

- IRIX: `/var/adm/cxfs_client`
- SGI ProPack: `/var/log/cxfs_client`

If no WWPNs are detected, you can manually specify the WWPNs in the `/etc/fencing.conf` fencing file for the SGI ProPack platform. This method does not work if the WWPNs are partially discovered.

The fencing file is not used on the IRIX platform.

The fencing file enumerates the worldwide port name for all of the HBAs that will be used to mount a CXFS filesystem. There must be a line for the HBA WWPN as a 64-bit hexadecimal number.

Note: The WWPN is that of the HBA itself, **not** any of the devices that are visible to that HBA in the fabric.

If used, the fencing file must contain a simple list of WWPNs, one per line.

If you use the fencing file, you must update it whenever the HBA configuration changes, including the replacement of an HBA.

Do the following:

1. Set up the switch and HBA.
2. Follow the Fibre Channel cable on the back of the node to determine the port to which it is connected in the switch. Ports are numbered beginning with 0. (For example, if there are 8 ports, they will be numbered 0 through 7.)
3. Use the `telnet` command to connect to the switch and log in as user `admin` (the password is `password` by default).
4. Execute the `switchshow` command to display the switches and their WWPN numbers.

For example:

```
brocade04:admin> switchshow
switchName:      brocade04
switchType:      2.4
switchState:     Online
switchRole:      Principal
switchDomain:    6
switchId:        fffc06
switchWwn:       10:00:00:60:69:12:11:9e
switchBeacon:    OFF
port 0: sw Online      F-Port 20:00:00:01:73:00:2c:0b
port 1: cu Online      F-Port 21:00:00:e0:8b:02:36:49
port 2: cu Online      F-Port 21:00:00:e0:8b:02:12:49
port 3: sw Online      F-Port 20:00:00:01:73:00:2d:3e
port 4: cu Online      F-Port 21:00:00:e0:8b:02:18:96
port 5: cu Online      F-Port 21:00:00:e0:8b:00:90:8e
port 6: sw Online      F-Port 20:00:00:01:73:00:3b:5f
port 7: sw Online      F-Port 20:00:00:01:73:00:33:76
port 8: sw Online      F-Port 21:00:00:e0:8b:01:d2:57
port 9: sw Online      F-Port 21:00:00:e0:8b:01:0c:57
port 10: sw Online     F-Port 20:08:00:a0:b8:0c:13:c9
port 11: sw Online     F-Port 20:0a:00:a0:b8:0c:04:5a
port 12: sw Online     F-Port 20:0c:00:a0:b8:0c:24:76
port 13: sw Online     L-Port 1 public
port 14: sw No_Light
port 15: cu Online     F-Port 21:00:00:e0:8b:00:42:d8
```

The WWPN is the hexadecimal string to the right of the port number. For example, the WWPN for port 0 is 2000000173002c0b (you must remove the colons from the WWPN reported in the switchshow output to produce the string to be used in the fencing file).

5. Create the `/etc/fencing.conf` fencing file and add the WWPN for the port determined in step 2. (Comment lines begin with #.)

For dual-ported HBAs, you must include the WWPNs of any ports that are used to access cluster disks. This may result in multiple WWPNs per HBA in the file; the numbers will probably differ by a single digit.

For example, if you determined that port 0 is the port connected to the switch, your fencing file should contain the following:

```
# WWPN of the HBA installed on this system
#
2000000173002c0b
```

6. After the node is added to the cluster, enable the fencing feature by using the CXFS GUI, `hafence cxfs_admin` on a CXFS administration node.

XFS Internal Errors in System Log File

After a filesystem has been defined in CXFS, running `mkfs` on it (or using "Make Filesystems with the GUI" on page 238) will cause XFS internal errors to appear in the system log file. For example (line breaks added for readability):

```
Aug 17 09:25:52 1A:yokohama-mdsl unix: ALERT: Filesystem "(NULL)": XFS internal error
xfs_mount_validate_sb(4) at line 237 of file ../fs/xfs/xfs_mount.c.
Caller 0xc000000000326ef4
```

```
Aug 17 09:14:52 6X:yokohama-mdsl clconfd[360]: < E clconf 11> CI_FAILURE, fsinfo_update(/dev/cxvm/work)
kernel returned 1010 (Filesystem is corrupted)
```

To avoid these errors, run `mkfs` before defining the filesystem in CXFS, or delete the CXFS filesystem before running `mkfs`.

Multiple Ethernet Interfaces on Altix Systems

In Altix systems with multiple Ethernet interfaces, the default behavior of the operating system is to dynamically assign interface names (such as `eth0`, `eth1`, and so on) at boot time. Therefore, the physical interface associated with the `eth0` device may change after a system reboot; if this occurs, it will cause a networking problem for CXFS. To avoid this problem, provide persistent device naming by using the `/etc/sysconfig/networking/eth0_persist` file to map specific Ethernet device names to specific MAC addresses. Adding lines of the format to the `eth0_persist` file:

```
ethN MAC_ID
```

For example:

```
eth0 08:00:69:13:dc:ec  
eth1 08:00:69:13:72:e8
```

For more information about persistent naming, see *SGI ProPack for Linux Start Here*.

Clients Unable to Remount Filesystems

If you have multiple metadata servers in the cluster but only one potential metadata server defined for a given filesystem and that server goes down, the now server-less filesystem goes into a shutdown state. Although the clients maintain membership in the cluster, they will not remount the filesystem automatically when the potential metadata server comes back up. You must manually unmount the filesystem.

If there had been only one potential metadata server in the cluster, the filesystem's clients would have lost membership and gone through a forced shutdown, which automatically unmounts the filesystems.

Forced Filesystem Shutdown Messages and XFS File Corruption

Forced filesystem shutdown messages **do not** necessarily imply that `xfs_repair` should be run. Following is an example of a message that does indicate an XFS file corruption:

```
XFS read error in file system metadata block 106412416
```

When a filesystem is forcibly shut down, the log is not empty — it contains valuable metadata. You must replay it by mounting the filesystem. The log is only empty if the filesystem is unmounted cleanly (that is, not a forced CXFS shutdown, not a crash). You can use the following command line to see an example of the transactions captured in the log file:

```
# xfs_logprint -t device
```

If you run `xfs_repair` before mounting the filesystem, `xfs_repair` will delete all of this valuable metadata.

You should run `xfs_ncheck` and capture the output to a file before running `xfs_repair`. If running `xfs_repair` results in files being placed in the `lost+found` directory, the saved output from `xfs_ncheck` may help you to identify the original names of the files.



Caution: Always contact SGI technical support before using `xfs_repair` on CXFS filesystems. See "Repair Filesystems with Care" on page 154.

If you think you have a filesystem with real corruption, do the following:

1. Mount the device in order to replay the log:

```
# mount device any_mount_point
```

2. Unmount the filesystem:

```
# umount device
```

3. Check the filesystem:

```
# xfs_check device
```

4. View the repairs that could be made, using `xfs_repair` in no-modify mode:

```
# xfs_repair -n device
```

5. Capture filesystem file name and inode pairs:

```
# xfs_ncheck device > xfs_ncheck.out
```

6. If you are certain that the repairs are appropriate, complete them:

```
# xfs_repair device
```

GUI Will Not Run

If the GUI will not run, check the following:

- Is the license key properly installed on the server-capable node? See the following:
 - "Verify the License" on page 164
 - "License Key Error" on page 464
- Are the cluster daemons running? See "Verify that the Cluster Daemons are Running" on page 164.

- Are the `tcpmux` and `tcpmux/sgi_sysadm` services enabled in the following files?
 - IRIX: `/etc/inetd.conf`
 - SGI ProPack: `/etc/xinetd.d/tcpmux` and `/etc/tcpmux.conf`
- Are the `inetd` or `tcp` wrappers interfering? This may be indicated by `connection refused` or `login failed` messages.
- Are you connecting to a CXFS administration node? The `cxfsmgr` command can only be executed on a CXFS administration node. The GUI may be run from another system via the Web if you connect the GUI to a CXFS administration node.

IPMI Issues

This section discusses the following IPMI issues:

- "BMC Does Not Respond to a `ping` Command" on page 455
- "`ipmitool` Command Fails" on page 455
- "Node is Not Reset" on page 457

BMC Does Not Respond to a `ping` Command

If the baseboard management controller (BMC) does not respond to a `ping(8)` command from a remote node, verify that the BMC has a valid IP address assigned. See step 4 in "BMC System Controller" on page 539.

Note: The BMC will not respond to the `ping` command when issued from the local node (the node containing the BMC).

`ipmitool` Command Fails

If an `ipmitool(1)` command issued to a local BMC device (the node containing the BMC) fails, check the following:

- Are the IPMI modules loaded? See step 2 in "BMC System Controller" on page 539.
- Does the IPMI device exist? The default device name is `/dev/ipmi0`.
- Has the `admin` user name and password been set on the BMC with the required ADMINISTRATOR privileges? See step 3 in "BMC System Controller" on page 539.

- Does the BMC have a valid IP address assigned? See step 4 in "BMC System Controller" on page 539.
- Does the `ipmitool` command line contain all of the required arguments, including the OEM identifier and the device path? The basic command line used for a local node is as follows:

```
ipmitool -o intelplus -d /dev/ipmi0 command
```

For example:

```
[root@linux root] ipmitool -o intelplus -d /dev/ipmi0 power status
Chassis Power is on
```

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

If an `ipmitool(1)` command issued to the BMC from a remote node fails, check the following:

- Does the BMC respond to the `ping(8)` command? See "BMC Does Not Respond to a ping Command" on page 455.
- Is the correct version of `ipmitool` installed? See step 1 in "BMC System Controller" on page 539.
- Have the admin user name and password been set on the BMC with the required ADMINISTRATOR privileges? See step 3 in "BMC System Controller" on page 539.
- Does the `ipmitool` command contain all of the required arguments, including the `lanplus` interface, the OEM identifier, and the IP address (or alias) for the BMC? The basic command line used from a remote node is as follows:

```
ipmitool -I lanplus -o intelplus -H bmc-nodename -U admin -P admin_password command
```

For example:

```
[root@linux root] ipmitool -I lanplus -o intelplus -H my-bmc-node \
-U admin -P mypassword power status
Chassis Power is on
```

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

- Does the BMC IP address (or alias) specified with the `ipmitool -H` command respond to a `ping(8)`?

- Does the BMC have address resolution protocol (ARP) and gratuitous ARP configured, with the ARP interval set to 5 seconds? (An interval of 5 seconds is supported for CXFS.) See step 4 in "BMC System Controller" on page 539.

Node is Not Reset

If a node is not properly reset by CXFS, check the following:

- Does the node's failpolicy contain `Reset` or `FenceReset`? See the following:
 - "Modify a Node Definition with the GUI" on page 217
 - "Create or Modify a Node with `cxfs_admin`" on page 270
- Does the BMC respond to a `ping(8)` command from the node defined as the `reset_node`? See "BMC Does Not Respond to a `ping` Command" on page 455.
- Does `ipmitool(1)` work correctly from the node defined as the `reset_node`? Check the system log files for relevant error messages and see the following:
 - "ipmitool Command Fails" on page 455
 - "BMC System Controller" on page 539

`cxfs_admin` Output is Not Current

If the `cxfs_admin` output appears to be stale (such as after you manually change the port status, in which case the CXFS database is not informed), you can update the CXFS database by running the following command:

```
# hafence -U
```

`clconfd` Is Not Running

Sending `clconfd` a `SIGTERM` signal, the default signal sent by the `kill(1)` command, will cause the `clconfd` process to terminate. When the `clconfd` process terminates on a `SIGTERM` signal, it is not restarted by `cmond` and the node will remain in the CXFS cluster membership. All filesystem activity will continue without interruption. However, if `clconfd` is not running on one or more server-capable nodes in the cluster, configuration changes cannot be made in the cluster and CXFS recovery may hang, preventing nodes from joining the cluster membership.

Understanding 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.

Sections are as follows:

- "Normal Messages" on page 458
- "Relocation Error" on page 461
- "Controller Disable Messages" on page 461
- "CMS Error Messages" on page 461
- "c1confd Daemon Death" on page 462
- "Out of Logical Swap Space" on page 462
- "No Cluster Name ID Error" on page 463
- "Lost CXFS Membership" on page 464
- "License Key Error" on page 464
- "IP Address Error" on page 465
- "System Log File Errors" on page 466
- "Log File Error Messages" on page 476
- "cxfs_admin Errors" on page 482
- "Mount Errors" on page 483
- "SYSLOG credid Warnings" on page 483

Normal Messages

You can expect to see the following messages. They are normal and do not indicate a problem.

```
NOTICE: Error reading mesg header 4 channel 1 cell 2
```

```
      Error number 4 (EINTR) on MEMBERSHIP message channel (channel  
      1; channel 0 is the main channel for CXFS and XVM data) for
```

connection with node 2. The 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 for cells 0x2

Membership has been delivered for the specified node. 0xXXX is a binary bitmask of cell numbers for which membership has been delivered; 0x2 equates to cell 1.

Cell(age): 0(4) 1(2) 2(9)

Shows the cell and its age (the number of memberships it has been part of). One or more of these messages always follows a Membership delivered message.

NOTICE: Cell 3 (client) has joined the membership

The node with the specified cell ID has joined the membership. This message precedes a Membership delivered message if a node joined the membership.

NOTICE: Cell 3 (client) has left the membership

This message precedes a Membership delivered message if a node has left the membership.

NOTICE: Resetting cells 0x4

The number here is a bitmask of node numbers on which a reset is being requested. 0xXXX is a binary bitmask of cells being reset. In this case, 0x4 equates to cell 2. This is an informative message; no corrective action is required.

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. If you do not have reset capability, this message can be ignored. System reset configuration is recommended for all potential metadata servers.

```
NOTICE: Error reading mesg header 4 channel 1 cell 2
```

The mesg header 4 text indicates that this is just an informative message.

```
clconfd[16574]: <<CI> E config 2> CI_ERR_NOTFOUND, Error
reading CMS status for machine tango, assuming machine is
FailSafe-disabled in cluster twango.
```

This indicates that the cluster is CXFS only and that you are not using FailSafe.

```
CI_CLCONFERR_INIT in ep_name() not binding socket
```

This message appears before the daemons start.

```
clconfd[16574]: <<CI> E clconf 0> CI_CLCONFERR_INIT, in
ep_name(): not binding socket
```

This clconfd message appears when daemons are starting up.

```
date <I0 clconfd clconf 610:0 clconfd_client.c:84> client
registration: clconfinfo, id 9119
date<I0 clconfd clconf 610:0 clconfd_service.c:781> sending reply
configuration and membership msg to client: clconfinfo, id 9119
date <I0 clconfd clconf 610:0 clconfd_client.c:96> client
un-registration: clconfinfo, id 9119
```

These messages are issued if you run the clconf_info command. The clconf_info command first registers as a CXFS client with clconfd; it then gets a reply message to its request for configuration and membership status; finally, it unregisters when it is done.

```
date <I0 clconfd clconf 610:0 clconfd_service.c:781 sending reply
configuration and membership msg to client: cad, id 602
```

This message indicates that the cad daemon is polling clconfd for status regularly. cad does not register and unregister each time like clconf_info because it is a daemon and it does not exit after each request. You will see register/unregister messages for cad only when cad or clconfd restarts.

```
dcvn_import_force: error 1502 from invk_dsvn_obtain_exist
```

This is a normal message sent during the recovery process.

```
kernel: cxfs_cconnect_loop: cxfs_connect_find returns error =
110
```

This message will be produced if a filesystem is not successfully mounted within the designated timeout period. The mount will be retried.

Relocation Error

If you try to relocate a filesystem and see an error similar to the following `cxfs_admin` example, it means that relocation has not been enabled:

```
Error returned from server: feature not enabled (12)
Command "relocate slicelC server=server1" failed during commit: feature not enabled
```

To allow the relocation to occur, you must enable relocation as specified in "Relocation" on page 20.

Controller Disable Messages

If you see messages such as the following on the console or in a message log, it means that the Fibre Channel switch is misconfigured:

```
controller disable is not supported on loop
```

CXFS fencing recovery operations do not support loop mode. Verify that all Fibre Channel switches are configured correctly. See the switch documentation for configuration information.

CMS Error Messages

The following messages may be logged by CMS.

```
CMS excluded cells 0xXXX with incomplete connectivity
```

Generated when CMS delivers a membership that excluded some **new** cells that had not established connections with enough cells yet to be admitted. `0xXXX` is a binary bitmask of excluded cells.

```
CMS calculation limited to last membership:configuration change incomplete on cells 0xXXX
```

Generated when the leader is attempting to make a configuration change current (that is, actually use the change on all nodes), but some cells in the cluster have not yet gotten the configuration change staged (uploaded and ready to be made current). 0xXXX is a binary bitmask of cells that do not yet have the change in their configuration. Changes make their way through the cluster asynchronously, so this situation is expected. It can take a few attempts by the CMS leader before all nodes have the change staged. As long as this situation resolves eventually, there is no problem. For more information, use `idbg cms_info`.

`CMS calculation limited to last membership:recovery incomplete`

Generated when new members were disallowed due to recovery from the last cell failure that is still being processed.

clconfd Daemon Death

If the `clconfd` daemon exits immediately after it starts up, it usually means that the CXFS license key has not been properly installed. Check the end of the `clconfd` log file (`/var/cluster/ha/log/clconfd_nodename`) for error messages. For information about licensing error messages, see "License Key Error" on page 464.

You must properly install the license keys before you can use CXFS. If you increase the number of CPUs in your system, you may need a new license key. See Chapter 4, "CXFS License Keys" on page 75.

Out of Logical Swap Space

The following example system log file message indicates an oversubscribed system:

```
ALERT: inetd [164] - out of logical swap space during fork while
allocating uarea - see swap(1M)
Availsmem 8207 availrmem 427 rlx freemem 10, real freemem 9
```

See "Use System Capacity Wisely" on page 158.

The cluster daemons could also be leaking memory in this case. You may need to restart them:

- On administration nodes:
 - IRIX:

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

– SGI ProPack:

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

- On client-only nodes:

```
# killall cxfs_client  
# /etc/init.d/cxfs_client start
```

No Cluster Name ID Error

For example:

```
Mar 1 15:06:18 5A:nt-test-07 unix: NOTICE: Physvol (name cip4) has no  
CLUSTER name id: set to ""
```

This message means the following:

- The disk labeled as an XVM physvol was probably labeled under IRIX 6.5.6f and the system was subsequently upgraded to a newer version that uses a new version of XVM label format. This does not indicate a problem.
- The cluster name had not yet been set when XVM encountered these disks with an XVM cluster physvol label on them. This is normal output when XVM performs the initial scan of the disk inventory, before node/cluster initialization has completed on this host.

The message indicates that XVM sees a disk with an XVM cluster physvol label, but that this node has not yet joined a CXFS membership; therefore, the cluster name is empty ("").

When a node or cluster initializes, XVM rescans the disk inventory, searching for XVM cluster physvol labels. At that point, the cluster name should be set for this host. An empty cluster name after node/cluster initialization indicates a problem with cluster initialization.

The first time any configuration change is made to any XVM element on this disk, the label will be updated and converted to the new label format, and these notices will go away.

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

Lost CXFS Membership

The following message in the system log file indicates a kernel-triggered revocation of CXFS membership:

```
Membership lost - withdrawing from cluster
```

You must allow CXFS membership for the local node in this situation. See "Allow Membership of the Local Node with the GUI" on page 232 or "Enable a Node with `cxfs_admin`" on page 279.

License Key Error

You will see the following error if you try to install CXFS on a server-capable node without a valid license key already in place:

```
Preparing... #####
[100%]
  1:cxfs_cluster #####
[100%]
cxfs          0:off 1:off 2:off 3:on  4:off 5:on  6:off
cluster_cx-exitop: Added CXFS keys to /var/cluster/cdb/cdb.db
cluster_cx-exitop: Added CXFS administration access keys to
/var/cluster/cdb/cdb.db
cxfs license check failed - use '/usr/cluster/bin/cxfslicense -d' for
details

* * * * * I M P O R T A N T * * * * *

CXFS is not properly licensed for this host.  Run
  '/usr/cluster/bin/cxfslicense -d'
for more detailed license information.
After fixing the license, please run
  '/bin/true; /etc/init.d/cxfs_cluster restart'.

cluster_cx-exitop: success
```

If you see the following message in the `/var/cluster/ha/log/clconfd_nodename` logfile, it means that the CXFS license key was not properly installed:

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

for detailed failure information.

If you do not have the CXFS license key properly installed, you will see an error on the console when trying to run CXFS. For example, on an SGI ProPack node:

```
Cluster services: CXFS not properly licensed for this host. Run
    '/usr/cluster/bin/cxfslicense -d'
for detailed failure information. After fixing the
license, please run '/etc/init.d/cxfs_cluster restart'.
```

An error such as the following example will appear in the system log file:

```
Mar  4 12:58:05 6X:typhoon-q32 crsd[533]: <<CI> N crs 0> Crsd restarted.
Mar  4 12:58:05 6X:typhoon-q32 clconfd[537]: <<CI> N clconf 0>
Mar  4 12:58:05 5B:typhoon-q32 CLCONFD failed the CXFS license check. Use the
Mar  4 12:58:05 5B:typhoon-q32    '/usr/cluster/bin/cxfslicense -d'
Mar  4 12:58:05 5B:typhoon-q32 command to diagnose the license problem.
```

If the `clconfd` daemon dies right after it starts up, this error may be present.

You must properly install the license key before you can use CXFS. See Chapter 4, "CXFS License Keys" on page 75.

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, first try the steps in "Eliminate a Residual Cluster" on page 440. If that does not work, reboot the nodes that have 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.

System Log File Errors

CXFS logs both normal operations and critical errors to the system log file, as well as to individual log files for each log group.

The system log files are:

- IRIX: `/var/adm/SYSLOG`
- SGI ProPack: `/var/log/messages`

In general, errors in the system log file file take the following form:

timestamp priority_&_facility : hostname process[ID]: <internal_info> CODE message_text

For example:

```
Sep  7 11:12:59 6X:cxfs0 cli[5830]: < E clconf 0> CI_IPCERR_NOSERVER, clconf
ipc: ipcclnt_connect() failed, file /var/cluster/ha/comm/clconfd-ipc_cxfs0
```

Table 19-2 shows the parts of the preceding message.

Table 19-2 System Log File Error Message Format

Content	Part	Meaning
Sep 7 11:12:59	Time Stamp	September 7 at 11:12 AM.
6X	Facility and level	6X indicates an informational message. See <code>syslogd</code> and the file <code>/usr/include/sys/syslog.h</code> .
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 Error Messages

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 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; see "Stopping and Restarting Cluster Administration Daemons" on page 489.

clconfd Error Messages

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

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

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

If you get an error or the dump is empty, re-create the database; for more information, see "Clearing the Cluster Database" on page 485.

If the database exists, restart the cluster daemons; see "Stopping and Restarting Cluster Administration Daemons" on page 489.

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; see "Stopping and Restarting Cluster Administration Daemons" on page 489.

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 CXFS GUI or `cxfs_admin` to modify the node's definition and add reset information. System reset configuration is recommended for all potential metadata servers. See "Define a Node with the GUI" on page 208, or "Create or Modify a Node with `cxfs_admin`" on page 270.

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` 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 the `mount` command are usually issued just before this message and provide more information about the reason of the failure.

```
CI_FAILURE, CMD(/sbin/clmount -o 'server_list=(cxfs2,cxfs0)'  
/dev/xvm/strip4 /xvm/strip4): exited with status 1 (0x1)
```

You have tried to mount a filesystem without first running `mkfs`. You must use `mkfs` to construct the filesystem before mounting it. For more information, see the `mkfs` man page.

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

There was a problem accessing the cluster database. Retry the operation. If the error persists, stop and restart the cluster daemons; see "Stopping and Restarting Cluster Administration Daemons" on page 489.

If the problem persists, clear the database, reboot, and re-create the database. See "Clearing the Cluster Database" on page 485.

```
CI_FAILURE, Exiting, monitoring agent should revive me.
```

The daemon 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 CXFS client node that is no longer in the cluster database. If you can identify the CXFS client 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 server node that is no longer in the cluster database. If you can identify the server 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_NOSEVER, 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.

CI_FAILURE, | > > SGI_CMS_CONFIG_ID_AUX_V2 error 22 (Invalid argument)

CI_FAILURE, | > > clconfd_kernel_config_thread: failed to update kernel config - retrying in 1 | > > second(s)

The previous configuration change has not fully propagated across the cluster and clconfd keeps trying until it succeeds. Possible causes include the following:

- The cxfss_client daemon is hung or is no longer running on one or more client-only nodes
- The clconfd daemon is hung or is no longer running on one or more administration nodes
- The cluster recovery is hung
- The local node is currently trying to join the cluster
- Other membership problems

If problems continue, you could try restarting cluster services.

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=(cxfs2,cxfs0)'
/dev/xvm/strip4 /xvm/strip4): /dev/xvm/strip4: Invalid
argument
```

You have tried to mount a filesystem without first running `mkfs`. You must use `mkfs` to construct the filesystem before mounting it. For more information, see the `mkfs` man page.

```
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 Error Messages

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 database. This can happen if you start cluster processes without creating the database. See "Recreating the Cluster Database" on page 489.

```
CI_ERR_RETRY, Could not find machine listing.
```

The `crsd` daemon could not find the local node in the cluster 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 `sgi-crsd` entry in the file. For information about adding this entry, see "`/etc/services` on CXFS Administration Nodes" on page 128.

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 "`/etc/services` on CXFS Administration Nodes" on page 128.

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.

CI_ERR_INTR, BMC is busy, delaying 5 seconds. Attempt 1 of 5.

The `crsd` daemon was unable to contact the baseboard management controller (BMC) of the system being reset. There will be 5 attempts to connect. You can ignore this message if the connection is successful upon a subsequent attempt. If the reset is not successful after all 5 attempts, see "IPMI Issues" on page 455.

cmond Error Messages

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

Could not register for notification.`cdb_error = 7`

An error number of 7 indicates that the cluster database was not initialized when the cluster process was started.

This may be caused if you execute the `cdbreinit` on one CXFS administration node while some other CXFS administration nodes in the pool are still running `fs2d` and already have the node listed in the database.

Do the following:

1. Execute the following command on the nodes that show the error:

```
# /usr/cluster/bin/cdb-init-std-nodes
```

This command will recreate the missing nodes without disrupting the rest of the database.

2. If the error persists, force the daemons to restart by executing the following command on IRIX:

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

On SGI ProPack:

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

Verify that cmond is restarted.

3. If the error persists, reinitialize the database on just the node that is having problems.

4. If the error still persists, reinitialize all nodes in the cluster.

See "Recreating the Cluster Database" on page 489.

```
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 (nonzero) status. Look at the corresponding daemon logs for error messages.

cxfslicense Error Message

The following message will be output by the `cxfslicense -d` command if you execute it before rebooting the system:

```
error reading kernel XVM cluster mirror status. Check if XVM module is started.
```

After you reboot the system and therefore load the XVM module, this message will no longer appear when you run `cxfslicense -d`.

fs2d Error Messages

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 database. Retry the operation. If the error persists, stop and restart the cluster daemons; see "Stopping and Restarting Cluster Administration Daemons" on page 489.

If the problem persists, clear the database, reboot, and re-create the database. See "Clearing the Cluster Database" on page 485.

```
Error 9 writing CDB string value for node
#cluster#elaine#machines#cxfs2#Cellular#status
```

An internal error occurred when writing to the cluster database. Retry the operation. If the error persists, stop and restart the cluster daemons; see "Stopping and Restarting Cluster Administration Daemons" on page 489.

If the problem persists, clear the database, reboot, and re-create the database. See "Clearing the Cluster Database" on page 485.

```
Failed to update CDB for node
#cluster#elaine#Cellular#FileSystems#fs1#FSStatus
```

An internal error occurred when writing to the cluster database. Retry the operation. If the error persists, stop and restart the cluster daemons; see "Stopping and Restarting Cluster Administration Daemons" on page 489.

If the problem persists, clear the database, reboot, and re-create the database. See "Clearing the Cluster Database" on page 485.

```
Failed to update CDB for node
#cluster#elaine#machines#cxfs2#Cellular#status
```

An internal error occurred when writing to the cluster database. Retry the operation. If the error persists, stop and restart the cluster daemons; see "Stopping and Restarting Cluster Administration Daemons" on page 489.

If the problem persists, clear the database, reboot, and re-create the database. See "Clearing the Cluster Database" on page 485.

```
Machine 101 machine_sync failed with lock_timeout error
```

The `fs2d` daemon was not able to synchronize the cluster database and the `sync` process timed out. This operation will be retried automatically by `fs2d`.

```
ALERT: CXFS Recovery: Cell 0: Server Cell 2 Died, Recovering
```

The server (cell 2) died and the system is now recovering a filesystem.

General Messages

```
CI_CONFERR_NOTFOUND, Logging configuration error: could not  
read cluster database /var/cluster/cdb/cdb.db, cdb error = 3.
```

The cluster database has not been initialized. See "Recreating the Cluster Database" on page 489.

```
WARNING: Error receiving messages from cell 2 tcpchannel 1
```

There has been an error on the CXFS 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 Error Messages

CXFS maintains logs for each of the CXFS daemons. For information about customizing these logs, see "Set Log Configuration with the GUI" on page 230.

Log file messages take the following form:

```
daemon_log timestamp 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 19-3 on page 477, shows the parts in the preceding message.

Table 19-3 Log File Error Message Format

Content	Part	Meaning
cad_log	Daemon identifier	The message pertains to the cad daemon
Sep 2 17:25:06.092	Time stamp 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; see "Stopping and Restarting Cluster Administration Daemons" on page 489.

```
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; see "Stopping and Restarting Cluster Administration Daemons" on page 489.

```
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.

```
csrm_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; see "Stopping and Restarting Cluster Administration Daemons" on page 489.

```
error 3 sending event notification to client [counter: 7 info:  
0x000000021010f078]"
```

GUI process exited without cleaning up. (The counter and info numbers are internal data structures.)

cli Messages

The following are examples of messages from
`/var/cluster/ha/log/cli_hostname:`

CI_CONFERR_NOTFOUND, No machines found in the CDB.

The local node is not defined in the cluster database.

CI_ERR_INVALID, Cluster (bob) not defined

The cluster called bob is not present in the cluster database.

CI_ERR_INVALID, CLI private command: failed (Cluster (bob) not defined)

The cluster called bob is not present in the cluster database.

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

The cluster state could not be determined. Check if the clconfd daemon is running.

CI_IPCERR_NOPULSE, ipccnt_pulse_internal(): server failed to pulse

The cluster state could not be determined. Check if the clconfd daemon is running.

CI_IPCERR_NOSERVER, clconf ipc: ipccnt_connect() failed, file /var/cluster/ha/comm/clconfd-ipc_cxfs0

The local node (cxfs0) is not defined in the cluster database.

CI_IPCERR_NOSERVER, Connection file /var/cluster/ha/comm/clconfd-ipc_cxfs0 not present.

The local node (cxfs0) is not defined in the cluster database.

crsd Errors

The following are examples of messages from /var/cluster/ha/log/crsd_hostname:

```
CI_CONFERR_INVALID, Nodeid -1 is invalid.  
I_CONFERR_INVALID, Error from ci_security_init().  
CI_ERR_SYS:125, bind() failed.  
CI_ERR_SYS:125, Initialization failed, exiting.  
CI_ERR_NOTFOUND, Nodeid does not have a value.  
CI_CONFERR_INVALID, Nodeid -1 is invalid.
```

For each of these messages, either 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 cluster database. 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_CRSEERR_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. 0x10087d48 is a pointer to an internal datastructure that uniquely identifies the request while it is being handled.

fs2d Errors

The following are examples of messages from `/var/cluster/ha/log/fs2d_log`:
Failed to copy global CDB to node cxfs1 (1), error 4

There are communication problems between the local node and node cxfs1. Check the control networks of the two nodes.

```
Communication failure send new quorum to machine cxfs2 (102)
(error 6003)
```

There are communication problems 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 problems between the local node and node *cxfs2*. Check the control networks of the two nodes.

```
Outgoing RPC to hostname : NULL
```

If you see this message, check your Remote Procedure Call (RPC) setup. For more information, see the *rpcinfo*, *rpcinfo*, and *portmap* man pages.

```
fs2d - RPC machine register: rejecting quorum from machine
hostname due to that machine not responding to our poll attempts
```

This message might indicate that the NIC for the private network has not been configured or has been configured incorrectly. It also might indicate that the cable has been unplugged.

cdbreinit Error Messages

```
Thu Jun 3 16:20:45.431 cxfsopus1.example.com cbe_fs2 - cbe_create_node: cannot create new node (RPC error = 9)
libcdb          - cdb_create_node: error 9 creating child of node 0x60000000000135c0 with subkey "ifd1"
```

This error means that some nodes have not been created in the cluster database. Error 9 usually means that *fs2d* has encountered an internal error while creating that node. To fix the problem, make sure that *fs2d* is not running on any administration-capable node and rerun *cdbreinit*.

Messages During Remote Installation

If you are performing a remote IRIX installation, you may see informational messages such as the following:

```
cdb-exitop: can't run remotely - scheduling to run later
```

cxfs_admin Errors

Following are common `cxfs_admin` errors.

```
Connecting to the local CXFS server...
receiving conflicting bootstrap packets from cluster(s) - cannot identify
server to connect to
gave up trying to connect to server
FATAL: exiting on fatal error
```

The `cxfs_admin` command can see multiple clusters. Reconfigure your network so that each cluster's private network subnet is independent of the private network subnet of other clusters. If you have multiple clusters using the same public network as the backup CXFS metadata network, use the `-i` option to identify the cluster name. See "Accessing the Correct Cluster at a Multiple-Cluster Site" on page 269.

```
Connecting to the CXFS server for the "mycluster" cluster...
Error returned from server: authorization error (8)
Inappropriate privileges to connect to the CXFS server
```

The host can see the cluster, but does not have permission to connect to it. Use the `access` command to give permission to connect.

```
Connecting to the CXFS server for the "mycluster" cluster...
Error returned from server: permissions error (9)
Insufficient privileges to acquire the administration lock
```

The host only has monitoring privileges and no administration privileges. Use the `permission=admin` attribute with the `access` command to grant the host administration rights, or use `-r` on the `cxfs_admin` command line.

```
Connecting to the CXFS server for the "mycluster" cluster...
not receiving bootstrap packets from any cluster - cannot identify server to connect to
gave up trying to connect to server
FATAL: exiting on fatal error
```

The host is not on the CXFS metadata private network and has not been granted explicit access to the cluster. Grant the host access by using the `access` command from a server-capable node or another host with admin access to the cluster.

Mount Errors

The following error indicates that one of the LUNs in this volume is inaccessible. A GPT-labeled LUN in the volume may cause this if GPT labels are not supported on the system:

```
# /sbin/mount -t cxfs -o 'client_timeout=30s,retry=0,server_list=(server1,server2)' \
/dev/cxvm/strip93 /mnt/strip93
cxfs.util get_subvol_stripe: open(/dev/rcxvm/strip93) returned -1, errno 19 (Operation not supported by device)
cxfs.util get_subvol_stripe: Some of the volumes needed for /dev/rcxvm/strip93 may have a main path that runs through
cxfs.util set_xfs_args: get_subvol_stripe failed
cxfs.util mount_main: set_xfs_args failed
```

For information about what platforms support GPT labels, see the release notes.

SYSLOG credid Warnings

Messages such as the following in the SYSLOG indicate that groups from another node are being dropped, and you may not be able to access things as expected, based on group permissions (line breaks added here for readability):

```
May 1 18:34:42 4A:nodeA unix: WARNING: credid_bundle_import: received cred for uid 5778 with 23 groups when \
configured for only 16 groups. Extra groups dropped.
May 1 18:34:59 4A:nodeB unix: WARNING: credid_getcred: received cred for uid 5778 with 23 groups when \
configured for only 16 groups. Extra groups dropped.
May 1 18:35:44 4A:nodeB unix: WARNING: credid_getcred: received cred for uid 5778 with 23 groups when \
configured for only 16 groups. Extra groups dropped.
May 1 18:36:29 4A:nodeA unix: WARNING: credid_bundle_import: received cred for uid 5778 with 23 groups \
when configured for only 16 groups. Extra groups dropped.
May 1 18:38:32 4A:nodeA unix: WARNING: credid_bundle_import: received cred for uid 5778 with 23 groups \
when configured for only 16 groups. Extra groups dropped.
May 1 18:38:50 4A:nodeB unix: WARNING: credid_getcred: received cred for uid 5778 with 23 groups when \
configured for only 16 groups. Extra groups dropped.
May 1 18:39:32 4A:nodeB unix: WARNING: credid_getcred: received cred for uid 5778 with 23 groups when \
configured for only 16 groups. Extra groups dropped.
May 1 18:40:13 4A:nodeB unix: WARNING: credid_getcred: received cred for uid 5778 with 23 groups when \
configured for only 16 groups. Extra groups dropped.
May 1 18:40:35 4A:nodeA unix: WARNING: credid_bundle_import: received cred for uid 5778 with 23 groups \
when configured for only 16 groups. Extra groups dropped.
May 1 19:04:52 4A:nodeA unix: WARNING: credid_bundle_import: received cred for uid 6595 with 21 groups \
when configured for only 16 groups. Extra groups dropped.
May 1 19:38:58 4A:nodeA unix: WARNING: credid_bundle_import: received cred for uid 6595 with 21 groups \
```

when configured for only 16 groups. Extra groups dropped.

The IRIX `ngroups_max` static system tunable parameter specifies the maximum number of multiple groups to which a user may simultaneously belong. You should increase the number of groups by running the following command and then rebooting:

```
irix# systune ngroups_max value
```

Corrective Actions

This section covers the following corrective actions:

- "Restarting CXFS Services" on page 484
- "Clearing the Cluster Database" on page 485
- "Rebooting" on page 486
- "Recovering a Two-Node Cluster" on page 486
- "Rebooting without Rejoining the Cluster" on page 488
- "Stopping and Restarting Cluster Administration Daemons" on page 489
- "Recreating the Cluster Database" on page 489
- "Verifying Connectivity in a Multicast Environment" on page 490
- "Perform a Power Cycle on a Node" on page 491
- "Reset a Node" on page 491

Restarting CXFS Services

If CXFS services do not restart after a reboot, it may be that the node was marked as `INACTIVE` in the cluster data base using the **Stop CXFS Services** function of the GUI, a `disable node:nodename` function of `cxfs_admin`. In this case, issuing a `/etc/init.d/cluster start` (IRIX), `/etc/init.d/cxfs_cluster start` (SGI ProPack), or `/etc/init.d/cxfs start` will not restart the services.

You must manually start CXFS services. If you use the GUI to restart the services, or enable with `cxfs_admin`, the configuration will be set so that future reboots will also restart CXFS services.

For information, see "Start CXFS Services with the GUI" on page 227 or "Enable a Node with `cxfs_admin`" on page 279.

Clearing the Cluster Database

To clear the cluster database on all of the administration nodes of the cluster, do the following, completing each step on each administration node before moving to the next step:



Caution: This procedure deletes all configuration information.

1. Enter the following on all server-capable administration nodes:

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

2. Enter the following on all administration nodes:

- IRIX:

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

- SGI ProPack:

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



Caution: Complete steps 1 and 2 on each node before moving to step 3 for any node.

3. Enter the following on all administration nodes:

```
# /usr/cluster/bin/cdbreinit
```

See also "Reboot Before Changing Node ID or Cluster ID" on page 158.

4. Enter the following on all administration nodes:

- IRIX:

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

- SGI ProPack:

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

5. Enter the following on all server-capable administration nodes:

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

See "Eliminate a Residual Cluster" on page 440, to get rid of possible stale cluster configuration in the kernel. If needed, reboot the nodes.

Rebooting

Enter the following individually on every node to reboot the cluster (other than Windows, which uses a different reboot mechanism) :

```
# reboot
```

For information about nodes running operating systems other than IRIX or SGI ProPack, see the *CXFS MultiOS Client-Only Guide for SGI InfiniteStorage*.

If you want CXFS services to restart whenever the node is rebooted, use the CXFS GUI to start CXFS services or `cxfs_admin` to enable the node. For information, see "Start CXFS Services with the GUI" on page 227 and "Enable a Node with `cxfs_admin`" on page 279.

The following are situations that may require a rebooting:

- If some CXFS clients are unable to unmount a filesystem because of a busy vnode and a reset of the node does not fix the problem, you may need to reboot every node in the cluster
- If there is no recovery activity within 10 minutes, you may need to reboot the node

Recovering a Two-Node Cluster

Suppose the following:

1. You have cluster named `clusterA` that has two server-capable nodes and there is no CXFS tiebreaker:
 - `node1`
 - `node2`

2. node1 goes down and will remain down for a while.
3. node2 recovers and clusterA remains up.

Note: An existing cluster can drop down to 50% of the remaining server-capable nodes **after** the initial CXFS kernel membership is formed. For more information, see "CXFS Kernel Membership, Quorum, and Tiebreaker" on page 512.

4. node2 goes down and therefore clusterA fails.
5. node2 comes back up. However, clusterA cannot form because the initialization of a cluster requires either:
 - **More than 50%** of the server-capable nodes
 - 50% of the server-capable nodes, **one of which is the CXFS tiebreaker**

To allow node2 to form a cluster by itself, you must do the following:

1. Set node2 to be the CXFS tiebreaker node, using the GUI or `cxfs_admin`:
 - See "Set Tiebreaker Node with the GUI" on page 229.
 - "Create or Modify a Cluster with `cxfs_admin`" on page 282
2. Revoke the CXFS kernel membership of node2:
 - See "Revoke Membership of the Local Node with the GUI" on page 232.
 - See "Disable a Node with `cxfs_admin`" on page 279.
3. Allow CXFS kernel membership of node2:
 - See "Allow Membership of the Local Node with the GUI" on page 232.
 - "Enable a Node with `cxfs_admin`" on page 279.
4. Unset the CXFS tiebreaker node capability.



Caution: All two-server-capable node clusters without a tiebreaker set must have fencing or reset configured. SGI recommends reset.

See:

- "Set Tiebreaker Node with the GUI" on page 229
- "Create or Modify a Node with `cxfs_admin`" on page 270

The cluster will attempt to communicate with the `node1` because it is still configured in the cluster, even though it is down. Therefore, it may take some time for the CXFS kernel membership to form and for filesystems to mount.

Rebooting without Rejoining the Cluster

The following arguments to `chkconfig` control the other cluster administration daemons and the replicated cluster database:

- IRIX: `cluster`
- SGI ProPack: `cxfs_cluster`

If they are 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 the argument off and booting in multiuser mode:

- IRIX:

```
irix# init 1
irix# /etc/chkconfig cluster off
irix# init 2
```
- SGI ProPack:

```
[root@linux root]# init 1
[root@linux root]# /bin/chkconfig cxfs_cluster off
[root@linux root]# init 3
```

For more information, see "CXFS `chkconfig` Arguments" on page 320.

Stopping and Restarting Cluster Administration Daemons

The commands to stop and restart cluster administration daemons depends upon the platform. See also "Restarting CXFS Services" on page 484. For general information about the daemons, see "Daemons" on page 495.

To stop and restart cluster administration daemons, enter the following:

- On administration nodes:
 - IRIX:

```
# /etc/init.d/cluster stop
# /etc/init.d/cluster start
```
 - SGI ProPack:

```
# /etc/init.d/cxfs_cluster stop
# /etc/init.d/cxfs_cluster start
```
- On client-only nodes:

```
# killall cxfs_client
# /etc/init.d/cxfs_client start
```

Note: You could also use the `restart` option to stop and start.

These commands affect the cluster administration daemons only.



Caution: When the cluster administration daemons are stopped, the node will not receive database updates and will not update the kernel configuration. This can have very unpleasant side effects. Under most circumstances, the administration daemons should remain running at all times. Use these commands only as directed.

Recreating the Cluster Database

To recreate the initial cluster database, do the following:

1. Ensure that the database membership quorum is held by nodes with a good database, in order to avoid propagating a bad database.

2. Enter the following:

```
# /usr/cluster/bin/cdbreinit
```

Note: See also "Reboot Before Changing Node ID or Cluster ID" on page 158.

Verifying Connectivity in a Multicast Environment

To verify general connectivity in a multicast environment, you can execute a `ping` command on the `224.0.0.1` IP address.

To verify the CXFS heartbeat, use the `224.0.0.250` IP address, which is the default CXFS heartbeat multicast address (because it is the default, this address does not have to appear in the `/etc/hosts` file).

Note: A node is capable of responding only when the administration daemons (`fs2d`, `cmnd`, `cad`, and `crsd`) or the `cxfs_client` daemon is running.

For example, to see the response for two packets sent from IRIX IP address `163.154.17.49` to the multicast address for CXFS heartbeat and ignore loopback, enter the following:

```
irixnodeA# ping -c 2 -I 163.154.17.49 -L 224.0.0.250
PING 224.0.0.250 (224.0.0.250): 56 data bytes
64 bytes from 163.154.17.140: icmp_seq=0 ttl=64 time=1.146 ms
64 bytes from 163.154.17.55: icmp_seq=0 DUP! ttl=255 time=1.460 ms
64 bytes from 163.154.17.52: icmp_seq=0 DUP! ttl=255 time=4.607 ms
64 bytes from 163.154.17.50: icmp_seq=0 DUP! ttl=255 time=4.942 ms
64 bytes from 163.154.17.140: icmp_seq=1 ttl=64 time=2.692 ms

----224.0.0.250 PING Statistics----
2 packets transmitted, 2 packets received, +3 duplicates, 0.0% packet
loss
round-trip min/avg/max = 1.146/2.969/4.942 ms
```

The above output indicates that there is a response from the following addresses:

```
163.154.17.140
163.154.17.55
163.154.17.52
```

163.154.17.50

To override the default address, you can use the `-c` and `-m` options or make the name `cluster_mcast` resolvable on all nodes (such as in the `/etc/hosts` file). For more information, see the `cxfs_client` man page.

Perform a Power Cycle on a Node

When CXFS is running, you can perform a powercycle on a node with the following command:

```
admin powerCycle node nodename
```

This command uses the CXFS daemons to shut off power to the node and then restart it.

You can perform a powercycle on a node in a cluster even when the CXFS daemons are not running by using the `standalone` option:

```
admin powerCycle standalone node nodename
```

The above command does not go through the `crsd` daemon.

If the node has not been defined in the cluster database, you can use the following command line (line breaks added here for readability, but it should be all on one line):

```
admin powerCycle dev_name port|IP_address_or_hostname_of_device of dev_type tty|network|ipmi  
with sysctrl_type mssc|mmsc|l2|l1|bmc
```

Reset a Node

When CXFS is running, you can reset a node with a system controller by using the following command:

```
admin reset node hostname
```

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

Even when the CXFS daemons are not running, you can reset a node with a system controller by using the `standalone` option of the `admin reset` command:

```
admin reset standalone node hostname
```

If you have defined the node but have not defined system controller information for it, you could use the following commands to connect to the system controller or reset the node:

```
admin ping dev_name port|IP_address_or_hostname_of_device of dev_type tty|network|ipmi with sysctrl_type msc|mmsc|I2|I1|bmc
admin reset dev_name port|IP_address_or_hostname_of_device of dev_type tty|network|ipmi with sysctrl_type msc|mmsc|I2|I1|bmc
```

For more information about the command elements, see "Define a Node with `cmgr`" on page 577.

The above command does not go through the `crsd` daemon.

Reporting Problems to SGI

When reporting a problem about a CXFS node to SGI, you should retain the information discussed in this section, depending upon the circumstances you experience.

Reporting IRIX Problems

Retain the following information for IRIX nodes:

- If a panic has occurred on an IRIX node, retain the system core files in `/var/adm/crash`, including the following:

```
analysis.number
unix.number
vmcore.number.comp
```

- For any type of problem:
 - Run the `/usr/cluster/bin/cxfsdump` utility on an IRIX node and retain the output. You can run this utility immediately after noticing a problem. The `cxfsdump` utility attempts to collect information from all nodes in the cluster by using the `rsh` command.
 - Determine the Fibre Channel HBA World Wide name mapping:

```
scsiha -w bus#
```

- Gather output from the following commands:

```
/usr/bin/hinv  
/usr/sbin/topology
```

Reporting SGI ProPack Problems

Retain the following information for SGI ProPack nodes:

- The kernel you are running:

```
[root@linux root]# uname -a
```

- The CXFS packages you are running:

```
[root@linux root]# rpm -q cxfs_client sgi-cxfs-kmp cxfs_utils cxfs-xvm-cmds
```

- The number and types of processors in your machine:

```
[root@linux root]# cat /proc/cpuinfo
```

- The hardware installed on your machine:

```
[root@linux root]# /sbin/lspci
```

- Modules that are loaded on your machine:

```
[root@linux root]# /sbin/lsmmod
```

- The `/var/log/cxfs_client` log file
- Any messages that appeared in the system logs immediately before the system exhibited the problem.
- Output about the cluster obtained from the `cxfsdump` utility run on an administration node.
- After a system kernel panic, the debugger information from the KDB built-in kernel debugger. See "Kernel Status Tools" on page 434.
- Fibre Channel HBA World Wide name mapping:

```
cat /sys/class/fc_transport/bus_ID/node_name
```

For example:

```
cat /sys/class/fc_transport/11:0:0:0/node_name
```

The *bus_ID* value is the output of `hwinfo --disk` in the SysFS `BusID` field.

- Output from the following commands:
 - Information from the following files:

```
/var/log/messages
/var/log/cxfs_client      (for client-only nodes)
/var/cluster/ha/log/*    (for administration nodes)
/etc/failover.conf       (for XVM failover version 1)
/etc/failover2.conf      (for XVM failover version 2)
/etc/hosts
/proc/discontig
```
 - Output from the following commands:

```
/usr/cluster/bin/cdbutil gettree '#'
/usr/bin/hinv
/usr/bin/topology
/sbin/xvm show -v phys
/sbin/xvm show -top -v vol
/bin/netstat -ia
```
- When a CXFS daemon or command aborts and creates core files, provide the core files and the following associated information:
 - The application that created the core file:

```
file core_filename
```
 - The binaries listed by the following command:

```
ldd application_path
```

CXFS Software Architecture

This appendix discusses the following for administration nodes:

- "Daemons"
- "Communication Paths" on page 498
- "Communication Paths in a Coexecution Cluster" on page 503
- "Flow of Metadata for Reads and Writes" on page 504

Also see the *CXFS MultiOS Client-Only Guide for SGI InfiniteStorage*.

Daemons

The following table lists the CXFS daemons and threads. CXFS shares with XFS the IRIX `xfsd` and SGI ProPack `xfsdatab` kernel threads to push buffered writes to disk.

If you are using a coexecution (of type `CXFS` and `FailSafe`) cluster, see the *FailSafe Administrator's Guide for SGI InfiniteStorage*, for information about FailSafe daemons.

Note: On SGI ProPack, the process names begin with a `*` (such as `[*mtcp_notify]`).

Table A-1 CXFS Daemons and Threads

Layer	Subsystem	Process	Description
CXFS daemons	cluster_services	clconfd	CXFS control daemon for administration nodes. Reads the cluster configuration from the CDB database and manages the local kernel's CXFS kernel membership services accordingly.
	cxfs_client	cxfs_client	CXFS client daemon for client-only nodes. Manages the local kernel's CXFS kernel membership services accordingly.
Cluster software infrastructure (cluster administrative processes)	cluster_admin	cad	Cluster administration daemon. Provides administration services.
	cluster_control	crsd	Node control daemon. Monitors the connection to other nodes. Has the ability to reset other nodes.
		cmond	Daemon that manages all other daemons. This process starts other processes in all nodes in the cluster and restarts them on failures.
Kernel Threads	IRIX sthreads	fs2d	Manages the database and keeps each copy in synchronization on all nodes in the pool.
		cmsd	Manages CXFS kernel membership and heartbeating. (The CXFS cmsd resides in the kernel; it differs from the IRIS FailSafe cmsd that resides in user space.)
		Recovery corpseleader	Manages recovery protocol for node. Coordinates recovery between nodes.

Layer	Subsystem	Process	Description
		dcshake	Purges idle CXFS vnodes on the CXFS client.
		cxfsd	Manages sending extent and size updates from the client to the server. This daemon (which runs on the CXFS client) takes modified inodes on the client and ships back any size and unwritten extent changes to the server.
	xthreads	mesgtcprcv	Reads messages (one per open message channel).
		mesgtcaccept	Responsible for accepting new connections.
		mesgtcpdiscovery	Responsible for monitoring and discovering other nodes.
		mesgtcpmulticast	Responsible for supplying heartbeat.

The `fs2d`, `clconfd`, and `crsd` daemons run at real-time priority. However, the `mount` and `umount` commands and scripts executed by `clconfd` are run at normal, time-shared priority.

Communication Paths

The following figures show communication paths in CXFS.

Note: The following figures do not represent the `cmnd` cluster manager daemon. The purpose of this daemon is to keep the other daemons running.

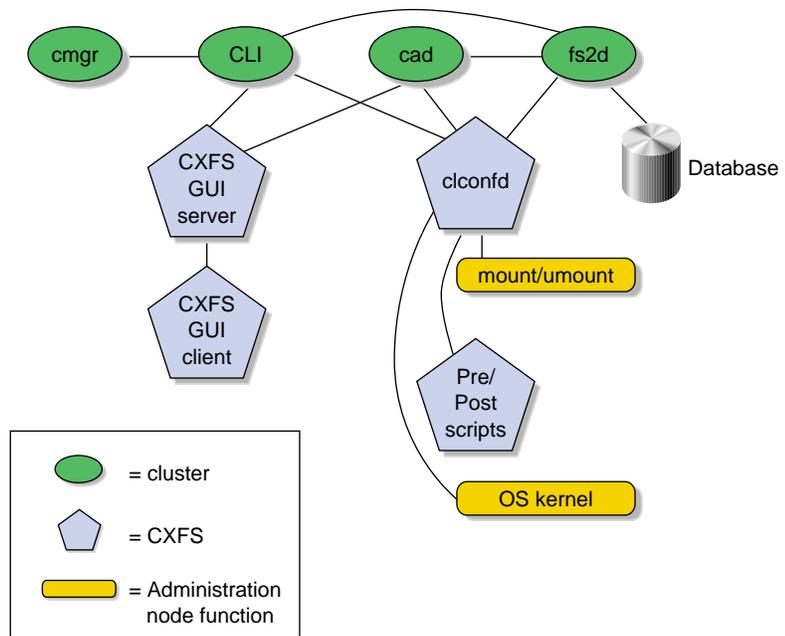


Figure A-1 Communication within One Administration Node

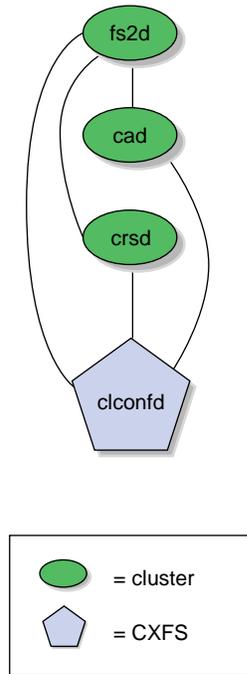


Figure A-2 Daemon Communication within One Administration Node

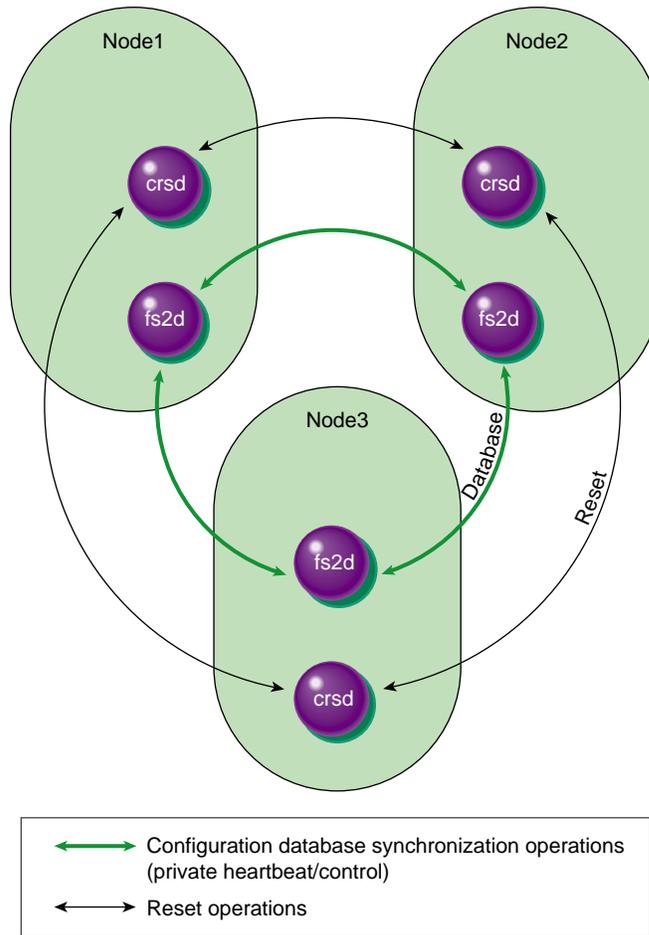


Figure A-3 Communication between Nodes in the Pool

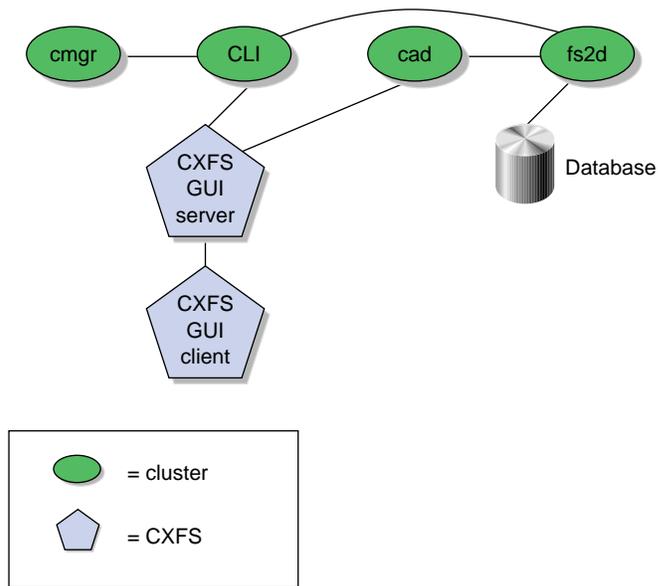


Figure A-4 Communication for an Administration Node Not in a Cluster

One of the administration nodes running the `fs2d` daemon is chosen to periodically multicast its IP address and the generation number of the cluster database to each of the client-only nodes. Each time the database is changed, a new generation number is formed and multicast. The following figure describes the communication among nodes, using a Solaris client-only node as an example.

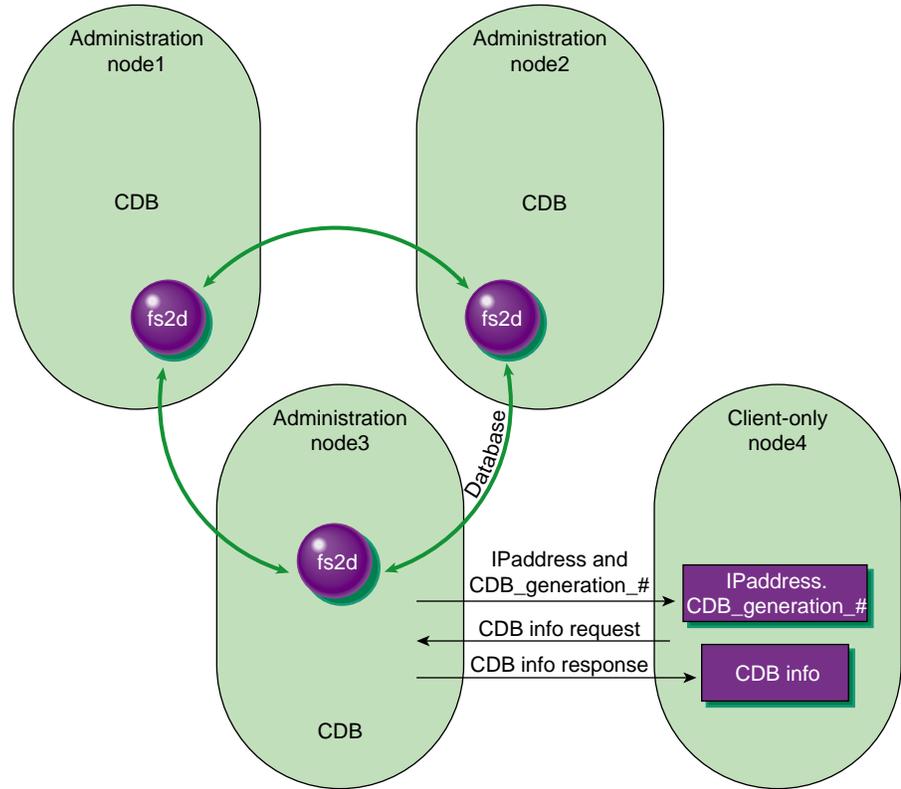


Figure A-5 Communication Among Administration Nodes and Client-Only Nodes

Communication Paths in a Coexecution Cluster

The following figures show the communication paths within one node in a coexecution cluster running CXFS and IRIS FailSafe.

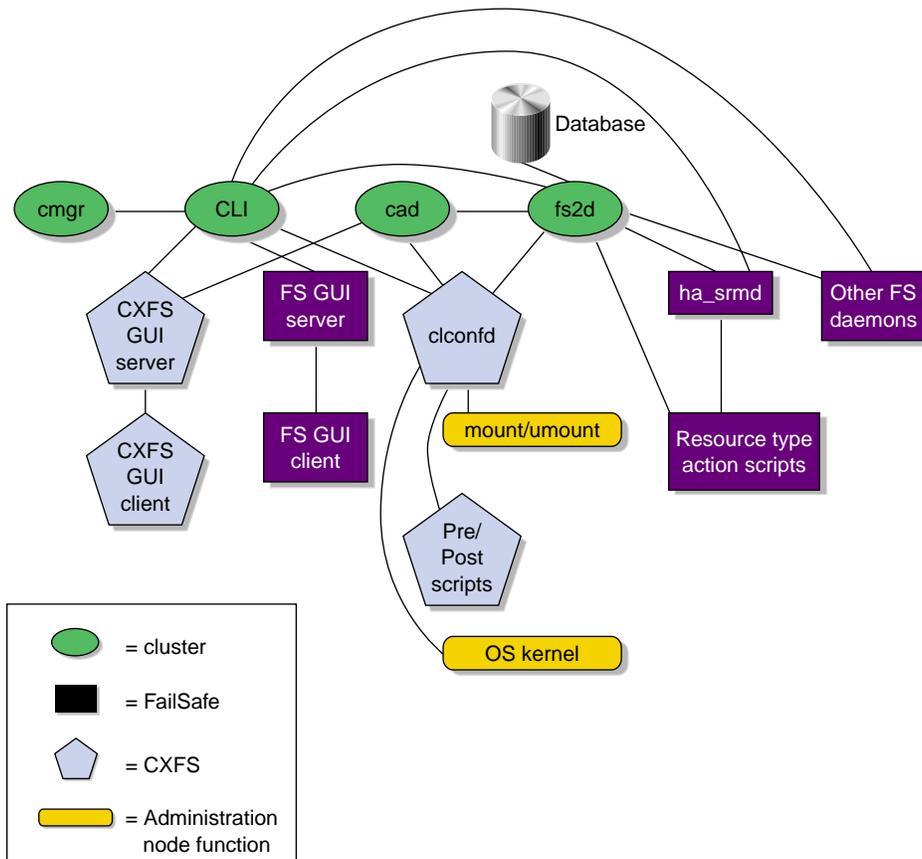


Figure A-6 Administrative Communication within One Administration Node under Coexecution

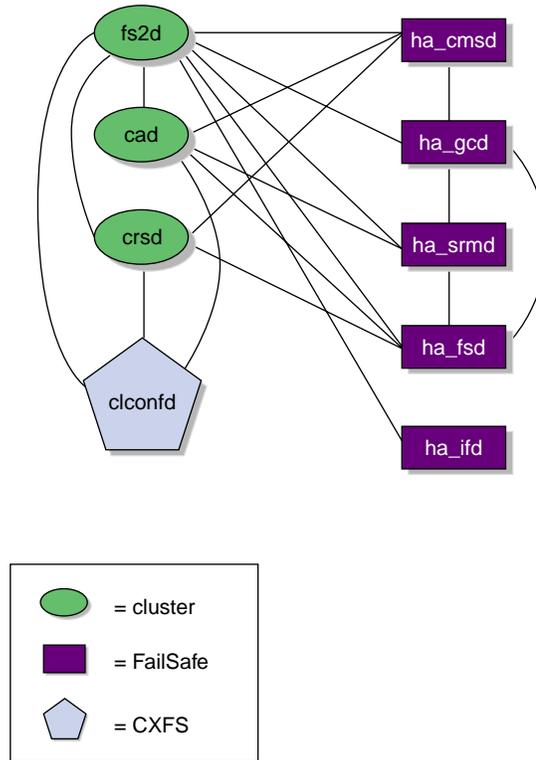


Figure A-7 Daemon Communication within One Administration Node under Coexecution

Flow of Metadata for Reads and Writes

The following figures show examples of metadata flow.

Note: A token protects a file. There can be multiple read tokens for a file at any given time, but only one write token.

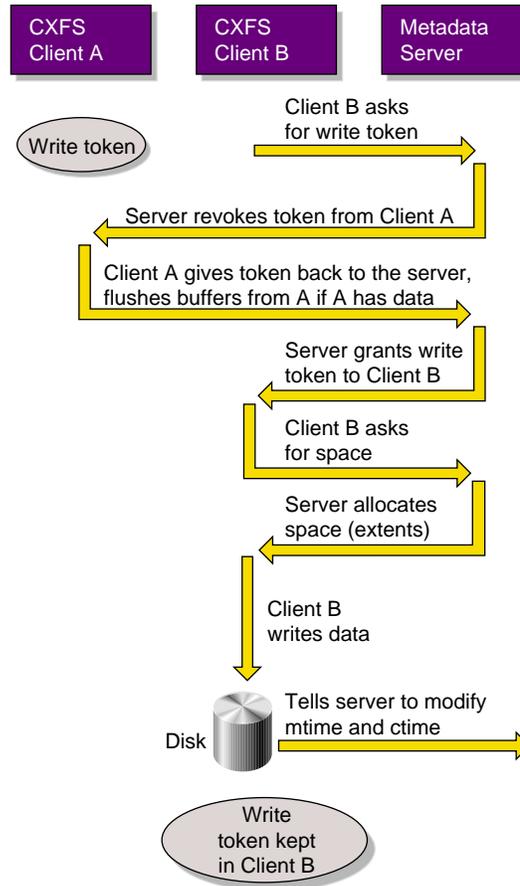


Figure A-8 Metadata Flow on a Write

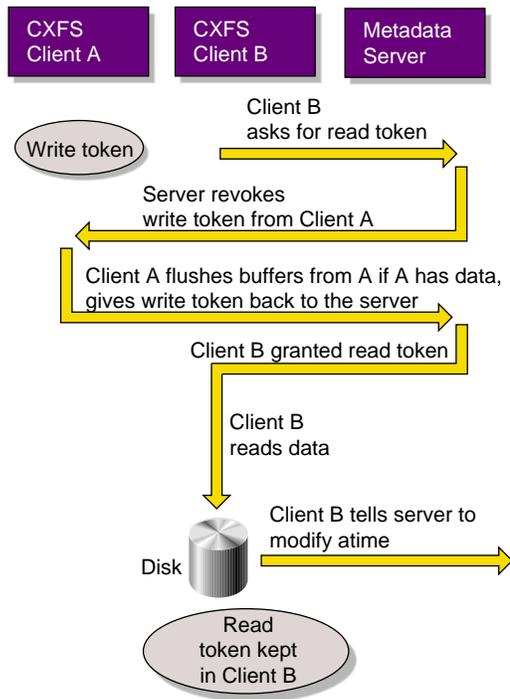


Figure A-9 Metadata Flow on a Read on Client B Following a Write on Client A

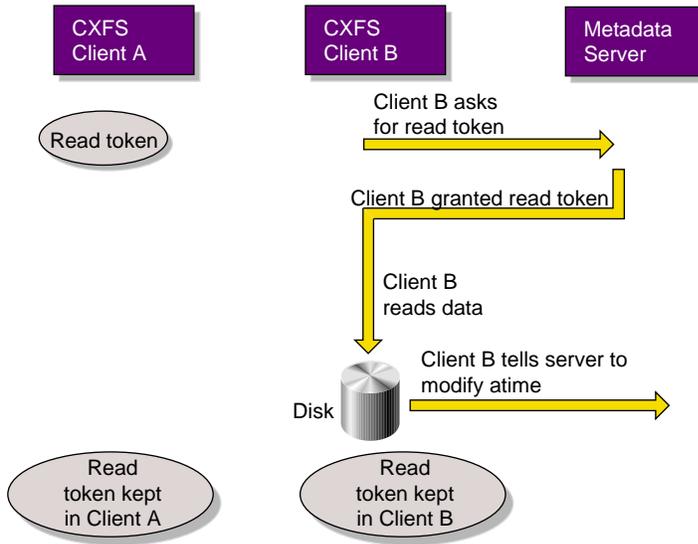


Figure A-10 Metadata Flow on a Read on Client B Following a Read on Client A

Memberships and Quorums

The nodes in a FailSafe or CXFS cluster must act together to provide a service. To act in a coordinated fashion, each node must know about all the other nodes currently active and providing the service. The set of nodes that are currently working together to provide a service is called a *membership*. Cluster activity is coordinated by a configuration database that is replicated or at least accessible on all nodes in the cluster. The cluster software sends heartbeat messages between the nodes to indicate that a node is up and running. Heartbeat messages for each membership type are exchanged via a private network so that each node can verify each membership.

Nodes within the cluster must have the correct memberships in order to provide services. This appendix discusses the different types of membership and the effect they have on the operation of your cluster.

Nodes might not be able to communicate for reasons such as the following:

- They are down
- The communication daemons have failed or have been turned off
- Software has not been configured, or has been misconfigured
- The network is misconfigured (in this case, some heartbeat messages may fail while others succeed)
- The network router or cable fails (in this case, all heartbeat messages will fail)

Nodes that cannot communicate must be excluded from the membership because the other nodes will not be able to verify their status.

It is critical that only one membership of each type exist at any one time, as confusion and corruption will result if two sets of nodes operate simultaneously but independently. There is a risk of this happening whenever a segmentation of the private network occurs, or any other network problem occurs that causes the nodes eligible for membership to be divided into two or more sets, where the nodes in each set can communicate with themselves, but not with nodes outside of the set. Thus, in order to form a membership, the nodes must have a quorum, the minimum number of nodes required to form a membership. The quorum is typically set at half the total eligible members.

For example, consider the case of six nodes eligible for a membership:

- If all six nodes can communicate with each other, they will form a membership of six and begin offering the membership's services.
- If a network segmentation occurs that causes four nodes to be in one set and two in another set, the two-node set will try to form its own membership but will be unable to do so because it does not have enough nodes to form a quorum; these nodes will therefore stop offering services. The four-node set will be able to form a new membership of four nodes and will continue to offer the membership's services.
- If a network segmentation occurs that divides the nodes into three sets of two nodes each, no set will be able to form a membership because none contains enough nodes to form a quorum. In this case, the membership services will be unavailable; this situation is unavoidable, as each set of two nodes thinks that the four other nodes may have formed a quorum, and so no set may safely offer the membership's services.
- If a network segmentation occurs that divides the nodes into two sets of three, then both could have a quorum, which could cause problems. To prevent this situation from occurring, some memberships may require a majority (>50%) of nodes or a tiebreaker node to form or maintain a membership. Tiebreaker nodes are used when exactly half of the server-capable administration nodes can communicate with each other.

The following sections provide more information about the specific requirements for membership.

Note: Because the nodes are unable to distinguish between a network segmentation and the failure of one or more nodes, the quorum must always be met, regardless of whether a partition has actually occurred or not.

Membership Types

There are three types of membership:

- "Cluster Database Membership and Quorum" on page 511
- "CXFS Kernel Membership, Quorum, and Tiebreaker" on page 512

- "FailSafe Membership, Quorum, and Tiebreaker" on page 514

Each provides a different service using a different heartbeat. Nodes are usually part of more than one membership.

Cluster Database Membership and Quorum

The nodes that are part of the the cluster database membership (also known as `fs2d membership`) work together to coordinate configuration changes to the cluster database:

- The *potential* cluster database membership is all of the administration nodes (installed with `cluster_admin` and running `fs2d`) that are defined using the GUI or `cxfs_admin` as nodes. (CXFS client-only nodes are not eligible for cluster database membership.)
- The *actual* membership is the subset of eligible nodes that are up and running and accessible to each other, as determined by heartbeats on the private network. If the primary private network is unavailable, the cluster database heartbeat will failover to the next available heartbeat network defined for the node, if any.

The cluster database heartbeat messages use remote procedure calls (RPCs). Heartbeats are performed among all nodes in the pool. You cannot change the heartbeat timeout or interval.

If a node loses its cluster database membership, the cluster database write-operations from the node will fail; therefore, FailSafe and CXFS configuration changes cannot be made from that node.

The *cluster database membership quorum* ensures atomic write-operations to the cluster database that `fs2d` replicates in all administration nodes in the pool.

The cluster database membership quorum allows an initial membership to be formed when at least half ($\geq 50\%$) of the eligible members are present. If there is a difference in the membership log between members, the cluster database tiebreaker node is used to determine which database is replicated. (See "Cluster Database Membership Logs" on page 515.) The cluster database tiebreaker node is always the administration node in the membership with the lowest node ID; you cannot reconfigure the tiebreaker for cluster database membership.

When the quorum is lost, the cluster database cannot be updated. This means that FailSafe and CXFS configuration changes cannot be made; although FailSafe and CXFS may continue to run, the loss of the cluster database quorum usually results in

the loss of quorum for FailSafe and/or CXFS, because the nodes that drop from the cluster database membership will probably also drop from other memberships.

CXFS Kernel Membership, Quorum, and Tiebreaker

The nodes that are part of the CXFS kernel membership can share CXFS filesystems:

- The *potential* CXFS kernel membership is the group of all CXFS nodes defined in the cluster and on which CXFS services have been enabled. Nodes are enabled when CXFS services are started. The enabled status is stored in the cluster database; if an enabled node goes down, its status will remain enabled to indicate that it is supposed to be in the membership.
- The *actual* membership consists of the eligible nodes on which CXFS services have been enabled and that are communicating with other nodes using the heartbeat/control network. CXFS supports only one private network, and that network is the only network used for CXFS kernel membership heartbeats (but remember that the CXFS nodes may use multiple networks for the cluster database membership heartbeats).

Note: CXFS metadata also uses the private network. The multiple heartbeats on the private network therefore reduce the bandwidth available for CXFS metadata.

During the boot process, a CXFS node applies for CXFS kernel membership. Once accepted, the node can actively share the filesystems in the cluster.

The CXFS heartbeat uses multicast. Heartbeats are performed among all CXFS-enabled nodes in the cluster.

If a node loses its CXFS kernel membership, it can no longer share CXFS filesystems.

The *CXFS kernel membership quorum* ensures that only one metadata server is writing the metadata portion of the CXFS filesystem over the storage area network:

- For the *initial* CXFS kernel membership quorum, a majority (>50%) of the server-capable administration nodes with CXFS services enabled must be available to form a membership. (*Server-capable administration* nodes are those that are installed with the `cluster_admin` product and are also defined with the CXFS GUI or `cxfs_admin` as capable of serving metadata.)

Note: Client-only nodes can be part of the CXFS kernel membership, but they are not considered when forming a CXFS kernel membership quorum. Only server-capable nodes are counted when forming the quorum.

- To *maintain* the existing CXFS kernel membership quorum requires at least half (50%) of the server-capable nodes that are eligible for membership. If CXFS kernel quorum is lost, the shared CXFS filesystems are no longer available.

No matter what the cluster components are, SGI recommends a system reset configuration on potential metadata servers to protect data integrity and improve server reliability. I/O fencing (or system reset when available) must be used on client-only nodes. See "Use a Client-Only Tiebreaker" on page 145. In clusters with an even number of potential metadata servers, a tiebreaker should be used in addition to I/O fencing or system reset; see "Isolating Failed Nodes: Failure Policies" on page 28.

You can set the CXFS tiebreaker node by using the GUI's **Set Tiebreaker Node** task or by using the `modify` command in `cxfs_admin`. See "Set Tiebreaker Node with the GUI" on page 229 and "Create a Tiebreaker with `cxfs_admin`" on page 284.

Note: If one of the server-capable nodes is the CXFS tiebreaker in a two server-capable cluster, failure of that node or stopping the CXFS services on that node will result in a cluster-wide forced shutdown. SGI recommends making a client-only node the tiebreaker to avoid losing the cluster if the tiebreaker node fails.

If I/O fencing or reset is used, the quorum is maintained by whichever side wins the reset/fence race.

If a tiebreaker node is set and the network being used for heartbeat/control is divided in half, only the group that has the CXFS tiebreaker node will remain in the CXFS kernel membership. Nodes on any portion of the heartbeat/control network that are not in the group with the tiebreaker node will exit from the membership. Therefore, if the heartbeat/control network is cut in half, you will not have an active metadata server on each half of the heartbeat/control network trying to access the same CXFS metadata over the storage area network at the same time.

Note: A tiebreaker node must be configured individually for CXFS and for FailSafe. In a coexecution cluster, these could be different nodes.

FailSafe Membership, Quorum, and Tiebreaker

The nodes that are part of the FailSafe membership provide highly available (HA) resources for the cluster:

- The *potential* FailSafe membership is the set of all FailSafe nodes that are defined in the cluster and on which HA services have been enabled. Nodes are enabled when HA services are started. The enabled status is stored in the cluster database; if an enabled node goes down, its status will remain enabled to indicate that it is supposed to be in the membership.
- The *actual* membership consists of the eligible nodes whose state is known and that are communicating with other FailSafe nodes using heartbeat and control networks. If the primary private network is unavailable, the FailSafe heartbeat will failover to the next available heartbeat network defined for the node.

The FailSafe heartbeat uses user datagram protocol (UDP). Heartbeats are performed among all FailSafe-enabled nodes in the cluster. You can change the FailSafe heartbeat timing with the GUI **Set FailSafe HA Parameters** task (the `node_timeout` parameter is the heartbeat timeout and the heartbeat is the heartbeat interval).

If a node loses its FailSafe membership, FailSafe will fail over its HA resources to another node in the cluster.

The *FailSafe membership quorum* ensures that a FailSafe resource is available only on one node in the cluster. The quorum requires that the state of a majority (>50%) of eligible nodes to be known and that half (50%) of the eligible nodes be present to form or maintain membership.

If a network partition results in a tied membership, in which there are two sets of nodes (each consisting of 50% of the potential FailSafe membership), then a node from the set containing the *FailSafe tiebreaker node* will attempt to perform a reset on a node in the other set. *Reset* is the failure action that performs a system reset via a serial line connected to the system controller.

If the node can verify that the other node was reset, then the membership will continue on the set with the tiebreaker. However, containing the tiebreaker is not a guarantee of membership; for more information, see the *FailSafe Administrator's Guide for SGI InfiniteStorage*. The default FailSafe tiebreaker is the node with the lowest node ID in the cluster.

When FailSafe membership quorum is lost, the resources will continue to run but they are no longer highly available.

Cluster Database Membership Logs

Each `fs2d` daemon keeps a *membership log* that contains a history of each database change (write transaction), along with a list of nodes that were part of the membership when the write transaction was performed. All nodes that are part of the cluster database membership will have identical membership logs.

When a node is defined in the database, it must obtain a current copy of the cluster database and the membership log from a node that is already in the cluster database membership. The method used to choose which node's database is replicated follows a hierarchy:

1. If the membership logs in the pool share a common transaction history, but one log does not have the most recent transactions and is therefore incomplete, the database from a node that has the complete log will be chosen to be replicated.
2. If there are two different sets of membership logs, the database from the set with the most number of nodes will be chosen.
3. If there are two different sets of membership logs, and each set has an equal number of nodes, then the set containing the node with the lowest node ID will be chosen.

To ensure that the complete transaction history is maintained, do not make configuration changes on two different administration nodes in the pool simultaneously. You should connect the CXFS or FailSafe GUI to a single administration node in the pool when making changes. However, you can use any node in the pool when requesting status or configuration information.

The following figures describe potential scenarios using the hierarchies.

Figure B-1 on page 517, shows:

- Time 1: An established pool of three administration nodes sharing heartbeats, with node IDs 1-3, represented by the node names N1-N3. The `fs2d` database tiebreaker node is the node in the membership with the lowest node ID. Each successive database write is identified by a letter in the membership log.
- Time 2: A new node, N4, is defined using `cxfs_admin` or the GUI connected to node N1. Node N4 (node ID = 4) joins the pool. Its membership log is empty.
- Time 3: Because N1/N2/N3 have identical membership logs, the database is replicated from one of them. In this case, N2 is randomly chosen.
- Time 4: All nodes in the pool have identical membership logs.

- Time 5: A network partition occurs that isolates N1. Therefore, N1 can no longer receive database updates. Configuration changes are made by connecting the GUI to N2 (or running `cxfs_admin` on node N2); this results in updates to the membership logs in N2, N3, and N4, but not to N1 because it is isolated.
- Time 6: The partition is resolved and N1 is no longer isolated. Because N2/N3/N4 have identical membership logs, and share the beginning history with N1, the database is replicated from one of them. N4 is chosen at random.
- Time 7: All nodes in the pool have identical membership logs.

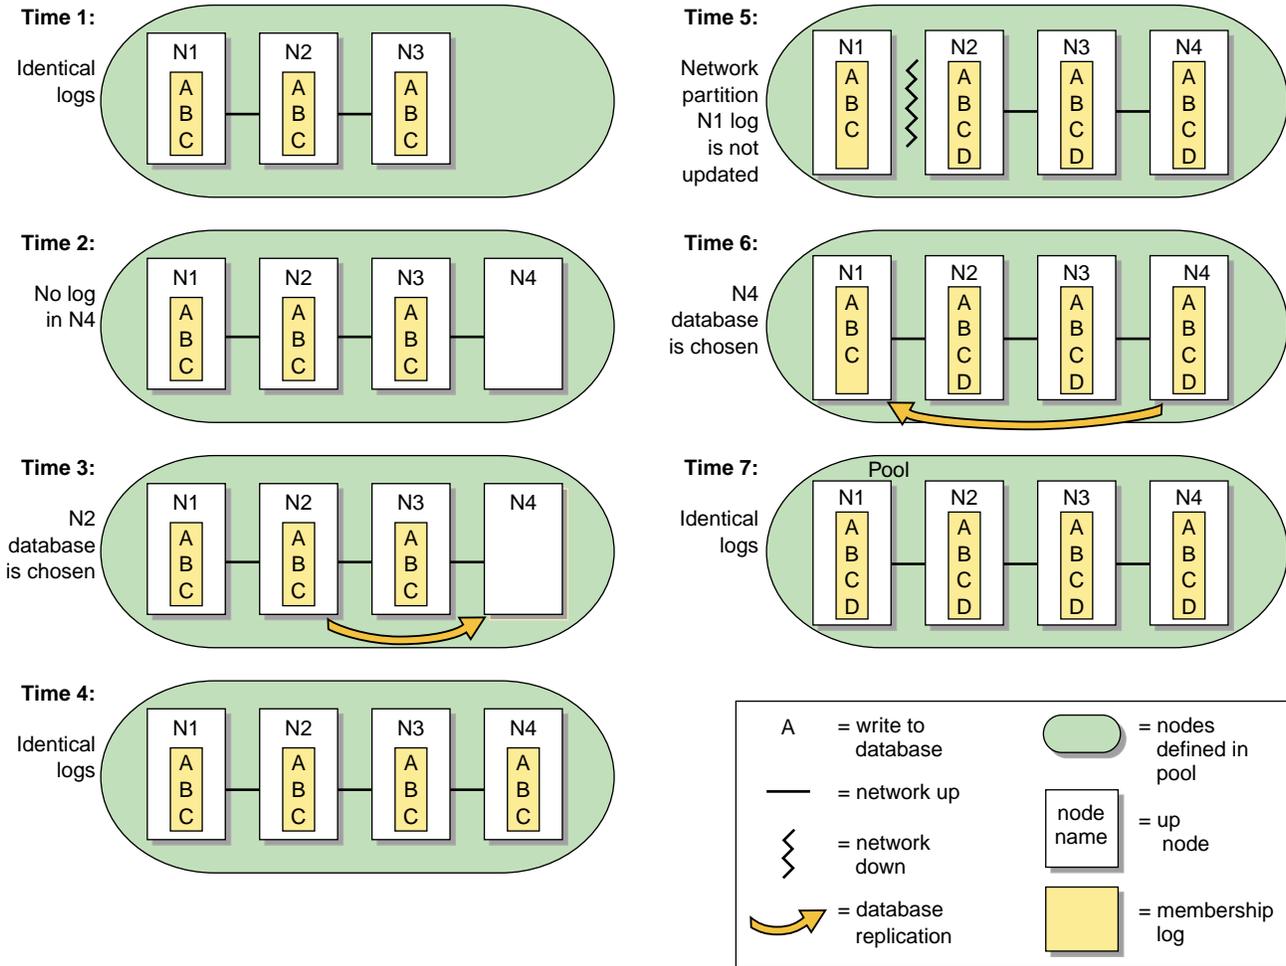


Figure B-1 One Node is Out of Date: Most Recent Log is Replicated

Recall that a node can be in only one pool at a time. If there are two separate pools, and from a node in one pool you define one or more nodes that are already in the other pool, the result will be that nodes from one of the pools will move into the other pool. **This operation is not recommended**, and determining which nodes will move into which other pool can be difficult. Figure B-2 on page 518 illustrates what to expect in this situation.

- Time 1: There are two pools that do not share membership log contents. One pool has two nodes (N1/N2), the other has three (N3/N4/N5).
- Time 2: N1 and N2 are defined as part of the second pool by running `cxfs_admin` or connecting the GUI to node N3, N4, or N5. This results in a new pool with five nodes with different membership logs.
- Time 3: The database from the larger set of nodes is the one that must be replicated. N3 is chosen at random from the N3/N4/N5 set.
- Time 4: All nodes in the pool have identical membership logs.

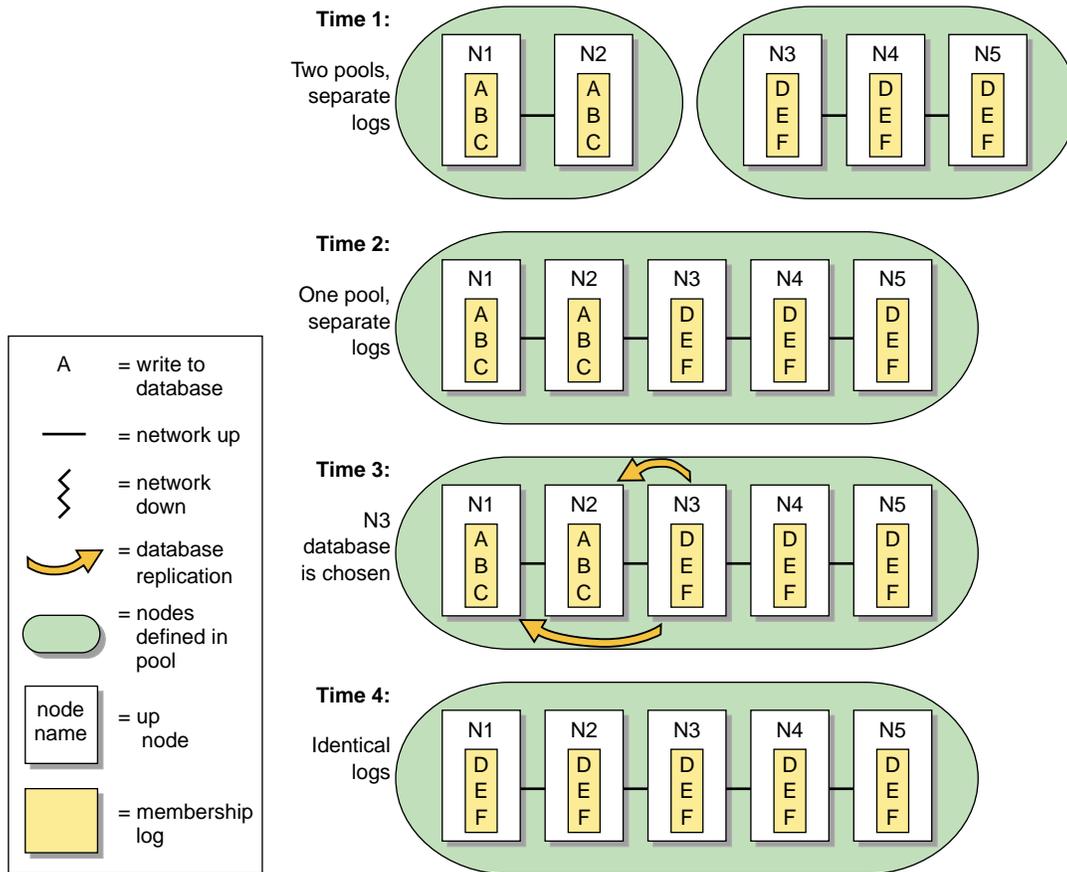


Figure B-2 Unequally Sized Pools are Joined: Log from Larger Pool is Replicated

Figure B-3 on page 520, shows a similar situation in which two nodes are defined in two pools, but the pools are of equal size:

- Time 1: There are two pools that do not share membership log contents. Each pool has two nodes (N1/N2 in pool 1, and N3/N4 in pool 2).
- Time 2: N1 and N2 are defined as part of the second pool by connecting the GUI or running `cxfs_admin` on node N3 or N4. This results in a new pool with four nodes with different membership logs.
- Time 3: Because each set has the same number of nodes, the tiebreaker node (the node with the lowest node ID in the membership) must be used to determine whose database will be chosen. Because node N1 is the lowest node ID (node ID=1), the database from N1 is chosen.
- Time 4: All nodes in the pool have identical membership logs.

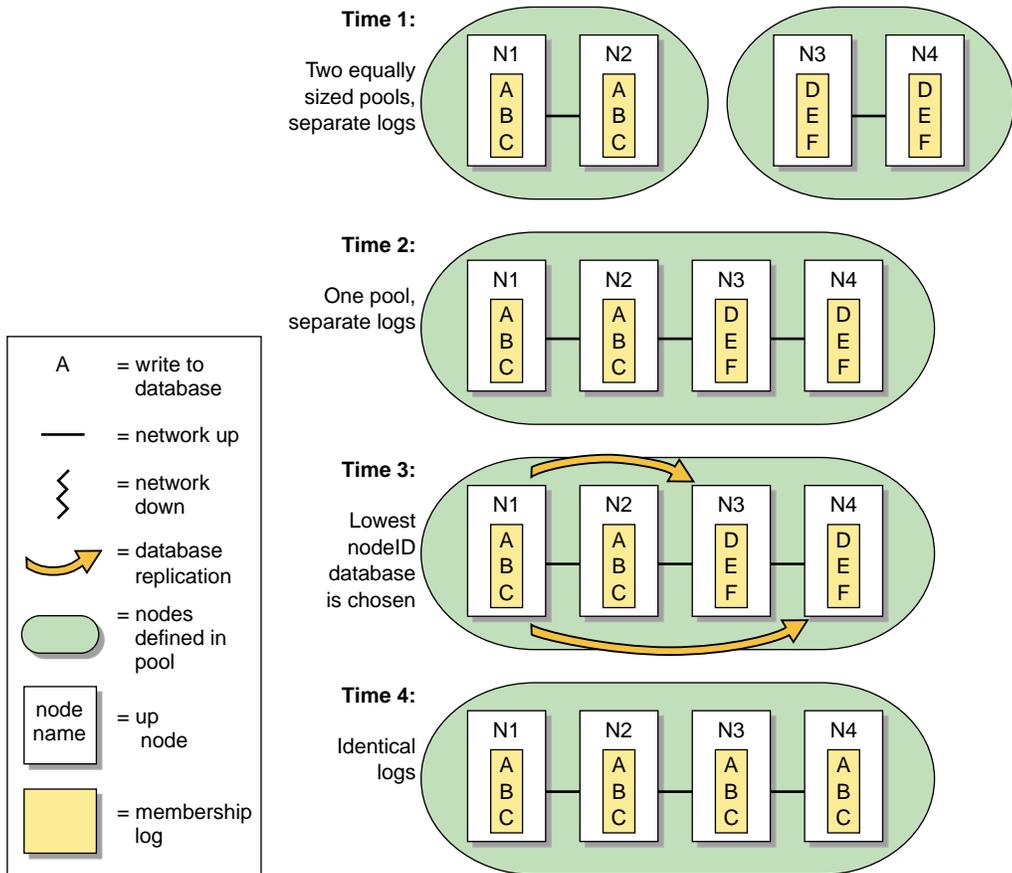


Figure B-3 Equally Sized Pools are Joined: Log from Node with Lowest Node ID is Replicated

Quorum and Tiebreaker Examples

Changing CXFS Kernel Membership Quorum Example

Figure B-4 on page 522, shows an example of a changing CXFS kernel membership quorum. It shows a pool of:

- Five CXFS server-capable administration nodes (A, B, C, D, and E)

- Two client-only nodes (F and G)
- One client admin node (H)

All nodes except E are defined as part of the cluster. Assume that CXFS services have been enabled on A, B, C, D, F, G, and H.

Of the seven nodes eligible for CXFS kernel membership, four are server-capable nodes (A, B, C, and D). Therefore, at least three of these four nodes must be able to communicate with each other to form an initial CXFS kernel quorum (>50% of the eligible server-capable nodes). Once the quorum has been reached, a membership will form with the nodes in the quorum plus all other eligible nodes that can communicate with the nodes in the quorum.

Figure B-4 on page 522, shows the following:

- Time 1: The CXFS kernel membership quorum is formed with three server-capable nodes, A, B, and C. The membership is A, B, C, F, G, and H.
- Time 2: Node B shuts down and leaves the membership. The remaining nodes in the quorum are A and C. The membership is still be available in this case because it satisfies the quorum requirement to maintain 50% of the eligible server-capable nodes (that is, two of the four server-capable nodes). The membership is A, C, F, G, and H.
- Time 3: Node A also shuts down and leaves the membership. Therefore, the quorum requirement is no longer met because quorum cannot be maintained with fewer than 50% of the eligible server-capable nodes. Without a quorum, the membership cannot continue, and so the CXFS filesystems in the cluster would not be available.

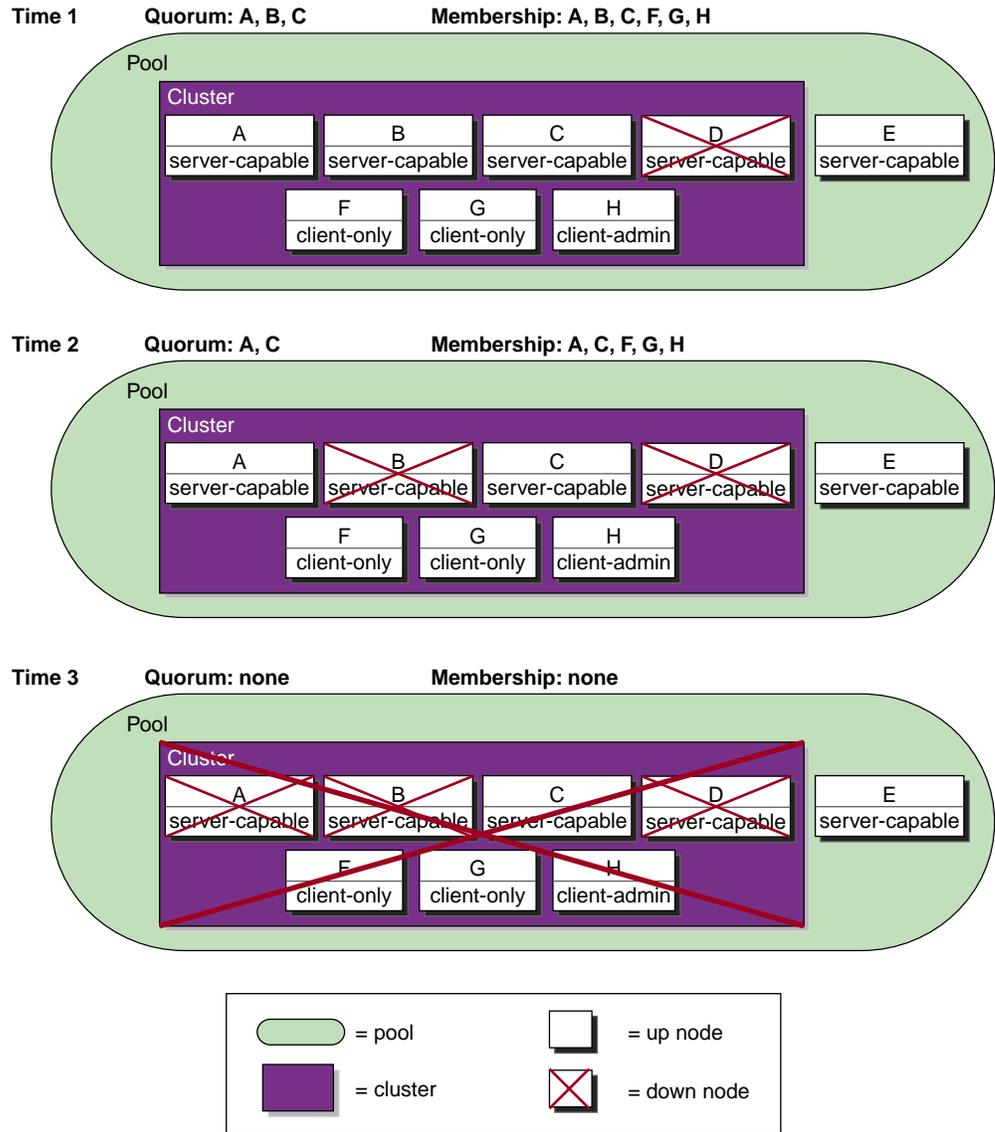


Figure B-4 Changing Quorum for CXFS Kernel Membership

Coexecution Example

Figure B-5 on page 524, shows an example of the different memberships in a cluster running CXFS and FailSafe. The pool contains 15 nodes (named N1 through N15). N1 has the lowest node ID number. There are CXFS nodes running IRIX, Solaris, and Windows; only the nodes running IRIX are administration nodes containing the cluster database. The FailSafe nodes are those where HA services are enabled; each of these is an administration node.

- Cluster database membership:
 - Eligible: N1, N2, N3, N5, N9, and N10 (that is, all nodes containing the cluster database)
 - Actual: N1, N2, N3, and N10 (because N5 and N9 are down)
 - Quorum: N1, N2, N3, and N10 (>50% of eligible nodes)
- FailSafe membership:
 - Eligible: N1, N2, and N3 (that is, those nodes with HA services enabled and defined as part of the cluster)
 - Actual: N1, N2, N3
 - Quorum: N1, N2, N3 (>50% of eligible nodes)
- CXFS kernel membership:
 - Eligible: N1-N8 and N11-N15 (N9 and N10 are not defined as part of the cluster)
 - Actual: N1, N2, N3, N4, N6, and N11-N15 (because N5, N7, and N8 are down)
 - Quorum: N1, N2 (>50% of server-capable eligible nodes)

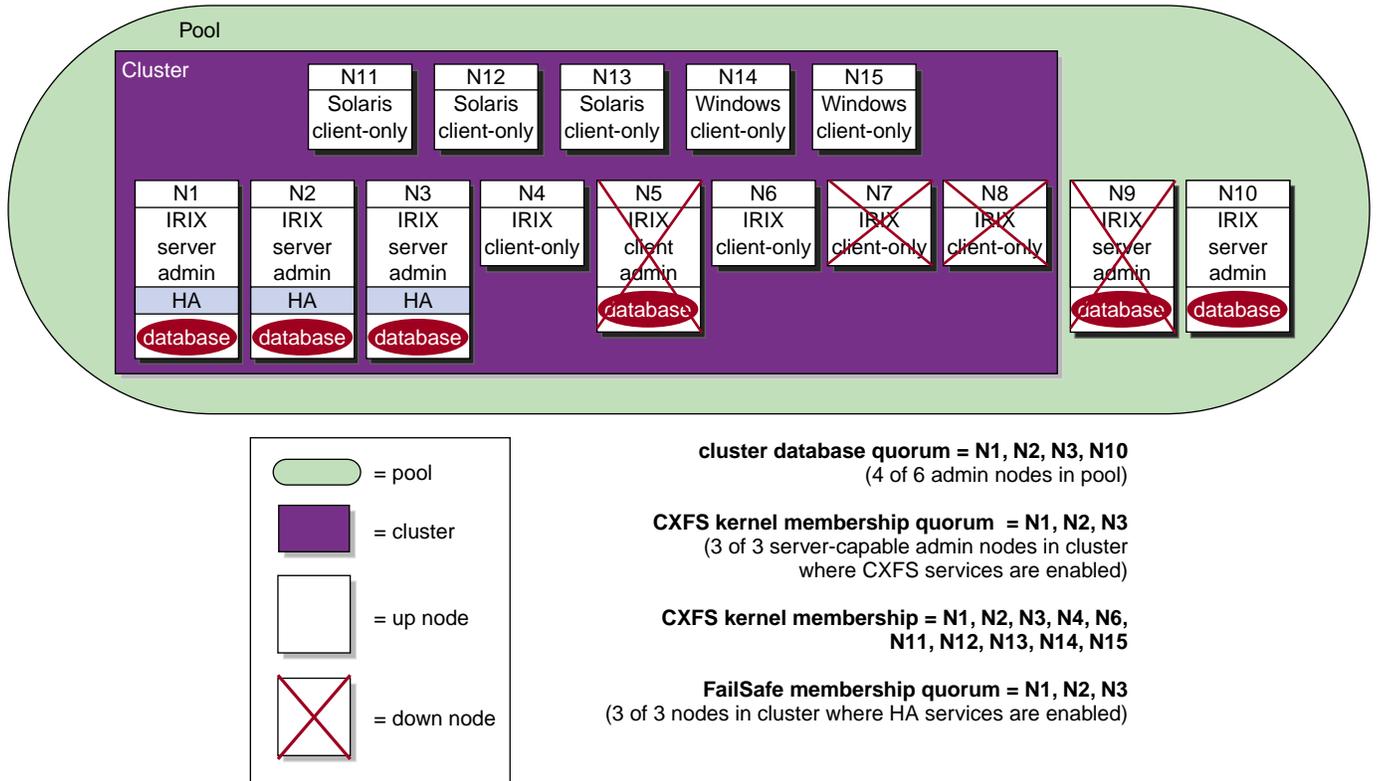


Figure B-5 Example Memberships in a Coexecution Cluster

CXFS Tiebreaker Node Example

Figure B-6 on page 525, displays a situation in which a router dies and the heartbeat/control network is effectively split in two. The potential CXFS kernel membership is defined to be nodes A, B, C, and D. The nodes on network segment 2 (nodes C and D) will leave the CXFS kernel membership because they do not contain the CXFS tiebreaker node, and therefore do not have a quorum. On network segment 1, one of the other two potential metadata servers will become active and the membership will only include the systems on network segment 1. The nodes that were on network segment 2 will remain out of the membership until CXFS services are restarted on them and the router is repaired.

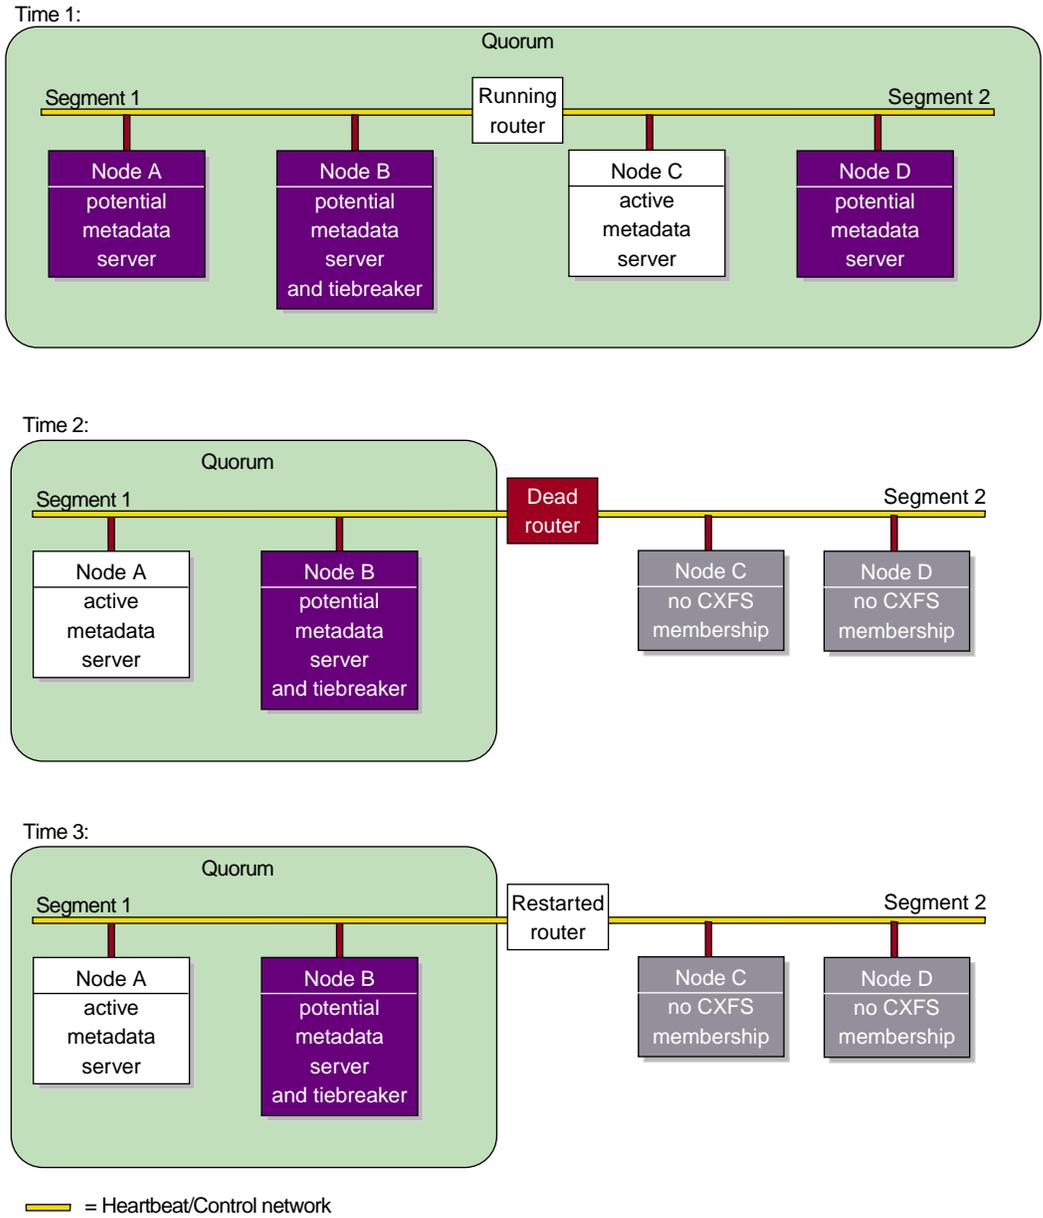


Figure B-6 CXFS Tiebreaker Node

Heartbeat Considerations

There are different heartbeats for each membership type, and each uses a different networking method. Therefore, certain network misconfiguration can cause one heartbeat to fail while another succeeds.

At least two networks should be designated as FailSafe heartbeat networks. FailSafe uses only the highest priority working network for heartbeats; the other network is for heartbeat failover. Usually the private network is used as the highest priority heartbeat network.

In a coexecution cluster, there must be two networks as required by FailSafe; at least one private network is recommended for FailSafe and a private network is required by CXFS.

In a coexecution cluster, CXFS metadata, CXFS heartbeat, and FailSafe heartbeat can use the same network. The heartbeat intervals and timeouts should be appropriately adjusted, if possible, so that all network traffic has sufficient bandwidth. You cannot change the heartbeat timeout or interval for the cluster database membership. Before you adjust the heartbeat settings for the FailSafe membership or CXFS kernel membership, you should consider the impact on the other heartbeats.

If the highest priority network fails, the FailSafe, cluster and CXFS kernel memberships will continue using the next priority network. (However, if private network failover has not been defined for CXFS, then the CXFS kernel membership will fail.)

CXFS Recovery Issues in a Cluster with Only Two Server-Capable Nodes

A cluster with an odd number of server-capable nodes is recommended for a production environment. However, if you use a production cluster with an even number of server-capable nodes (especially only two server-capable nodes), you must do one of the following:

- Use system reset configuration on potential metadata servers to protect data integrity and improve server reliability. I/O fencing (or system reset when available) must be used on client-only nodes. Clusters should have an odd number of server-capable nodes or use a client-only tiebreaker node.
- Set a CXFS tiebreaker node. If the tiebreaker node is a server-capable administration node, there will be a loss of CXFS kernel membership, and

therefore CXFS filesystems, if the tiebreaker node goes down. If the tiebreaker is an administration node, the cluster database membership may also be lost.

However, even with these methods, there are recovery and relocation issues inherent to a cluster with only two server-capable nodes.

IP Filtering for the CXFS Private Network

The method used for IP filtering for the CXFS private network varies by OS platform:

- IRIX can use `ipfilter` or `ipfilterd`
- SGI ProPack uses `SuSEfirewall2` or `iptables`

Note: The `SuSEfirewall2` utility silently makes changes to the `/proc/sys/net` system tunable parameters.

If you use I/O fencing, the configuration must allow communication between the node and the `telnet` port on the switch. For details about port use, see "CXFS Port Usage" on page 320.

For more information about these utilities, see their documentation.

Following is an example `/etc/ipfilterd.conf` file for configuring CXFS. Note the following:

- There must be an `/etc/ipfilterd.conf` file configured on each node on which you want to filter IP traffic. The files will be similar except for the first set of lines, which are node-dependent; that is, the lines in the file for NodeA must match the networking interfaces on which the network traffic may pass for NodeA.
- The `sysctl` variable `ipfilterd_inactive_behavior` must be set to 0, which means that the filter will be disabled as soon as `ipfilterd` is terminated using the `killall` command.
- The `ipfilterd` argument to `chkconfig` must be turned on for each node where `ipfilterd` will run. For example:

```
nodeA# chkconfig ipfilterd on
```

- If any network interface name is changed on a system, you must update the `/etc/ipfilterd.conf` file to include the change in the appropriate `accept` line. That is:

```
accept -i changed_or_new_interface
```

- For debugging purposes, each dropped packet will log a message similar to the following in the `syslog` file:

```
May 24 16:44:44 5A:rodin unix: NOTICE: ipfilter(cache) - packet dropped:
10.1.1.5 SPT=137 DPT=137 UDP
```

If you want to disable the filtering, such as in the case where it is blocking wanted traffic, do the following:

1. Kill the `ipfilterd` daemon:

```
nodeA# killall ipfilterd
```

2. Turn off the `ipfilterflag` argument:

```
nodeA# chkconfig ipfilterd off
```

Following is a sample file for nodeA:

```
nodeA# cat ipfilterd.conf
#
# ipfilterd.conf for nodeA
#
#
# Filters follow:
#
# Do not restrict traffic on any of the interfaces for NodeA,
# except from efl (CXFS heartbeat)
#
accept -i lo0
accept -i ef0
accept -i eg0
accept -i egl
accept -i lb0

#
# Restrict access over the CXFS heartbeat network
# Interface efl
#

# Accept any fragment, reassembly won't work if first fragment filtered out.
accept -i efl ip.off>0

# CXFS is using RPC, need portmapper.
```

```
accept -i ef1 udp.port 111
accept -i ef1 tcp.port 111

# fs2d daemon is dynamically assigning ports in range 600-1023.
# We need port definition (sport + dport for both directions).
accept -i ef1 tcp.sport>=600 and tcp.sport<=1023
accept -i ef1 tcp.dport>=600 and tcp.dport<=1023

# sgi-cad defaults to 9000/tcp
accept -i ef1 tcp.port 9000

# sgi-crsd
# Each node opens 7500/udp, both directions needed
accept -i ef1 udp.port 7500

# Uncomment the line below for CXFS client-only node.
# accept -i ef1 udp.port 5449

# CXFS kernel ports 5450-5453
# Connections in both directions so open dport and sport.
accept -i ef1 tcp.port 5450
accept -i ef1 tcp.port 5451
accept -i ef1 udp.port 5452
accept -i ef1 udp.port 5453

# fs2d client are using ports in range 7000-8500
accept -i ef1 tcp.dport>7000
accept -i ef1 udp.dport>7000

# Uncomment the line below for IO fencing only if switches are on CXFS private network
# (ip.src is the switch address)
# accept -i ef1 tcp.sport=23 and ip.src=10.1.1.6

# Let icmp traffic pass, especially 'PORT UNREACHABLE ICMP packet'
accept -i ef1 icmp
# Reject the rest (-l will log any rejected packet to the SYSLOG)
reject -i ef1 -l
```

Operating System Path Differences

This appendix lists the locations for commonly used commands. For information about other client-only operating systems, see the *CXFS MultiOS Client-Only Guide for SGI InfiniteStorage*.

Table D-1 IRIX Paths

Command/File	IRIX
chkconfig	/etc/chkconfig
cxfs_client log	/var/adm/cxfs_client
cxfs_client.options	/etc/config/cxfs_client.options
df	/usr/sbin/df
grioadmin	/usr/sbin/grioadmin
griogqs	/usr/sbin/griogqs
hinv	/sbin/hinv
hostname	/usr/bsd/hostname
mount	/sbin/mount
netstat	/usr/etc/netstat
ping	/usr/etc/ping
ps	/usr/bin/ps
scsiha	/usr/sbin/scsiha
xvm	/sbin/xvm
Cluster daemon configuration files	/etc/config/
System log	/var/adm/SYSLOG

Command/File	IRIX
CXFS/cluster daemon initialization	/etc/init.d/cluster
CXFS license verification command (server-capable nodes):	/usr/cluster/bin/cxfslicense

Table D-2 SGI ProPack Paths

Command/File	SGI ProPack
chkconfig	/bin/chkconfig
cxfs_client log	/var/log/cxfs_client
cxfs_client.options	/etc/cluster/config/cxfs_client.options
df	/bin/df
grioadmin	/usr/sbin/grioadmin
griogqs	/usr/sbin/griogqs
hinv	/usr/bin/hinv
hostname	/bin/hostname
mount	/bin/mount
netstat	/bin/netstat
ping	/bin/ping
ps	/bin/ps
updatedb	/etc/sysconfig/locate
xscsiha	(No equivalent)
xvm	/sbin/xvm
Cluster daemon configuration files	/etc/cluster/config/
System log	/var/log/messages

Command/File	SGI ProPack
CXFS/cluster daemon initialization	<code>/etc/init.d/cxfs_cluster</code>
CXFS license verification command (server-capable nodes):	<code>/usr/cluster/bin/cxfslicense</code>

Filesystem Specifications

Item	IRIX	SGI ProPack
Maximum filesystem size	2^{64} bytes (about 18 million terabytes)	2^{64} bytes
Maximum files size/offset	$2^{63}-1$ bytes (about 9 million terabytes)	$2^{63}-1$ bytes
Filesystem block size (in bytes) ¹	512, 1024, 2048, 4096, 8192, 16384, 32768, or 65536	512, 1024, 2048, 4096, 8192, or 16384
XVM device block size (in bytes)	512	512
Physical LUN limit for DVH-labeled disks	2 TB	2 TB
Physical LUN limit for GPT-labeled disks ²	2^{63} device blocks	2^{63} device blocks
Maximum concatenated slices	65536 ³	65536

¹ If the filesystem is to be accessible by other platforms in a multiOS cluster, its block size must be supported on all platforms in the cluster

² Note the following about physical LUN limits for GPT-labeled disks:

- Physical LUNs with GPT labels are not constrained by XVM or CXFS to be smaller than the largest possible filesystem.
- Cluster nodes may constrain the LUN size to be smaller due to driver or other operating system constraints. A LUN used in the cluster may not be larger than the maximum size allowed by any node.
- All nodes that mount a filesystem using LUNs larger than 2 TB must be upgraded to CXFS 4.2 or later.

³ 65536 concatenated 2-TB slices equals 130 PetaBytes

System Reset Configuration

This appendix discusses system controllers that can be used in CXFS system reset configurations:

- "BMC System Controller" on page 539
- "L2 System Controller" on page 545
- "L1 System Controller" on page 551
- "MSC System Controller" on page 553
- "MMSC System Controller" on page 556

Note: Serial cables are provided with SAN server configurations for TTY ports. Other configurations that use TTY ports require that you purchase serial cables.

BMC System Controller

SGI recommends that you use the baseboard management controller (BMC) system controller when it is available.

The BMC must not be on the primary CXFS private network. Ideally, the BMC should be on a different private network that is reachable by all server-capable nodes in the cluster. A public network is not ideal for security reasons, but is acceptable.

Altix XE systems contain an integrated BMC. CXFS uses Intelligent Platform Management Interface (IPMI) to communicate with the BMC.

To use the BMC, you must create an admin user ID and assign the BMC a static IP address. This can be done using `ipmitool(1)` on the system containing the BMC. Do the following:

1. Verify that you have `ipmitool(1)` version 1.8.9 or later (which is required for use with the `lanplus` interface):

```
[root@linux root]# ipmitool -v
ipmitool version 1.8.9
```

2. Load the following IPMI modules:

```
[root@linux root]# modprobe ipmi_msghandler
[root@linux root]# modprobe ipmi_devintf
[root@linux root]# modprobe ipmi_si
```

3. Create a user named admin with a password and ADMINISTRATOR privileges:

- a. Find the next available user ID:

```
ipmitool -o intelplus -d /dev/ipmi0 user list 1|2
```

- b. Assign the user name admin to the next available user ID:

```
ipmitool -o intelplus -d /dev/ipmi0 user set name userID admin
```

- c. Set the password for user admin:

```
ipmitool -o intelplus -d /dev/ipmi0 user set password userID admin_password
```

- d. Enable the access modes and set the privilege level to ADMINISTRATOR:

```
ipmitool -o intelplus -d /dev/ipmi0 channel setaccess 1|2 userID callin=on ipmi=on link=on privilege=4
```

Note: You must apply the privilege change separately for channel 1 and for channel 2.

- e. Verify that the correct settings were applied:

```
ipmitool -o intelplus -d /dev/ipmi0 user list 1|2
ipmitool -o intelplus -d /dev/ipmi0 channel getaccess 1|2 userID
```

For example (line breaks shown here for readability):

```
[root@linux root]# ipmitool -o intelplus -d /dev/ipmi0 user list 1
ID Name          Callin Link Auth IPMI Msg Channel Priv Limit
1           true   false   true   ADMINISTRATOR
```

```
[root@linux root]# ipmitool -o intelplus -d /dev/ipmi0 user list 2
ID Name          Callin Link Auth IPMI Msg Channel Priv Limit
1           true   false   true   ADMINISTRATOR
```

```
[root@linux root]# ipmitool -o intelplus -d /dev/ipmi0 user set name 2 admin
[root@linux root]# ipmitool -o intelplus -d /dev/ipmi0 user set password 2 password
[root@linux root]# ipmitool -o intelplus -d /dev/ipmi0 channel setaccess 1 2 callin=on \
ipmi=on link=on privilege=4
[root@linux root]# ipmitool -o intelplus -d /dev/ipmi0 channel setaccess 2 2 callin=on \
ipmi=on link=on privilege=4
```

```
[root@linux root]# ipmitool -o intelplus -d /dev/ipmi0 user list 1
ID Name Callin Link Auth IPMI Msg Channel Priv Limit
1 true false true ADMINISTRATOR
2 admin true true true ADMINISTRATOR
```

```
[root@linux root]# ipmitool -o intelplus -d /dev/ipmi0 user list 2
ID Name Callin Link Auth IPMI Msg Channel Priv Limit
1 true false true ADMINISTRATOR
2 admin true true true ADMINISTRATOR
```

```
[root@linux root]# ipmitool -o intelplus -d /dev/ipmi0 channel getaccess 1 2
Maximum User IDs : 15
Enabled User IDs : 2
User ID : 2
User Name : admin
Fixed Name : No
Access Available : call-in / callback
Link Authentication : enabled
IPMI Messaging : enabled
Privilege Level : ADMINISTRATOR
```

```
[root@linux root]# ipmitool -o intelplus -d /dev/ipmi0 channel getaccess 2 2
Maximum User IDs : 15
Enabled User IDs : 2
User ID : 2
User Name : admin
Fixed Name : No
Access Available : call-in / callback
Link Authentication : enabled
IPMI Messaging : enabled
Privilege Level : ADMINISTRATOR
```

4. Apply the following local area network (LAN) settings for the BMC on the Altix XE system, for which the IPMI device is `/dev/ipmi0`. The BMC LAN settings apply to LAN channels 1 and 2.

Note: You must apply each change separately for channel 1 and for channel 2.

- Set the IP Address (use the same IP address for both channels):

```
ipmitool -o intelplus -d /dev/ipmi0 lan set 1|2 ipaddr IP_address
```

- Set the subnet mask (use the same value for both channels):

```
ipmitool -o intelplus -d /dev/ipmi0 lan set 1|2 netmask netmask
```

- Enable address resolution protocol (ARP) responses:

```
ipmitool -o intelplus -d /dev/ipmi0 lan set 1|2 arp respond on
```

- Enable *gratuitous ARP*, which broadcasts the MAC address to IP address mappings on a specified interface:

```
ipmitool -o intelplus -d /dev/ipmi0 lan set 1|2 arp generate on
```

- Set the gratuitous ARP interval (in seconds):

Note: An interval of 5 seconds is supported for CXFS.

```
ipmitool -o intelplus -d /dev/ipmi0 lan set 1|2 arp interval 5
```

For example:

```
[root@linux root]# ipmitool -o intelplus -d /dev/ipmi0 lan set 1 ipaddr nodename-bmc.company.com
Setting LAN IP Address to nodename-bmc.company.com
[root@linux root]# ipmitool -o intelplus -d /dev/ipmi0 lan set 2 ipaddr nodename-bmc.company.com
Setting LAN IP Address to nodename-bmc.company.com
[root@linux root]# ipmitool -o intelplus -d /dev/ipmi0 lan set 1 netmask 255.255.0.0
Setting LAN Subnet Mask to 255.255.0.0
[root@linux root]# ipmitool -o intelplus -d /dev/ipmi0 lan set 2 netmask 255.255.0.0
Setting LAN Subnet Mask to 255.255.0.0
[root@linux root]# ipmitool -o intelplus -d /dev/ipmi0 lan set 1 arp respond on
Enabling BMC-generated ARP responses
```

```
[root@linux root]# ipmitool -o intelplus -d /dev/ipmi0 lan set 2 arp respond on
Enabling BMC-generated ARP responses
[root@linux root]# ipmitool -o intelplus -d /dev/ipmi0 lan set 1 arp generate on
Enabling BMC-generated Gratuitous ARPs
[root@linux root]# ipmitool -o intelplus -d /dev/ipmi0 lan set 2 arp generate on
Enabling BMC-generated Gratuitous ARPs
[root@linux root]# ipmitool -o intelplus -d /dev/ipmi0 lan set 1 arp interval 5
BMC-generated Gratuitous ARP interval: 5.0 seconds
[root@linux root]# ipmitool -o intelplus -d /dev/ipmi0 lan set 2 arp interval 5
BMC-generated Gratuitous ARP interval: 5.0 seconds
```

5. Verify your changes by using the following command:

```
ipmitool -o intelplus -d /dev/ipmi0 lan print 112
```

For example:

```
[root@linux root]# ipmitool -o intelplus -d /dev/ipmi0 lan print 1
Set in Progress          : Set Complete
Auth Type Support       : NONE MD5 PASSWORD
Auth Type Enable       : Callback :
                        : User      :
                        : Operator  :
                        : Admin    : MD5 PASSWORD
                        : OEM      :
IP Address Source      : Static Address
IP Address              : nodename-bmc.company.com
Subnet Mask             : 255.255.0.0
MAC Address             : 00:04:23:d5:af:3c
SNMP Community String  :
IP Header               : TTL=0x40 Flags=0x40 Precedence=0x00 TOS=0x10
BMC ARP Control        : ARP Responses Enabled, Gratuitous ARP Enabled
Gratituous ARP Intrvl  : 5.0 seconds
Default Gateway IP     : 0.0.0.0
Default Gateway MAC    : 00:00:00:00:00:00
Backup Gateway IP      : 0.0.0.0
Backup Gateway MAC     : 00:00:00:00:00:00
RMCP+ Cipher Suites    : 0,1,2,3,4,5,6,7,8,9,10,11,12,13,14
Cipher Suite Priv Max  : XXXXXXXXXXXXXXXX
                        : X=Cipher Suite Unused
                        : c=CALLBACK
                        : u=USER
```

F: System Reset Configuration

```
: o=OPERATOR
: a=ADMIN
: O=OEM
```

6. Verify the BMC configuration and connectivity from a remote node by issuing `ipmitool(1)` commands remotely using the `lanplus` interface:

```
ping IP_address_or_hostname
```

```
ipmitool -I lanplus -o intelplus -H IP_address_or_hostname -U admin -P admin_passwd lan print 1|2
```

For example (line breaks shown here for readability):

```
[root@linux root]# ping nodename-bmc.company.com
```

```
[root@linux root]# ipmitool -I lanplus -o intelplus -H nodename-bmc.company.com -U admin \
-P mypassword lan print 1
```

```
Set in Progress      : Set Complete
Auth Type Support    : NONE MD5 PASSWORD
Auth Type Enable     : Callback :
                    : User      :
                    : Operator  :
                    : Admin    : MD5 PASSWORD
                    : OEM      :

IP Address Source    : Static Address
IP Address           : nodename-bmc.company.com
Subnet Mask          : 255.255.0.0
MAC Address          : 00:04:23:d5:af:3c
SNMP Community String :
IP Header            : TTL=0x40 Flags=0x40 Precedence=0x00 TOS=0x10
BMC ARP Control      : ARP Responses Enabled, Gratuitous ARP Enabled
Gratuitous ARP Intrvl : 5.0 seconds
Default Gateway IP    : 0.0.0.0
Default Gateway MAC   : 00:00:00:00:00:00
Backup Gateway IP     : 0.0.0.0
Backup Gateway MAC    : 00:00:00:00:00:00
RMCP+ Cipher Suites  : 0,1,2,3,4,5,6,7,8,9,10,11,12,13,14
Cipher Suite Priv Max : XXXXXXXXXXXXXXXX
                    : X=Cipher Suite Unused
                    : c=CALLBACK
                    : u=USER
                    : o=OPERATOR
                    : a=ADMIN
```

: O=OEM

For more information, see the `ipmitool(1)` man page, *Guide to Programming Environments and Tools Available on SGI Altix XE System*, and the user guide or quick start guide for your system.

L2 System Controller

The L2 system controller and the required USB cables are optional equipment available for purchase. The L2 method is recommended when available.

The L2 controller must not be on the primary CXFS private network. Ideally, the L2 controller should be on a different private network that is reachable by all server nodes in the cluster. A public network is not ideal for security reasons, but is acceptable.

Altix systems with an integrated L2 (such as a NUMALink 4 R-brick), Altix 3000 Bx2 systems, Altix 450 systems and Altix 4700 systems use the L2 over Ethernet. See Figure F-1

In Altix 350, use IO10 and a *multiport serial adapter cable*, which is a device that provides four DB9 serial ports from a 36-pin connector; see Figure F-2.

Use the modem port on the L2 system controller as shown in Figure F-3. Use DB9 serial ports on an IX-brick on Altix 3000 and Origin 3000. Connect the serial cable to the modem port on one end and the serial port on the IX-brick (for example, serial port connector 0), as shown in Figure F-4.

Figure F-5, Figure F-6, and Figure F-7 show serial connections for two machines with an L2 system controller. (These figure shows direct attached storage. Serial connections for other storage configurations will be the same.)

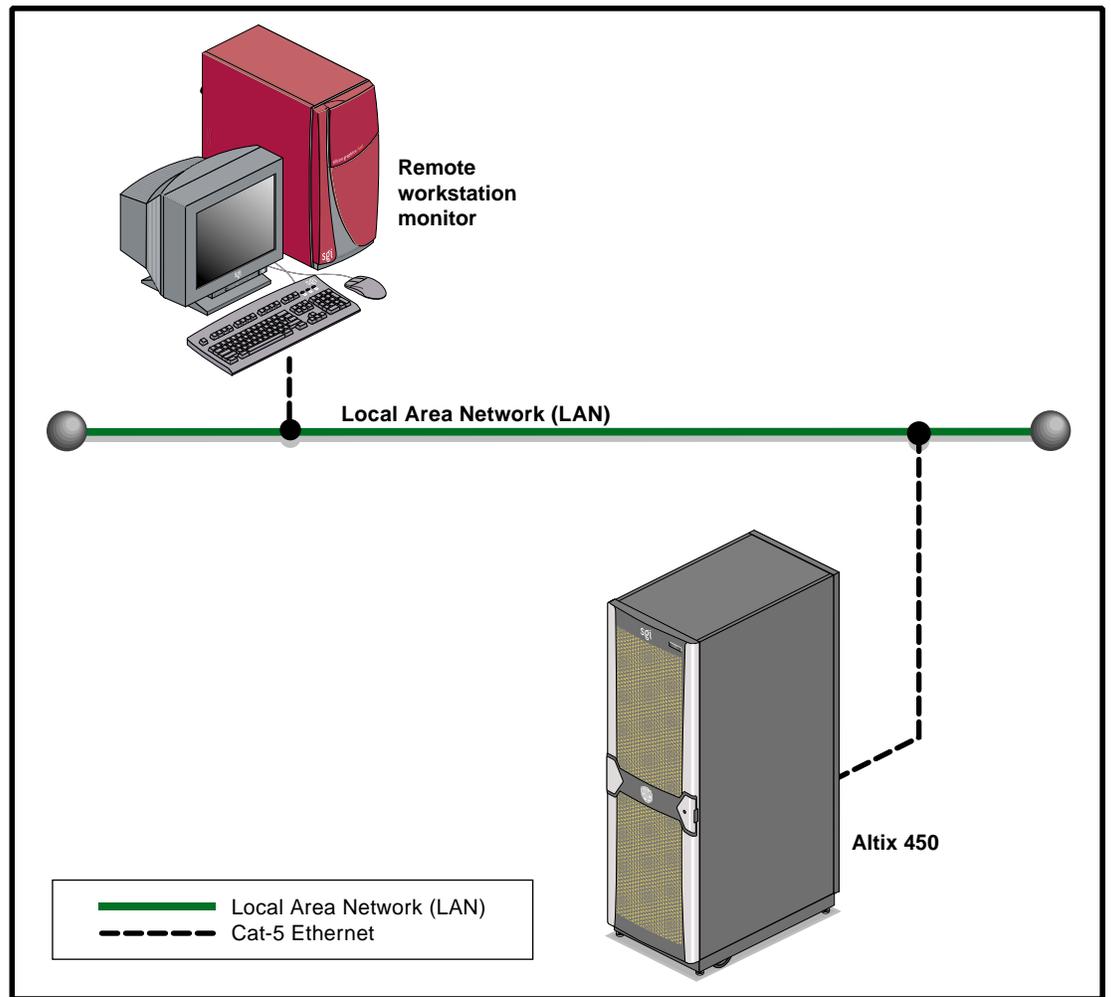


Figure F-1 SGI Altix 450 System Control Network

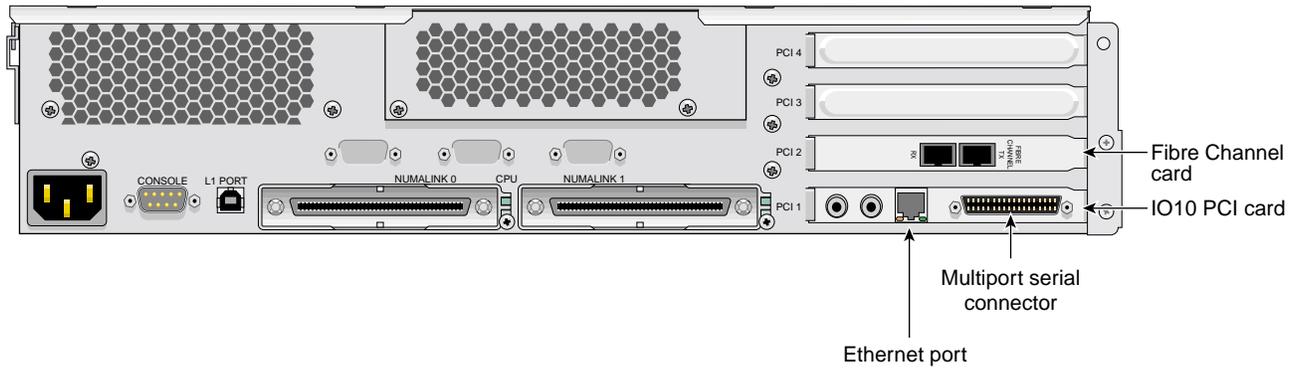


Figure F-2 Altix 350 Rear Panel

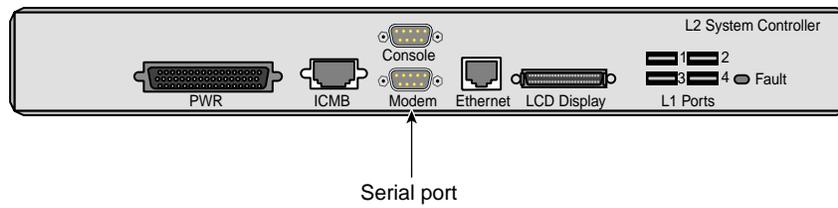


Figure F-3 L2 Rear Panel

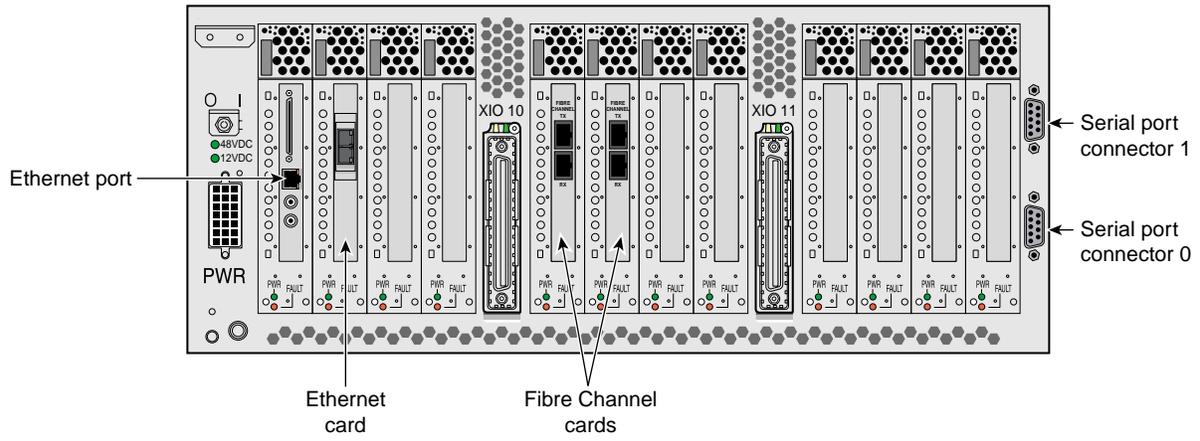


Figure F-4 IX-brick Rear Panel

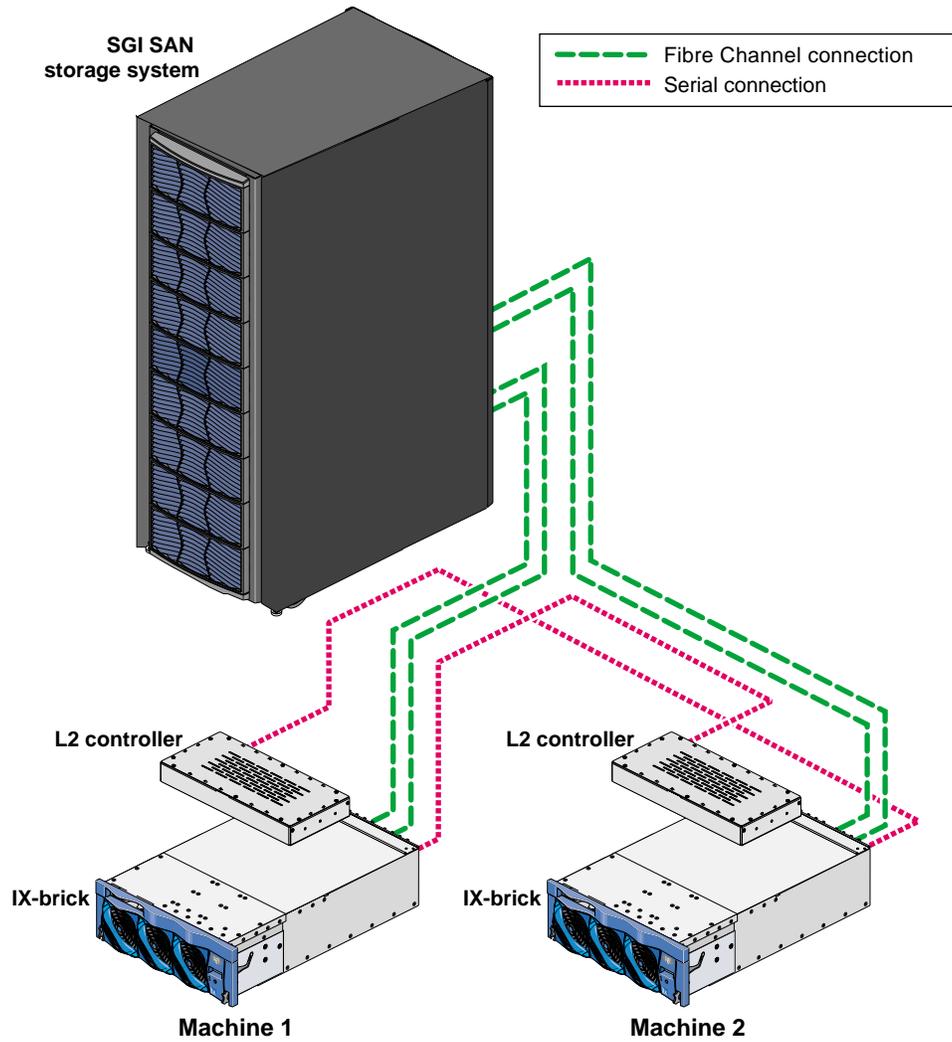


Figure F-5 Altix 3000 and Origin 3000 Serial Connections

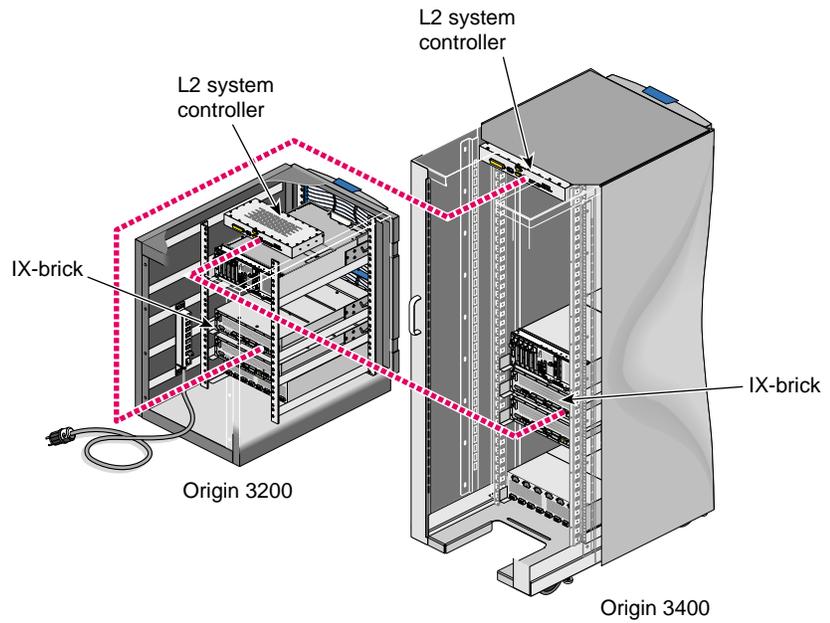


Figure F-6 Serial Connection Between SGI Origin 3200 and Origin 3400/3800 Servers

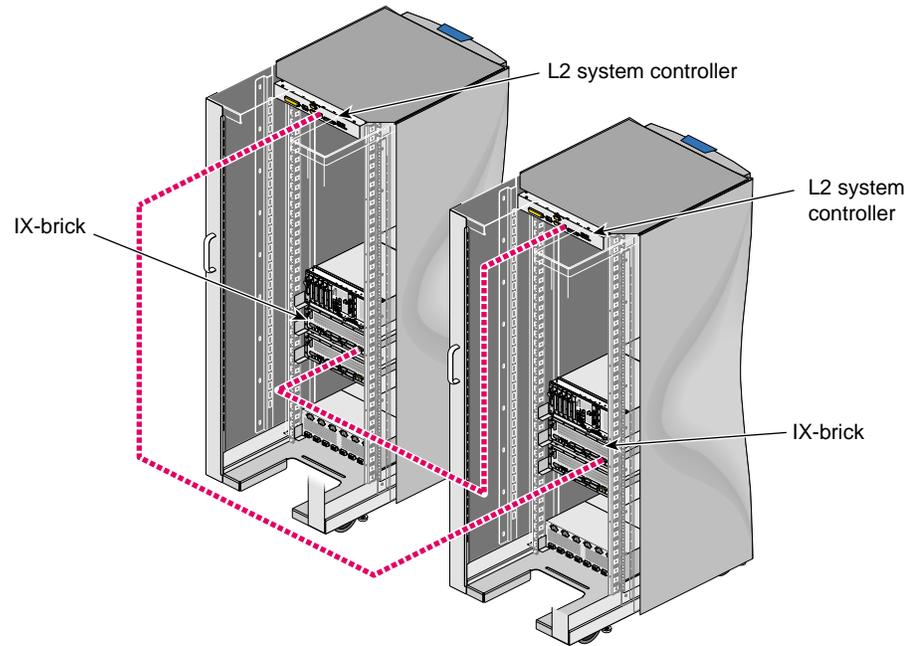


Figure F-7 Serial Connection Between Two SGI Origin 3400 or SGI Origin 3800 Servers

L1 System Controller

L1 system controller can be used for reset, however, SGI recommends the use of L2 when available.

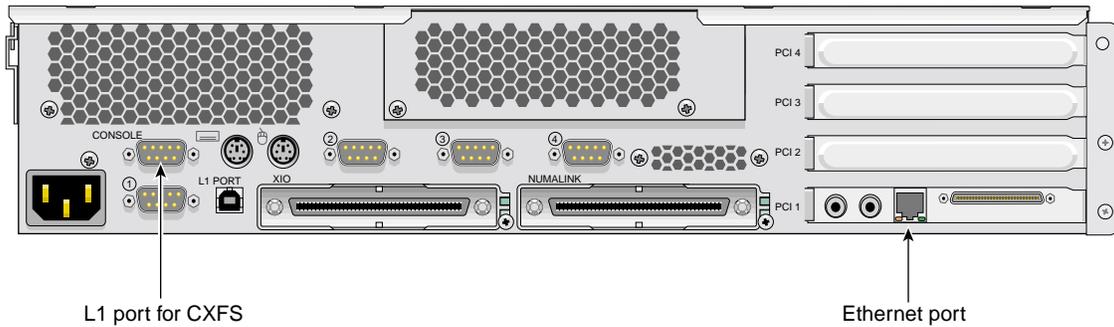


Figure F-8 Origin 350 Rear Panel

Connect the serial cable to the console port (port labeled **CONSOLE**) on one end and the serial port of other node on the other end. The serial ports on Origin 350 are labeled as **1**, **2**, **3**, and **4**; see Figure F-8.

Note: The USB port on the Origin 350 is labeled **L1 PORT**. Do not use this port for system reset. Use the port labeled **CONSOLE** as shown in Figure F-8.

Redirecting the Console for Origin 300, Origin 350, Origin 3200C, Onyx 300, Onyx 350, and Onyx 3200C

On Origin 300, Origin 350, Origin 3200C, Onyx 300, Onyx 350, and Onyx 3200C systems, there is only one serial/USB port that provides both L1 system controller and console support for the machine. In a CXFS configuration, this port (the DB9 connector) is used for system reset. It is connected to a serial port in another node or to the Ethernet multiplexer.

To get access to console input and output, you must redirect the console to another serial port in the machine.

Use the following procedure to redirect the console:

1. Edit the `/etc/inittab` file to use an alternate serial port.
2. Either issue an `init q` command or reboot.

For example, suppose you had the following in the `/etc/inittab` file (line breaks added for readability):

```
# on-board ports or on Challenge/Onyx MP machines, first IO4 board ports
t1:23:respawn:/sbin/suattr -C CAP_FOWNER,CAP_DEVICE_MGT,CAP_DAC_WRITE+ip
-c "exec /sbin/getty ttyd1 console"      # alt console
t2:23:off:/sbin/suattr -C CAP_FOWNER,CAP_DEVICE_MGT,CAP_DAC_WRITE+ip
-c "exec /sbin/getty -N ttyd2 co_9600"   # port 2
```

You could change it to the following:

```
# on-board ports or on Challenge/Onyx MP machines, first IO4 board ports
t1:23:off:/sbin/suattr -C CAP_FOWNER,CAP_DEVICE_MGT,CAP_DAC_WRITE+ip
-c "exec /sbin/getty ttyd1 co_9600"     # port 1
t2:23:respawn:/sbin/suattr -C CAP_FOWNER,CAP_DEVICE_MGT,CAP_DAC_WRITE+ip
-c "exec /sbin/getty -N ttyd2 console"  # alt console
```



Caution: Redirecting the console by using the above method works only when the IRIX operating system is running. To access the console when the operating system is not running (miniroot), you must physically reconnect the machine: unplug the serial hardware reset cable from the console/L1 port and then connect the console cable.

MSC System Controller

Figure F-9 and Figure F-10 show the serial connection between two deskside servers.

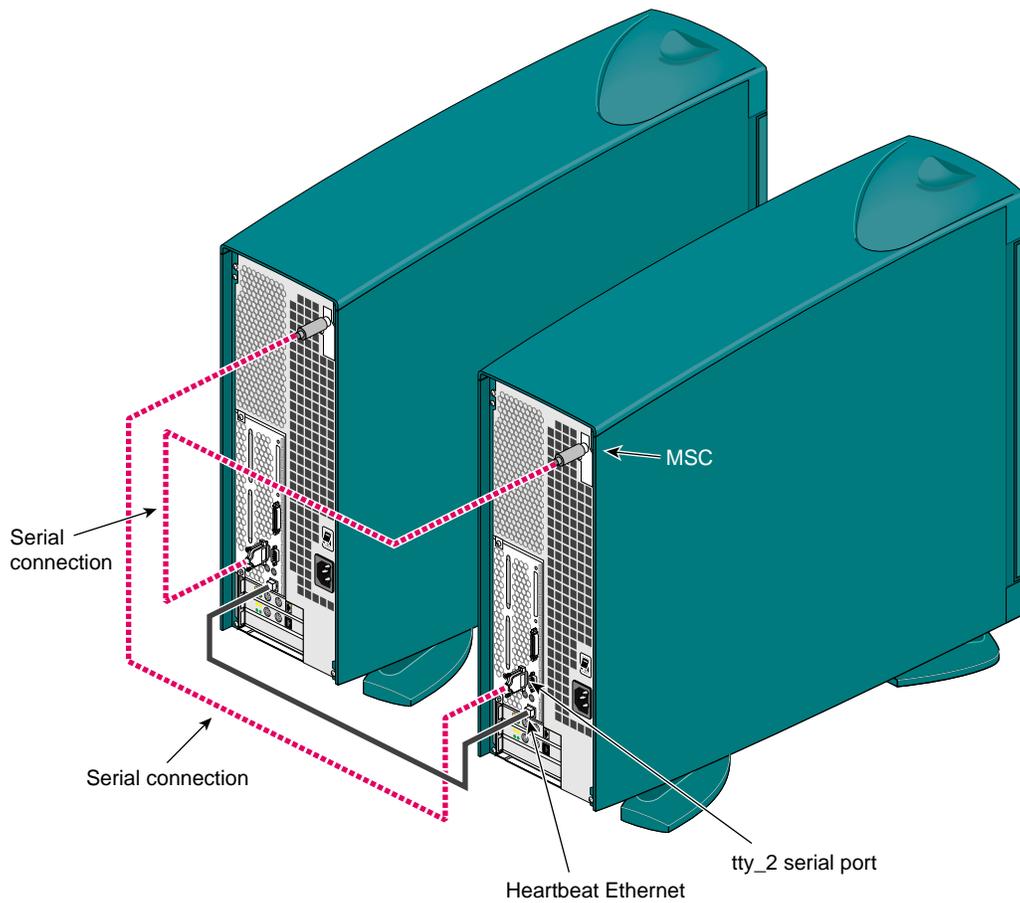


Figure F-9 Serial Connection Between Two Origin 200 Deskside Servers

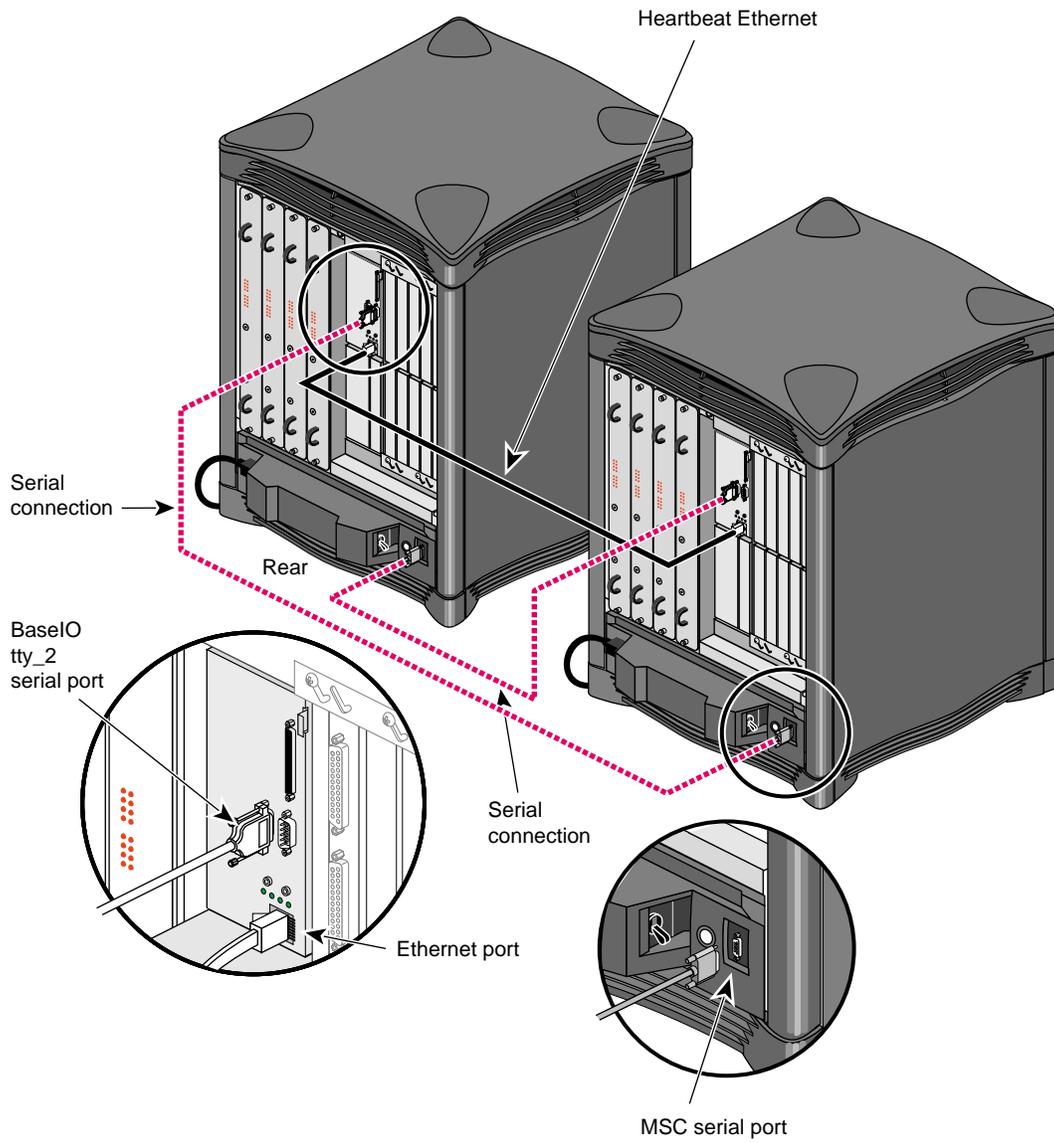


Figure F-10 Serial Connection Between Two SGI 2200 Deskside Servers

MMSC System Controller

Figure F-11 shows the MMSC. The alternate console port should be connected to the serial port on another machine using a serial cable.

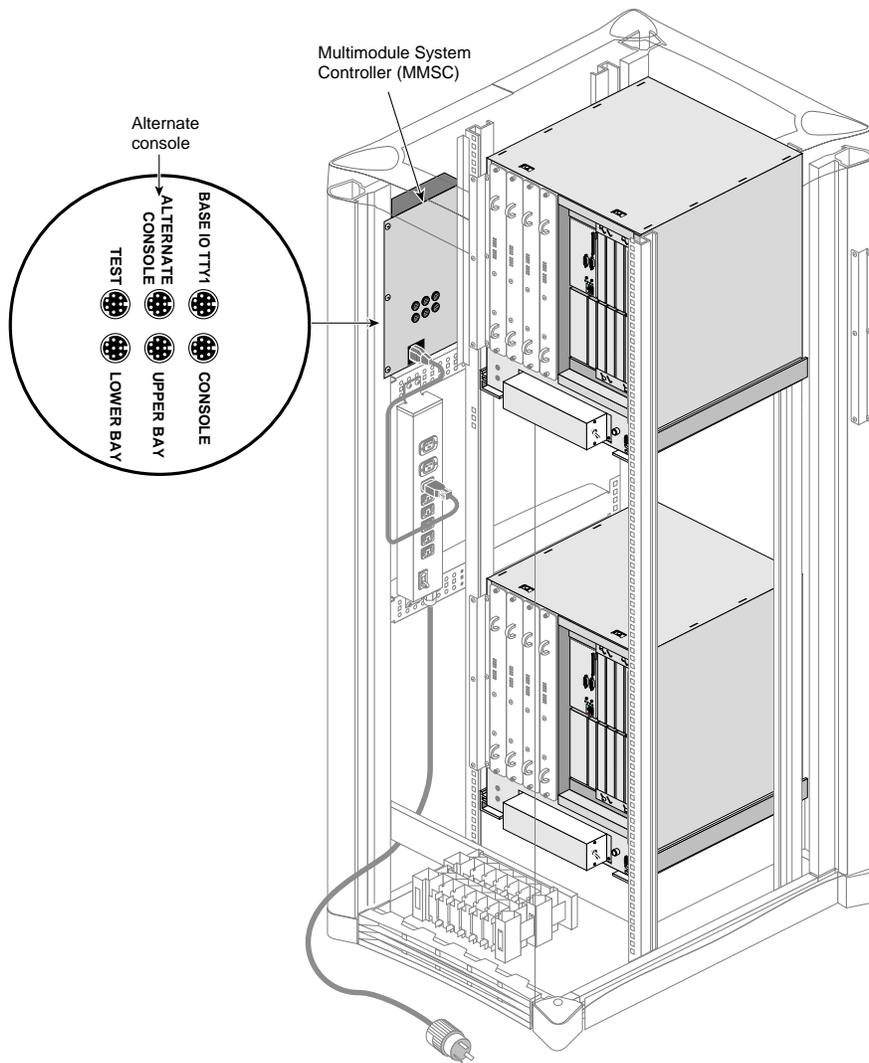


Figure F-11 MMSC Serial Port

Reference to `cmgr` Tasks

This appendix discusses the following:

- "cmgr Overview" on page 558
- "Initial Setup with the `cmgr` Command" on page 564
- "Set Configuration Defaults with `cmgr`" on page 576
- "Node Tasks with `cmgr`" on page 577
- "Cluster Tasks with `cmgr`" on page 597
- "Cluster Services Tasks with `cmgr`" on page 605
- "CXFS Filesystem Tasks with `cmgr`" on page 611
- "Switches and I/O Fencing Tasks with `cmgr`" on page 624
- "Script Example" on page 629
- "Creating a `cmgr` Script Automatically" on page 632
- "Troubleshooting `cmgr`" on page 636
- "Additional `cmgr` Examples" on page 636

For an overview of the tasks that must be performed to configure a cluster, see "Initial Setup with the `cmgr` Command" on page 564.

Tasks must be performed using a certain hierarchy. For example, to modify a partition ID, you must first identify the node name.

You can also use the `clconf_info` tool to view status. See Chapter 17, "Monitoring Status" on page 399.

Note: CXFS requires a license key to be installed on each server-capable node. If you install the software without properly installing the license key, you cannot use the `cmgr` command. For more information about licensing, see Chapter 4, "CXFS License Keys" on page 75.

For information about using the preferred `cxfs_admin` command rather than `cmgr`, see Chapter 12, "Reference to `cxfs_admin` Tasks" on page 255 and Appendix H, "Migration from `cmgr` to `cxfs_admin`" on page 659

cmgr Overview

To use the `cmgr` command, you must be logged in as `root` on a CXFS administration node. Then enter either of the following:

```
# /usr/cluster/bin/cmgr
```

or

```
# /usr/cluster/bin/cluster_mgr
```

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

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

```
cmgr>
```

For more information, see the `cmgr(1M)` man page]¹

Making Changes Safely

Do not make configuration changes on two different administration nodes in the pool simultaneously, or use the CXFS GUI, `cmgr`, and `xvm` commands simultaneously to make changes. You should run one instance of the `cmgr` command or the CXFS GUI on a single administration node in the pool when making changes at any given time. However, you can use any node in the pool when requesting status or configuration information.

Getting Help

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

¹ This man page is also accessible by `man cluster_mgr` for historical purposes.

Using Prompt Mode

The `-p` option to `cmgr` displays prompts for the required inputs of administration commands that define and modify CXFS components. You can run in prompt mode in either of the following ways:

- Specify a `-p` option on the command line:

```
# cmgr -p
```

- Execute a `set prompting on` command after you have brought up `cmgr`, as in the following example:

```
cmgr> set prompting on
```

This method allows you to toggle in and out of prompt mode as you execute individual subcommands. To get out of prompt mode, enter the following:

```
cmgr> set prompting off
```

The following shows an example of the questions that may be asked in prompting mode (the actual questions asked will vary depending upon your answers to previous questions):

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

Hostname[optional] ?
Is this a FailSafe node <true|false> ?
Is this a CXFS node <true|false> ?
Operating System <IRIX|Linux32|Linux64|AIX|HPUX|Solaris|MacOSX|Windows> ?
Node Function <server_admin|client_admin|client_only> ?
Node ID ?[optional]
Partition ID ?[optional] (0)
Do you wish to define failure hierarchy[y/n]:
Reset type <powerCycle|reset|nmi> ? (powerCycle)
Do you wish to define system controller info[y/n]:
Sysctrl Type <msc|mmsc|l2|l1|bmc>? (msc)
Sysctrl Password[optional] ? ( )
Sysctrl Status <enabled|disabled> ?
Sysctrl Owner ?
Sysctrl Device ?
Sysctrl Owner Type <tty|network|ipmi> ? (tty)
Number of Network Interfaces ? (1)
```

NIC 1 - IP Address ?

For details about this task, see "Define a Node with `cmgr`" on page 577.

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 executes the current definitions or modifications and returns to the `cmgr>` prompt

Using Script Files

You can execute a series of `cmgr` commands by using the `-f` option and specifying an input file:

```
cmgr -f input_file
```

Or, you could include the following as the first line of the file and then execute it as a bash script:

```
#!/usr/cluster/bin/cmgr -f
```

Each line of the file must be a valid `cmgr` command line, comment line (starting with #), or a blank line.

Note: You must include a `done` command line to finish a multilevel command and end the file with a `quit` command line.

If any line of the input file fails, `cmgr` will exit. You can choose to ignore the failure and continue the process by using the `-i` option with the `-f` option, as follows:

```
cmgr -if input_file
```

Or include it in the first line for a script:

```
#!/usr/cluster/bin/cmgr -if
```

Note: If you include `-i` when using a `cmgr` command line as the first line of the script, you must use this exact syntax (that is, `-if`).

For example, suppose the file `/tmp/showme` contains the following:

```
cxfs6# more /tmp/showme
show clusters
show nodes in cluster cxfs6-8
quit
```

You can execute the following command, which will yield the indicated output:

```
cxfs6# /usr/cluster/bin/cmgr -if /tmp/showme

1 Cluster(s) defined
    cxfs6-8
```

```
Cluster cxfs6-8 has following 3 machine(s)
    cxfs6
    cxfs7
    cxfs8
```

Or you could include the `cmgr` command line as the first line of the script, give it execute permission, and execute `showme` itself:

```
cxfs6# more /tmp/showme
#!/usr/cluster/bin/cmgr -if
#
show clusters
show nodes in cluster cxfs6-8
quit
```

```
cxfs6# /tmp/showme

1 Cluster(s) defined
    cxfs6-8
```

```
Cluster cxfs6-8 has following 3 machine(s)
    cxfs6
    cxfs7
    cxfs8
```

For an example of defining a complete cluster, see "Script Example" on page 629.

Invoking a Shell from within `cmgr`

To invoke a shell from within `cmgr`, enter the following:

```
cmgr> sh  
cxfs6#
```

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

```
cxfs6# exit  
cmgr>
```

Entering Subcommands on the Command Line

You can enter some `cmgr` subcommands directly from the command line using the following format:

```
cmgr -c "subcommand"
```

where *subcommand* can be any of the following with the appropriate operands:

- `admin`, which allows you to perform certain actions such as resetting a node
- `delete`, which deletes a cluster or a node
- `help`, which displays help information
- `show`, which displays information about the cluster or nodes
- `start`, which starts CXFS services and sets the configuration so that CXFS services will be automatically restarted upon reboot
- `stop`, which stops CXFS services and sets the configuration so that CXFS services are not restarted upon reboot
- `test`, which tests connectivity

For example, to display information about the cluster, enter the following:

```
# cmgr -c "show clusters"  
1 Cluster(s) defined  
    eagan
```

See the `cmgr` man page for more information.

Template Scripts

The `/var/cluster/cmgr-templates` directory contains template `cmgr` scripts that you can modify to configure the different components of your system.

Each template file contains lists of `cmgr` commands required to create a particular object, as well as comments describing each field. The template also provides default values for optional fields.

The `/var/cluster/cmgr-templates` directory contains the following templates to create a cluster and nodes:

- `cmgr-create-cluster`
- `cmgr-create-node`

To create a CXFS configuration, you can concatenate multiple templates into one file and execute the resulting script.

Note: If you concatenate information from multiple template scripts to prepare your cluster configuration, you must remove the `quit` at the end of each template script, except for the final `quit`. A `cmgr` script must have only one `quit` line.

For example, for a three-node configuration, you would concatenate three copies of the `cmgr-create-node` file and one copy of the `cmgr-create-cluster` file.

Initial Setup with the `cmgr` Command

Note: For the initial installation, SGI highly recommends that you use the GUI guided configuration tasks. See "Initial Setup with the CXFS GUI" on page 167.

To initially configure the cluster with the `cmgr` command, do the following:

1. Follow the directions in "Preliminary Cluster Configuration Steps" on page 164.
2. Define the nodes that are eligible to be part of the cluster. The hostname/IP-address pairings and priorities of the networks must be the same for each node in the cluster. See "Define a Node with `cmgr`" on page 577.

For large clusters, SGI recommends that you define only the first three CXFS administration nodes and then continue on to the next step; add the remaining nodes after you have a successful small cluster.

The following example sequence defines three nodes. (To use the default value for a prompt, press the Enter key. The Enter key is not shown in the examples in this guide.)

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



Caution: It is critical that you enter the primary name for the first node defined in the pool.

```
cxfs6 # /usr/cluster/bin/cmgr -p
Welcome to SGI Cluster Manager Command-Line Interface

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

Hostname[optional] ?
Is this a FailSafe node <true|false> ? false
Is this a CXFS node <true|false> ? true
Operating System <IRIX|Linux32|Linux64|AIX|HPUX|Solaris|MacOSX|Windows> ? irix
Node Function <server_admin|client_admin|client_only> ? server_admin
Node ID[optional]?
Partition ID[optional] ? (0)
Do you wish to define failure hierarchy[y/n]:n
Reset type <powerCycle|reset|nmi> ? (powerCycle)
Do you wish to define system controller info[y/n]:y
Sysctrl Type <msc|mmsc|l2|l1|bmc> ? (msc)
Sysctrl Password[optional] ? ( )
Sysctrl Status <enabled|disabled> ? enabled
Sysctrl Owner ? cxfs8
Sysctrl Device ? /dev/ttyd2
Sysctrl Owner Type <tty|network|ipmi> ? (tty)
Number of Network Interfaces ? (1)
NIC 1 - IP Address ? cxfs6
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
```

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] ?
Is this a FailSafe node <true|false> ? false
Is this a CXFS node <true|false> ? true
Node Function <server_admin|client_admin|client_only> ? server_admin
Operating System <IRIX|Linux32|Linux64|AIX|HPUX|Solaris|MacOSX|Windows> ? irix
Node ID[optional] ?
Partition ID[optional] ? (0)
Do you wish to define failure hierarchy[y/n]:n
Reset type <powerCycle|reset|nmi> ? (powerCycle)
Do you wish to define system controller info[y/n]:y
Sysctrl Type <msc|mmsc|l2|l1|bmc> ? (msc)
Sysctrl Password[optional] ? ( )
Sysctrl Status <enabled|disabled> ? enabled
Sysctrl Owner ? cxfs6
Sysctrl Device ? /dev/ttyd2
Sysctrl Owner Type <tty|network|ipmi> ? (tty)
Number of Network Interfaces ? (1)
NIC 1 - IP Address ? cxfs7
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
```

Successfully defined node cxfs7

To define the third node, named cxfs8, enter the following:

```
cmgr> define node cxfs8
Enter commands, you may enter "done" or "cancel" at any time to exit

Hostname[optional] ?
Is this a FailSafe node <true|false> ? false
Is this a CXFS node <true|false> ? true
```

```

Node Function <server_admin|client_admin|client_only> ? server_admin
Operating System <IRIX|Linux32|Linux64|AIX|HPUX|Solaris|MacOSX|Windows> ? irix
Node ID[optional] ?
Partition ID[optional] ? (0)
Do you wish to define failure hierarchy[y/n]:n
Reset type <powerCycle|reset|nmi> ? (powerCycle)
Do you wish to define system controller info[y/n]:y
Sysctrl Type <msc|mmsc|l2|l1|bmc> ? (msc)
Sysctrl Password[optional] ? ( )
Sysctrl Status <enabled|disabled> ? enabled
Sysctrl Owner ? cxfs7
Sysctrl Device ? /dev/ttyd2
Sysctrl Owner Type <tty|network|ipmi> ? (tty)
Number of Network Interfaces ? (1)
NIC 1 - IP Address ? cxfs8
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

Successfully defined node cxfs8

```

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

```

cmgr> show nodes in pool

3 Machine(s) defined
    cxfs6
    cxfs7
    cxfs8

```

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

```

cmgr> show node cxfs6

Logical Machine Name: cxfs6
Hostname: cxfs6.example.com
Operating System: irix
Node Is FailSafe: false
Node Is CXFS: true
Node Function: server_admin
Nodeid: 13203

```

```
Partition id: 0
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: cxfs6
ControlNet HB: true
ControlNet Control: true
ControlNet Priority: 1
```

3. Define the cluster and add the nodes to it. See "Define a Cluster with cmgr" on page 597.

For example, to define a cluster named `cxfs6-8` and add the nodes that are already defined, enter the following:

```
cmgr> define cluster cxfs6-8
Enter commands, you may enter "done" or "cancel" at any time to exit

Is this a FailSafe cluster <true|false> false ?
Is this a CXFS cluster <true|false> true ?
Cluster Notify Cmd [optional] ?
Cluster Notify Address [optional] ?
Cluster mode <normal|experimental>[optional]
Cluster ID ? 22

No nodes in cluster cxfs6-8

Add nodes to or remove nodes from cluster cxfs6-8
Enter "done" when completed or "cancel" to abort

cxfs6-8 ? add node cxfs6
cxfs6-8 ? add node cxfs7
cxfs6-8 ? add node cxfs8
cxfs6-8 ? done
Successfully defined cluster cxfs6-8

Added node <cxfs6> to cluster <cxfs6-8>
Added node <cxfs7> to cluster <cxfs6-8>
Added node <cxfs8> to cluster <cxfs6-8>
```

The fail action hierarchy is the set of instructions that determines which method is used in case of failure. If you set a hierarchy including fencing, you could define the switch at this point. For more information, see "Switches and I/O Fencing Tasks with `cmgr`" on page 624.

To define a list of private networks that can be used in case the highest priority network (consisting by default of the priority 1 NICs) fails, use the `add net` command; see "Define a Node with `cmgr`" on page 577.

For more information, see "Define a Cluster with `cmgr`" on page 597.

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

```
cmgr> show clusters

1 Cluster(s) defined
    cxfs6-8

cmgr> show cluster cxfs6-8
Cluster Name: cxfs6-8
Cluster Is FailSafe: false
Cluster Is CXFS: true
Cluster ID: 22
Cluster CX mode: normal

Cluster cxfs6-8 has following 3 machine(s)
    cxfs6
    cxfs7
    cxfs8

CXFS Failover Networks:
    default network 0.0.0.0, mask 0.0.0.0
```

For an example of this step using a script, see "Script Example" on page 629.

4. Start CXFS services for each node in the cluster by entering the following:

```
start cx_services for cluster clustername
```

For example:

```
cmgr> start cx_services for cluster cxfs6-8
```

CXFS services have been activated in cluster cxf6-8

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

Note: If you stop CXFS services using either the GUI or cmgr, the automatic restart capability is turned off. You must start CXFS services again to reinstate the automatic restart capability.

To verify that CXFS services have been started in the cluster and there is a membership formed, you can use the following cmgr command:

```
show status of cluster clustername
```

For example:

```
cmgr> show status of cluster cxf6-8
```

```
Cluster (cxf6-8) is not configured for FailSafe
```

```
CXFS cluster state is ACTIVE.
```

You can also use the clconf_info command. For example:

```
cxf6 # /usr/cluster/bin/clconf_info
```

```
Event at [2004-04-16 09:20:59]
```

```
Membership since Fri Apr 16 09:20:56 2004
```

Node	NodeID	Status	Age	CellID
cxf7	12812	up	0	1
cxf6	13203	up	0	0
cxf8	14033	up	0	2

```
0 CXFS FileSystems
```

For more information, see "Display a Cluster with cmgr" on page 604.

5. Obtain a shell window for one of the CXFS administration nodes in the cluster and use the IRIX `fx(1M)` command or the Linux `parted(8)` command to create a volume header on the disk drive. For information, see the man pages, *IRIX Admin: Disks and Filesystems*, and *Linux Configuration and Operations Guide*.
6. Create the XVM logical volumes. In the shell window, use the `xvm` command line interface. For information, see the *XVM Volume Manager Administrator's Guide*.
7. Make the XFS filesystems. In the shell window, use the `mkfs` command. For information, see the *XVM Volume Manager Administrator's Guide* and *IRIX Admin: Disks and Filesystems*.
8. Define the filesystems by using the `define cxfs_filesystem` subcommand to `cmgr`. See "CXFS Filesystem Tasks with `cmgr`" on page 611.

The following example shows two potential metadata servers for the `fs1` 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 `cxfs7` (rank 1) will be used. It also shows the filesystem being mounted by default on all nodes in the cluster (Default Local Status enabled) but explicitly not mounted on `cxfs8`.

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

Do the following:

```
cmgr> define cxfs_filesystem fs1 in cluster cxfs6-8
```

```
(Enter "cancel" at any time to abort)
```

```
Device ? /dev/cxvm/d761un0s0
```

```
Mount Point ? /mnts/fs1
```

```
Mount Options[optional] ?
```

```
Use Forced Unmount ? <true|false> ? false
```

```
Default Local Status <enabled|disabled> ? (enabled)
```

```
DEFINE CXFS FILESYSTEM OPTIONS
```

```
0) Modify Server.
```

```
1) Add Server.
```

G: Reference to cmgr Tasks

- 2) Remove Server.
- 3) Add Enabled Node.
- 4) Remove Enabled Node.
- 5) Add Disabled Node.
- 6) Remove Disabled Node.
- 7) Show Current Information.
- 8) Cancel. (Aborts command)
- 9) Done. (Exits and runs command)

Enter option:1

No current servers

Server Node ? **cxfs6**

Server Rank ? **0**

- 0) Modify Server.
- 1) Add Server.
- 2) Remove Server.
- 3) Add Enabled Node.
- 4) Remove Enabled Node.
- 5) Add Disabled Node.
- 6) Remove Disabled Node.
- 7) Show Current Information.
- 8) Cancel. (Aborts command)
- 9) Done. (Exits and runs command)

Enter option:1

Server Node ? **cxfs7**

Server Rank ? **1**

- 0) Modify Server.
- 1) Add Server.
- 2) Remove Server.
- 3) Add Enabled Node.
- 4) Remove Enabled Node.
- 5) Add Disabled Node.
- 6) Remove Disabled Node.
- 7) Show Current Information.
- 8) Cancel. (Aborts command)

9) Done. (Exits and runs command)

Enter option:5

No disabled clients

Disabled Node ? **cxfs8**

- 0) Modify Server.
- 1) Add Server.
- 2) Remove Server.
- 3) Add Enabled Node.
- 4) Remove Enabled Node.
- 5) Add Disabled Node.
- 6) Remove Disabled Node.
- 7) Show Current Information.
- 8) Cancel. (Aborts command)
- 9) Done. (Exits and runs command)

Enter option:7

Current settings for filesystem (fs1)

CXFS servers:

Rank 0	Node cxfs6
Rank 1	Node cxfs7

Default local status: enabled

No explicitly enabled clients

Explicitly disabled clients:

Disabled Node: cxfs8

- 0) Modify Server.
- 1) Add Server.
- 2) Remove Server.
- 3) Add Enabled Node.
- 4) Remove Enabled Node.
- 5) Add Disabled Node.
- 6) Remove Disabled Node.

- 7) Show Current Information.
- 8) Cancel. (Aborts command)
- 9) Done. (Exits and runs command)

Enter option:9

Successfully defined cxfs_filesystem fs1

cmgr> **define cxfs_filesystem fs2 in cluster cxfs6-8**

(Enter "cancel" at any time to abort)

Device ? **/dev/cxvm/d771un0s0**

Mount Point ? **/mnts/fs2**

Mount Options[optional] ?

Use Forced Unmount ? <true|false> ? **false**

Default Local Status <enabled|disabled> ? (enabled)

DEFINE CXFS FILESYSTEM OPTIONS

- 0) Modify Server.
- 1) Add Server.
- 2) Remove Server.
- 3) Add Enabled Node.
- 4) Remove Enabled Node.
- 5) Add Disabled Node.
- 6) Remove Disabled Node.
- 7) Show Current Information.
- 8) Cancel. (Aborts command)
- 9) Done. (Exits and runs command)

Enter option:1

Server Node ? **cxfs8**

Server Rank ? **0**

- 0) Modify Server.
- 1) Add Server.
- 2) Remove Server.
- 3) Add Enabled Node.
- 4) Remove Enabled Node.
- 5) Add Disabled Node.

- 6) Remove Disabled Node.
- 7) Show Current Information.
- 8) Cancel. (Aborts command)
- 9) Done. (Exits and runs command)

Enter option:7

Current settings for filesystem (fs2)

CXFS servers:

Rank 0 Node cxfs8

Default local status: enabled

No explicitly enabled clients

No explicitly disabled clients

- 0) Modify Server.
- 1) Add Server.
- 2) Remove Server.
- 3) Add Enabled Node.
- 4) Remove Enabled Node.
- 5) Add Disabled Node.
- 6) Remove Disabled Node.
- 7) Show Current Information.
- 8) Cancel. (Aborts command)
- 9) Done. (Exits and runs command)

Enter option:9

Successfully defined cxfs_filesystem fs2

To see the modified contents of cluster cxfs6-8, enter the following:

```
cmgr> show cxfs_filesystems in cluster cxfs6-8
```

```
fs1  
fs2
```

9. Mount the filesystems on all nodes in the cluster by using the `admin cxfs_mount cxfs_filesystem fs1 in cluster cxfs6-8` subcommand to `cmgr`. See "Mount a CXFS Filesystem with `cmgr`" on page 618. For example:

```
cmgr> admin cxfs_mount cxfs_filesystem fs1 in cluster cxfs6-8
cxfs_mount operation successful
```

```
cmgr> admin cxfs_mount cxfs_filesystem fs2 in cluster cxfs6-8
cxfs_mount operation successful
```

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

```
cmgr> quit
```

Set Configuration Defaults with `cmgr`

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

```
set cluster clustername
set node hostname
```

clustername and *hostname* are logical names. Logical names cannot begin with an underscore (`_`) or include any whitespace characters, and can be at most 255 characters.

To view the current defaults, use the following:

```
show set defaults
```

For example:

```
cmgr> set cluster cxfs6-8
cmgr> set node cxfs6
cmgr> show set defaults
Default cluster set to: cxfs6-8
```

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

Node Tasks with `cmgr`

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

Note: The entire cluster status information is sent to each CXFS administration node each time a change is made to the cluster database; therefore, the more CXFS administration nodes in a configuration, the longer it will take.

Define a Node with `cmgr`

To define a node, use the following commands:

```
define node logical_hostname
  set hostname to hostname
  set is_failsafe to true|false
  set is_cxfs to true|false
  set operating_system to irix|linux32|linux64|aix|solaris|macosx|windows
  set node_function to server_admin|client_admin|client_only
  set nodeid to nodeID
  set partition_id to partitionID
  set hierarchy to [system][fence][reset][fencereset][shutdown]
  set reset_type to powerCycle|reset|nmi (for SGI hardware)
  set sysctrl_type to misc|mmsc|i2|i1|bmc (based on node hardware)
  set sysctrl_password to password
  set sysctrl_status to enabled|disabled
  set sysctrl_owner to node_sending_reset_command
  set sysctrl_device to port|IP_address_or_hostname_of_device
  set sysctrl_owner_type to tty|network|ipmi
  add nic IP_address_or_hostname (if DNS)
    set heartbeat to true|false
    set ctrl_msgs to true|false
    set priority to integer
  remove nic IP_address_or_hostname (if DNS)
  set weight to 0|1 (no longer needed)
```

Usage notes:

- *logical_hostname* is a simple hostname (such as `lilly`) or a fully qualified domain name (such as `lilly.example.com`) 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.

- `hostname` is the fully qualified hostname unless the simple hostname is resolved on all nodes. Use the `ping` to display the fully qualified hostname. Do not enter an IP address. The default for `hostname` is the value for `logical_hostname`; therefore, you must supply a value for this command if you use a value other than the hostname or an abbreviation of it for `logical_hostname`.
- If you are running just CXFS on this node, set `is_cxfs` to `true` and `is_failsafe` to `false`. If you are running both CXFS and FailSafe on this node in a coexecution cluster, set both values to `true`.
- `operating_system` can be set to `irix`, `linux32`, `linux64`, `aix`, `solaris`, `macosx`, or `windows`. Choose `windows` for Windows 2000, Windows 2003, or Windows XP. Choose `linux64` when defining an `x86_64` or `ia64` architecture node. (Use the `uname -i` command to determine the architecture type.)

Note: For support details, see the *CXFS MultiOS Client-Only Guide for SGI InfiniteStorage*.

If you specify `aix`, `solaris`, `macosx` or `windows`, the `weight` is assumed to be 0. If you try to specify incompatible values for `operating_system` and `is_failsafe` or `weight`, the `define` command will fail.

- `node_function` specifies the function of the node. Enter one of the following:
 - `server_admin` is an IRIX or SGI ProPack node on which you will execute cluster administration commands and that you also want to be a CXFS metadata server. (You will use the **Define a CXFS Filesystem** task to define the specific filesystem for which this node can be a metadata server.) Use this node function only if the node will be a metadata servers. You must have installed the `cluster_admin` product on this node.
 - `client_admin` is an IRIX node on which you will execute cluster administration commands but that you do not want to use as a CXFS metadata server. Use this node function only if the node will run FailSafe but you do not want it to be a metadata server. You must install the `cluster_admin` product on this node.
 - `client_only`, is a node that shares CXFS filesystems but on which you will not execute cluster administration commands and that will not be a CXFS

metadata server. Use this node function for all nodes other than those that will be metadata servers, or those that will run FailSafe without being a metadata server. You must install the `cxfs_client` product on this node. This node can run IRIX, SGI ProPack, Linux third-party, AIX, Solaris, Mac OS X, or Windows. (Nodes other than IRIX and SGI ProPack are **required** to be client-only nodes.)

AIX, Solaris, Mac OS X, and Windows nodes are automatically specified as `client-only`. You should specify `client-only` with `linux32`.

- `nodeid` is an integer in the range 1 through 32767 that is unique among the nodes in the pool. You must not change the node ID number after the node has been defined.
 - For administration nodes, this value is optional. If you do not specify a number for an administration node, CXFS will calculate an ID for you. The default ID is a 5-digit number based on the machine's serial number and other machine-specific information; it is not sequential.
 - For client-only nodes, you must specify a unique value.
- `partition_id` uniquely defines a partition in a partitioned Origin 3000 system, Altix 3000 series system, or Altix 4700 system. The `set partition_id` command is optional; if you do not have a partitioned Origin 3000 system, you can skip this command or enter 0.

Note: In an Origin 3000 system, use the `mkpart` command to determine this value:

- The `-n` option lists the partition ID (which is 0 if the system is not partitioned).
- The `-l` option lists the bricks in the various partitions (use `rack#.slot#` format in `cmgr`)

For example (output truncated here for readability):

```
# mkpart -n
Partition id = 1
# mkpart -l
partition: 3 = brick: 003c10 003c13 003c16 003c21 003c24 003c29 ...
partition: 1 = brick: 001c10 001c13 001c16 001c21 001c24 001c29 ...
```

You could enter one of the following for the **Partition ID** field:

```
1
001.10
```

To unset the partition ID, use a value of 0 or none.

On an Altix 3000, you can find the partition ID by reading the `proc` file. For example:

```
[root@linux root]# cat /proc/sgi_sn/partition_id
0
```

The 0 indicates that the system is not partitioned. If the system is partitioned, the number of partitions (such as 1, 2, etc.) is displayed.

- `hierarchy` defines the failpolicy hierarchy, which determines what happens to a failed node. You can specify up to three options. The second option will be completed only if the first option fails; the third option will be completed only if both the first and second options fail. Options must be separated by commas and no whitespace.

The option choices are as follows:

- `system` deletes all hierarchy information about the node from the database, causing the system defaults to be used. The system defaults are the same as entering `reset`, `shutdown`. This means that a reset will be performed on a node with a system controller; if the reset fails or if the node does not have a system controller, CXFS services will be forcibly shut down.
- `fence` disables access to the SAN from the problem node. Fencing provides faster recovery of the CXFS kernel membership than `reset`.
- `fencereset` performs a fence and then, if the node is successfully fenced, **also** performs an asynchronous reset of the node via a system controller; recovery begins without waiting for reset acknowledgement.

Note: SGI recommends that a server-capable node include `reset` in its `hierarchy` (unless it is the only server-capable node in the cluster). See "Isolating Failed Nodes: Failure Policies" on page 28.

- `reset` performs a system reset via a system controller.
- `shutdown` tells the other nodes in the cluster to wait for a period of time (long enough for the node to shut itself down) before reforming the CXFS kernel membership. (However, there is no notification that the node's shutdown has actually taken place.)



Caution: Because there is no notification that a shutdown has occurred, if you have a cluster with no tiebreaker, you must not use the `shutdown` setting for any server-capable node in order to avoid multiple clusters being formed. See "Shutdown" on page 35.

You cannot use `shutdown` on client nodes if you choose dynamic monitoring.

Note: If the failure hierarchy contains `reset` or `fencereset`, the reset might be performed before the system kernel core-dump can complete, resulting in an incomplete core-dump.

For a list of valid fail action sets, see "Isolating Failed Nodes: Failure Policies" on page 28.

To perform a reset only if a fencing action fails, specify the following:

```
set hierarchy fence,reset
```

Note: If `shutdown` is not specified and the other actions fail, the node attempting to deliver the CXFS kernel membership will stall delivering the membership until either the failed node attempts to re-enter the cluster or the system administrator intervenes using `cms_intervene`. Objects held by the failed nodes stall until membership finally transitions and initiates recovery.

To perform a fence and an asynchronous reset, specify the following:

```
set hierarchy fencereset
```

To return to system defaults (`reset`, `shutdown`), specify the following:

```
set hierarchy system
```

- `reset_type` applies to SGI hardware and can be one of the following:
 - `powerCycle` shuts off power to the node and then restarts it
 - `reset` simulates the pressing of the reset button on the front of the machine

- nmi (nonmaskable interrupt) performs a core-dump of the operating system kernel, which may be useful when debugging a faulty machine
- `sysctrl_type` is the system controller type based on the node hardware, as show in Table 12-1 on page 278.
- `sysctrl_password` is the password for the system controller port, not the node's root password or PROM password. On some nodes, the system administrator may not have set this password. If you wish to set or change the system controller password, consult the hardware manual for your node.
- `sysctrl_status` allows you to provide information about the system controller but temporarily disable reset by setting this value to `disabled` (meaning that CXFS cannot reset the node). To allow CXFS to reset the node, enter `enabled`. For nodes without system controllers, set this to `disabled`; see "Requirements" on page 38.
- `sysctrl_device` is one of the following:

- For systems with serial ports (`reset_comms=tty`), this is the name of the terminal port (TTY) on the owner node (the node issuing the reset). A serial cable connects the terminal port on the owner node to the system controller of the node being reset. `/dev/ttyd2` is the most commonly used port, except on Origin 300 and Origin 350 systems (where `/dev/ttyd4` is commonly used) and on Altix 350 systems (where `/dev/ttyIOC0` is commonly used).

Note: Check the owner node's specific hardware configuration to verify which device to use.

- For systems with network-attached L2 system controllers (`reset_comms=network`), this is the IP address or hostname of the L2 controller on the node being reset. For example, `reset_device=nodename-l2.company.com`. For systems with network-attached baseboard management controller (BMC) system controllers (`reset_comms=ipmi`), this is the IP address or hostname of the BMC controller on the node being reset. For example:

```
reset_device=nodename-bmc.company.com
```

- `sysctrl_owner` is the name of the node that sends the reset command. If you use `tty` for `sysctrl_owner_type`, 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.

- `sysctrl_owner_type` is either `tty` for TTY serial devices or `network` for network reset (available for systems with L2 system controllers).

Note: If you are running in coexecution with FailSafe, the `network` and `ipmi` selections are not supported.

For example:

- For an Origin 3000 system:

```
Sysctrl Device? /dev/ttyd2
Sysctrl Owner Type? <tty|network|ipmi> tty
```

- For an Altix system with an integrated L2, such as a NUMALink 4 R-brick, or SGI Altix 3000 Bx2 systems:

```
Sysctrl Device? nodename-l2.company.com
Sysctrl Owner Type? <tty|network|ipmi> network
```

- For an Altix 350:

```
Sysctrl Device? /dev/ttyIOC0
Sysctrl Owner Type? <tty|network|ipmi> tty
```

- For an Altix XE system with an integrated baseboard management controller (BMC) using the Intelligent Platform Management Interface (IPMI):

```
Sysctrl Device? nodename-bmc.company.com
Sysctrl Owner Type? ipmi
```

- `nic` is the IP address or hostname of the private network. (The hostname must be resolved in the `/etc/hosts` file.)

There can be up to 8 network interfaces. The NIC number is not significant. Priority 1 is the highest priority. By default, only the priority 1 NICs will be used as the CXFS private network and they must be on the same subnet. However, you can use the `add net` command to configure the NICs of a given priority into a network; each network takes its priority from the set of NICs it contains. In this case, if the highest priority network fails, the second will be used, and so on; see "Define a Cluster with `cmgr`" on page 597.

Note: You cannot add a NIC or a network grouping while CXFS services are active (that is, when `start cx_services` has been executed); doing so can lead to cluster malfunction. If services have been started, they should be stopped with `stop cx_services`.

If you do not use the `add net` command to group the NICs into a set of networks, all NICs other than priority 1 are ignored.

SGI requires that this network be private; see "Private Network" on page 19.

For more information about using the hostname, see "Hostname Resolution and Network Configuration Rules" on page 93.

- `weight`, which is automatically set internally to either 0 or 1 to specify how many votes a particular CXFS administration node has in CXFS kernel membership decisions. This information is now set by the `Node Function` field and this command is no longer needed.

Note: Although it is possible to use the `set weight` command to set a weight other than 0 or 1, SGI recommends that you do not do so. There is no need for additional weight.

For more information, see "CXFS Kernel Membership, Quorum, and Tiebreaker" on page 512, and "Define a Node with the GUI" on page 208.

In prompting mode, 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" on page 208. Following is a summary of the prompts.

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

Hostname[optional] ? hostname
Is this a FailSafe node <true|false> ? true|false
Is this a CXFS node <true|false> ? true
Operating System <IRIX|Linux32|Linux64|AIX|HPUX|Solaris|MacOSX|Windows> ? OS_type
Node Function <server_admin|client_admin|client_only> ? node_function
Node ID ?[optional] node_ID
Partition ID ?[optional] (0)partition_ID
Do you wish to define failure hierarchy[y/n]:y|n
```

```

Do you wish to define system controller info[y/n]:y|n
Reset type <powerCycle|reset|nmi> ? (powerCycle)
Do you wish to define system controller info[y/n]:y|n
Sysctrl Type <msc|mmsc|l2|l1|bmc>? (msc) type (based on node hardware)
Sysctrl Password[optional] ? ( )password
Sysctrl Status <enabled|disabled> ? enabled|disabled
Sysctrl Owner ? node_sending_reset_command
Sysctrl Device ? port|IP_address_or_hostname_of_device
Sysctrl Owner Type <tty|network|ipmi> ? (tty) tty|network|ipmi
Number of Network Interfaces ? (1) number
NIC 1 - IP Address ? IP_address_or_hostname (if DNS)

```

For example, in normal mode:

```

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

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

? set is_failsafe to false
? set is_cxfs to true
? set operating_system to irix
? set node_function to server_admin
? set hierarchy to fencerreset,reset
? add nic 111.11.11.111
Enter network interface commands, when finished enter "done" or "cancel"

NIC 1 - set heartbeat to true
NIC 1 - set ctrl_msgs to true
NIC 1 - set priority to 1
NIC 1 - done
? done

```

For example, in prompting mode:

```

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

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

```

```
Hostname[optional] ?
Is this a FailSafe node <true|false> ? false
Is this a CXFS node <true|false> ? true
Operating System <IRIX|Linux32|Linux64|AIX|HPUX|Solaris|MacOSX|Windows> ? irix
Node Function <server_admin|client_admin|client_only> server_admin
Node ID[optional]?
Partition ID ? [optional] (0)
Do you wish to define failure hierarchy[y|n]:y
Hierarchy option 0 <System|FenceReset|Fence|Reset|Shutdown>[optional] ? fencereset
Hierarchy option 1 <System|FenceReset|Fence|Reset|Shutdown>[optional] ? reset
Hierarchy option 2 <System|FenceReset|Fence|Reset|Shutdown>[optional] ?
Reset type <powerCycle|reset|nmi> ? (powerCycle)
Do you wish to define system controller info[y/n]:n
Number of Network Interfaces ? (1)
NIC 1 - IP Address ? 111.11.11.111
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
```

Following is an example of defining a Solaris node in prompting mode (because it is a Solaris node, no default ID is provided, and you are not asked to specify the node function because it must be `client_only`).

```
cmgr> define node solaris1
Enter commands, you may enter "done" or "cancel" at any time to exit

Hostname[optional] ?
Is this a FailSafe node <true|false> ? false
Is this a CXFS node <true|false> ? true
Operating System <IRIX|Linux32|Linux64|AIX|HPUX|Solaris|MacOSX|Windows> ? solaris
Node ID ? 7
Do you wish to define failure hierarchy[y|n]:y
Hierarchy option 0 <System|FenceReset|Fence|Reset|Shutdown>[optional] ? fence
Hierarchy option 1 <System|FenceReset|Fence|Reset|Shutdown>[optional] ?
Number of Network Interfaces ? (1)
NIC 1 - IP Address ? 163.154.18.172
```

Modify a Node with `cmgr`

To modify an existing node, use the following commands:

```

modify node logical_hostname
  set hostname to hostname
  set partition_id to partitionID
  set reset_type to powerCycle|reset|nmi
  set sysctrl_type to msc|mmsc|I2|I1|bmc (based on node hardware)
  set sysctrl_password to password
  set sysctrl_status to enabled|disabled
  set sysctrl_owner to node_sending_reset_command
  set sysctrl_device to port|IP_address_or_hostname_of_device
  set sysctrl_owner_type to tty|network|ipmi
  set is_failsafe to true|false
  set is_cxfs to true|false
  set weight to 0|1
  add nic IP_address_or_hostname (if DNS)
    set heartbeat to true|false
    set ctrl_msgs to true|false
    set priority to integer
  remove nic IP_address_or_hostname (if DNS)
  set hierarchy to [system] [fence][reset][fencereset][shutdown]

```

The commands are the same as those used to define a node. You can change any of the information you specified when defining a node except the node ID. For details about the commands, see "Define a Node with `cmgr`" on page 577.



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

You cannot add a NIC or a network grouping while CXFS services are active (that is, when `start cx_services` has been executed); doing so can lead to cluster malfunction. If services have been started, they should be stopped with `stop cx_services`.

You cannot modify the `operating_system` setting for a node; trying to do so will cause an error. If you have mistakenly specified the incorrect operating system, you must delete the node and define it again.

You cannot modify the node function. To change the node function, you must delete the node and redefine it (and reinstall software products, as needed); the node function for a Solaris or Windows node is always `client_only`.

Example of Partitioning

The following shows an example of partitioning an Origin 3000 system:

```
# cmgr
Welcome to SGI Cluster Manager Command-Line Interface

cmgr> modify node n_preston
Enter commands, when finished enter either "done" or "cancel"

n_preston ? set partition_id to 1
n_preston ? done

Successfully modified node n_preston
```

To perform this function with prompting, enter the following:

```
# cmgr -p
Welcome to SGI Cluster Manager Command-Line Interface

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

Hostname[optional] ? (preston.example.com)
Is this a FailSafe node <true|false> ? (true)
Is this a CXFS node <true|false> ? (true)
Node ID[optional] ? (606)
Partition ID[optional] ? (0) 1
Do you wish to modify failure hierarchy[y/n]:n
Reset type <powerCycle|reset|nmi> ? (powerCycle)
Do you wish to modify system controller info[y/n]:n
Number of Network Interfaces? (1)
NIC 1 - IP Address ? (preston)
NIC 1 - Heartbeat HB (use network for heartbeats) ? (true)
NIC 1 - (use network for control messages) ? (true)
NIC 1 - Priority <1,2,...> ? (1)

Successfully modified node n_preston

cmgr> show node n_preston
Logical Machine Name: n_preston
Hostname: preston.example.com
```

```

Operating System: IRIX
Node Is FailSafe: true
Node Is CXFS: true
Node Function: client_admin
Nodeid: 606
Partition id: 1
Reset type: powerCycle
ControlNet Ipaddr: preston
ControlNet HB: true
ControlNet Control: true
ControlNet Priority: 1

```

To unset the partition ID, use a value of 0 or none.

Changing Failure Hierarchy

The following shows an example of changing the failure hierarchy for the node `perceval` from the system defaults to `fencerreset`, `reset`, `shutdown` and back to the system defaults.



Caution: If you have a cluster with **an even number of server-capable nodes** and **no tiebreaker**: to avoid a split-brain scenario, you should not use the `shutdown` setting for any server-capable node. For a more detailed explanation, see "Shutdown" on page 35.

```

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

Hostname[optional] ? (perceval.example.com)
Is this a FailSafe node <true|false> ? (false)
Is this a CXFS node <true|false> ? (true)
Node ID[optional] ? (803)
Partition ID[optional] ? (0)
Do you wish to modify failure hierarchy[y/n]:y
Hierarchy option 0 <System|FenceReset|Fence|Reset|Shutdown>[optional] ?fencerreset
Hierarchy option 1 <System|FenceReset|Fence|Reset|Shutdown>[optional] ?reset
Hierarchy option 2 <System|FenceReset|Fence|Reset|Shutdown>[optional] ?shutdown
Reset type <powerCycle|reset|nmi> ? (powerCycle)
Do you wish to modify system controller info[y/n]:n

```

G: Reference to cmgr Tasks

```
Number of Network Interfaces ? (1)
NIC 1 - IP Address ? (163.154.18.173)
```

Successfully modified node perceval

```
cmgr> show node perceval
Logical Machine Name: perceval
Hostname: perceval.example.com
Operating System: IRIX
Node Is FailSafe: false
Node Is CXFS: true
Node Function: client_admin
Nodeid: 803
Node Failure Hierarchy is: FenceReset Reset Shutdown
Reset type: powerCycle
ControlNet Ipaddr: 163.154.18.173
ControlNet HB: true
ControlNet Control: true
ControlNet Priority: 1
```

To return to system defaults:

```
cmgr> modify node perceval
```

Enter commands, you may enter "done" or "cancel" at any time to exit

```
Hostname[optional] ? (perceval.example.com)
Is this a FailSafe node <true|false> ? (false)
Is this a CXFS node <true|false> ? (true)
Node ID[optional] ? (803)
Partition ID[optional] ? (0)
Do you wish to modify failure hierarchy[y/n]:y
Hierarchy option 0 <System|FenceReset|Fence|Reset|Shutdown>[optional] ?
(FenceReset) system
Reset type <powerCycle|reset|nmi> ? (powerCycle)
Do you wish to modify system controller info[y/n]:n
Number of Network Interfaces ? (1)
NIC 1 - IP Address ? (163.154.18.173)
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)
```

```
cmgr> show node perceval
Logical Machine Name: perceval
Hostname: perceval.example.com
Operating System: IRIX
Node Is FailSafe: false
Node Is CXFS: true
Node Function: client_admin
Nodeid: 803
Reset type: powerCycle|reset|nmi
ControlNet Ipaddr: 163.154.18.173
ControlNet HB: true
ControlNet Control: true
ControlNet Priority: 1
```

Note: When the system defaults are in place for failure hierarchy, no status is displayed with the show command.

Perform an NMI on a Node with `cmgr`

When CXFS daemons are running, you can perform a nonmaskable interrupt (NMI) on a node with the following command:

```
admin nmi node nodename
```

Note: The `nmi` option is not available for a `dev_type` of `ipmi` and `sysctrl_type` of `bmc`.

This command uses the CXFS daemons to perform an NMI on the specified node.

You can perform an NMI on a node in a cluster even when the CXFS daemons are not running by using the `standalone` option:

```
admin nmi standalone node nodename
```

This command does not go through the CXFS daemons.

If the node has not been defined in the cluster database, you can use the following command line:

```
admin nmi dev_name port|IP_address_or_hostname_of_device of dev_type tty|network|ipmi with sysctrl_type msc|mmsc|l2|l1
```

For example:

```
admin nmi dev_name /dev/ttyIOC0 of dev_type tty with sysctrl_type l2
```

If crsd does not see the response it expects within a certain time, it will issue another NMI, which invalidates the dump. This is determined by the setting of the following values in the cluster database (defaults shown):

```
CrsResetInterval = "20000"  
CrsRetryInterval = "1000"  
CrsResendTimeout = "10000"  
CrsResendRetries = "2"
```

Note: These values can only be changed using advanced tools. If you feel that these values need to be changed, please contact your local SGI support provider.

Convert a Node to CXFS or FailSafe with cmgr

To convert an existing FailSafe node so that it also applies to CXFS, use the `modify` command to change the setting.

Note: You cannot turn off FailSafe or CXFS for a node if the respective high availability (HA) or CXFS services are active. You must first stop the services for the node.

For example, in normal mode:

```
cmgr> modify node cxfs6  
Enter commands, when finished enter either "done" or "cancel"
```

```
cxfs6 ? set is_FailSafe to true  
cxfs6 ? done
```

Successfully modified node cxfs6

For example, in prompting mode:

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

```

Hostname[optional] ? (cxfs6.example.com)
Is this a FailSafe node <true|false> ? (false) true
Is this a CXFS node <true|false> ? (true)
Node ID[optional] ? (13203)
Partition ID[optional] ? (0)
Do you wish to modify failure hierarchy[y/n]:n
Reset type <powerCycle|reset|nmi> ? (powerCycle)
Do you wish to modify system controller info[y/n]:n
Number of Network Interfaces ? (1)
NIC 1 - IP Address ? (163.154.18.172)
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)

Successfully modified node cxfs6

```

Delete a Node with **cmgr**

To delete a node, use the following command:

```
delete node hostname
```

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 `cxfs6-8` with the following configuration:

```

cmgr> show cluster cxfs6-8
Cluster Name: cxfs6-8
Cluster Is FailSafe: true
Cluster Is CXFS: true
Cluster ID: 20
Cluster HA mode: normal
Cluster CX mode: normal

```

```

Cluster cxfs6-8 has following 3 machine(s)
  cxfs6

```

```
cxfs7
cxfs8
```

To delete node cxfs8, you would do the following in prompting mode (assuming that CXFS services have been stopped on the node):

```
cmgr> modify cluster cxfs6-8
Enter commands, when finished enter either "done" or "cancel"
```

```
Is this a FailSafe cluster <true|false> ? (false)
Is this a CXFS cluster <true|false> ? (true)
Cluster Notify Cmd [optional] ?
Cluster Notify Address [optional] ?
Cluster CXFS mode <normal|experimental>[optional] ? (normal)
Cluster ID ? (20)
```

```
Current nodes in cluster cxfs6-8:
Node - 1: cxfs6
Node - 2: cxfs7
Node - 3: cxfs8
```

```
Add nodes to or remove nodes/networks from cluster
Enter "done" when completed or "cancel" to abort
```

```
cxfs6-8 ? remove node cxfs8
cxfs6-8 ? done
Successfully modified cluster cxfs6-8
```

```
cmgr> show cluster cxfs6-8
Cluster Name: cxfs6-8
Cluster Is FailSafe: false
Cluster Is CXFS: true
Cluster ID: 20
Cluster CX mode: normal
```

```
Cluster cxfs6-8 has following 2 machine(s)
cxfs6
cxfs7
```

To delete `cxfs8` from the pool, enter the following:

```
cmgr> delete node cxfs8
```

```
Deleted machine (cxfs6).
```

Display a Node with `cmgr`

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

```
show node hostname
```

For example:

```
cmgr> show node cxfs6
Logical Machine Name: cxfs6
Hostname: cxfs6.example.com
Operating System: IRIX
Node Is FailSafe: false
Node Is CXFS: true
Node Function: server_admin
Nodeid: 13203
Reset type: powerCycle
ControlNet Ipaddr: 163.154.18.172
ControlNet HB: true
ControlNet Control: true
ControlNet Priority: 1
```

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

```
show nodes in pool
```

For example:

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

```
3 Machine(s) defined
  cxfs8
  cxfs6
  cxfs7
```

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

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

For example:

```
cmgr> show nodes in cluster cxf6-8
```

```
Cluster cxf6-8 has following 3 machine(s)
  cxf6
  cxf7
  cxf8
```

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 cxf6-8
cmgr> show nodes
```

```
Cluster cxf6-8 has following 3 machine(s)
  cxf6
  cxf7
  cxf8
```

Test Node Connectivity with cmgr

You can use `cmgr` 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. This test requires that the `/etc/.rhosts` file be configured properly; see "IRIX Modifications for CXFS Connectivity Diagnostics" on page 111, "SGI ProPack Modifications for CXFS Connectivity Diagnostics" on page 124.

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

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

For example:

```
cmgr> test connectivity in cluster cxfs6-8 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.

Test the Serial Connections with `cmgr`

See "System Reset Connection for CXFS Administration Nodes" on page 178.

Cluster Tasks with `cmgr`

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

Define a Cluster with `cmgr`

When you define a cluster with `cmgr`, you define a cluster and add nodes to the cluster with the same command. For general information, see "Define a Cluster with the GUI" on page 224.

Use the following commands to define a cluster:

```
define cluster clustername
    set is_failsafe to true|false
    set is_cxfs to true|false
    set clusterid to clusterID
    set notify_cmd to notify_command
    set notify_addr to email_address
    set ha_mode to normal|experimental
    set cx_mode to normal|experimental
    add node node1name
    add node node2name
    add net network network_address mask netmask
```

Usage notes:

- `cluster` is the logical name of the cluster. Logical names cannot begin with an underscore (`_`) or include any whitespace characters, and can be at most 255 characters. Clusters must have unique names
- If you are running just CXFS, set `is_cxfs` to `true` and `is_failsafe` to `false`. If you are running a coexecution cluster, set both values to `true`.
- `clusterid` is a unique number within your network in the range 1 through 255. The cluster ID is used by the operating system 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 must not change a cluster ID after the cluster has been defined. Clusters must have unique IDs.
- `notify_cmd` is the command to be run whenever the status changes for a node or cluster.
- `notify_addr` is the address to be notified of cluster and node status changes. 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. If you use the `notify_addr` command, you must specify the e-mail program (such as `/usr/sbin/Mail`) as the `notify_command`.
- The `set ha_mode` and `set cx_mode` commands should usually be set to `normal`. The `set cx_mode` command applies only to CXFS, and the `set ha_mode` command applies only to IRIS FailSafe.

- `net` defines a set of NICs into a network. If the highest priority network (beginning with NIC priority 1) fails, the next highest will be used. All NICs within one network must be at the same priority. NICs of a given priority (such as priority 2) cannot be in two separate `net` networks. Although the primary network must be private, the backup network may be public.

If you do not specify a `net` list, the set of priority 1 NICs are used by default as the CXFS heartbeat network and there will be no failover to any other set of NICs.

The `network` parameter specifies an IP network address (such as 1.2.3.0) and the `mask` parameter specifies the subnet mask (such as 255.255.255.0) in decimal notation. The order in which you specify `network` or `mask` is not important.

Note: You cannot add a NIC or a network grouping while CXFS services are active (that is, when `start cx_services` has been executed); doing so can lead to cluster malfunction. If services have been started, they should be stopped with `stop cx_services`.

The following shows the commands with prompting:

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

Is this a FailSafe cluster <true|false> ? true|false
Is this a CXFS cluster <true|false> ? true|false
Cluster Notify Cmd [optional] ?
Cluster Notify Address [optional] ?
Cluster CXFS mode <normal|experimental>[optional] use_default_of_normal
Cluster ID ? cluster_ID
No nodes in cluster clustername

No networks in cluster clustername

Add nodes to or remove nodes/networks from cluster clustername
Enter "done" when completed or "cancel" to abort

clustername ? add node node1name
clustername ? add node node2name
...
clustername ? done
```

Successfully defined cluster *clustername*

Added node *<node1name>* to cluster *<clustername>*

Added node *<node2name>* to cluster *<clustername>*

...

You should set the cluster to the default `normal` mode. Setting the mode to `experimental` turns off heartbeating in the CXFS kernel membership code so that you can debug the cluster without causing node failures. For example, this can be useful if you just want to disconnect the network for a short time (provided that there is no other cluster networking activity, which will also detect a failure even if there is no heartbeating). However, you should never use `experimental` mode on a production cluster and should only use it if directed to by SGI customer support. SGI does not support the use of `experimental` by customers.

For example:

```
cmgr> define cluster cxfs6-8
Enter commands, you may enter "done" or "cancel" at any time to exit

Is this a FailSafe cluster <true|false> ? false
Is this a CXFS cluster <true|false> ? true
Cluster Notify Cmd [optional] ?
Cluster Notify Address [optional] ?
Cluster CXFS mode <normal|experimental>[optional]
Cluster ID ? 20

No nodes in cluster cxfs6-8

No networks in cluster cxfs6-8

Add nodes to or remove nodes/networks from cluster cxfs6-8
Enter "done" when completed or "cancel" to abort

cxfs6-8 ? add node cxfs6
cxfs6-8 ? add node cxfs7
cxfs6-8 ? add node cxfs8
cxfs6-8 ? done
Successfully defined cluster cxfs6-8

Added node <cxfs6> to cluster <cxfs6-8>
```

Added node <cxfs7> to cluster <cxfs6-8>
 Added node <cxfs8> to cluster <cxfs6-8>

To do this without prompting, enter the following:

```
cmgr> define cluster cxfs6-8
Enter commands, you may enter "done" or "cancel" at any time to exit

cluster cxfs6-8? set is_cxfs to true
cluster cxfs6-8? set clusterid to 20
cluster cxfs6-8? add node cxfs6
cluster cxfs6-8? add node cxfs7
cluster cxfs6-8? add node cxfs8
cluster cxfs6-8? done
Successfully defined cluster cxfs6-8
```

Modify a Cluster with cmgr

The commands are as follows:

```
modify cluster clustername
  set is_failsafe to true
  set is_cxfs to true
  set clusterid to clusterID
  set notify_cmd to command
  set notify_addr to email_address
  set ha_mode to normal|experimental
  set cx_mode to normal|experimental
  add node node1name
  add node node2name
  remove node node1name
  remove node node2name
  add net network network_address mask netmask
  remove net network network_address
```

These commands are the same as the `define cluster` commands. For more information, see "Define a Cluster with `cmgr`" on page 597, and "Define a Cluster with the GUI" on page 224.

Note: If you want to rename a cluster, you must delete it and then define a new cluster. If you have started CXFS services on the node, you must either reboot it or reuse the cluster ID number when renaming the cluster.

However, be aware that if you already have CXFS filesystems defined and then rename the cluster, CXFS will not be able to mount the filesystems. For more information, see "Cannot Mount Filesystems" on page 449.

Convert a Cluster to CXFS or FailSafe with `cmgr`

To convert a cluster, use the following commands:

```
modify cluster clustername
  set is_failsafe to true|false
  set is_cxfs to true|false
  set clusterid to clusterID
```

- `cluster` is the logical name of the cluster. Logical names cannot begin with an underscore (`_`) or include any whitespace characters, and can be at most 255 characters.
- If you are running just CXFS, set `is_cxfs` to `true` and `is_failsafe` to `false`. If you are running a coexecution cluster, set both values to `true`.
- `clusterid` is a unique number within your network in the range 1 through 255. The cluster ID is used by the operating system 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 must not change a cluster ID after the cluster has been defined.

For example, to convert CXFS cluster `cxfs6-8` so that it also applies to FailSafe, enter the following:

```
cmgr> modify cluster cxfs6-8
Enter commands, when finished enter either "done" or "cancel"

cxfs6-8 ? set is_failsafe to true
```

The cluster must support all of the functionalities (FailSafe and/or CXFS) that are turned on for its nodes; that is, if your cluster is of type CXFS, then you cannot modify a node that is part of the cluster so that it is of type FailSafe or of type CXFS and FailSafe. However, the nodes do not have to support all the functionalities of the cluster; that is, you can have a node of type CXFS in a cluster of type CXFS and FailSafe.

Delete a Cluster with **cmgr**

To delete a cluster, use the following command:

```
delete cluster clustername
```

However, you cannot delete a cluster that contains nodes; you must first stop CXFS services on the nodes and then redefine the cluster so that it no longer contains the nodes.

For example, in normal mode:

```
cmgr> modify cluster cxfs6-8  
Enter commands, when finished enter either "done" or "cancel"  
  
cxfs6-8 ? remove node cxfs6  
cxfs6-8 ? remove node cxfs7  
cxfs6-8 ? remove node cxfs8  
cxfs6-8 ? done  
Successfully modified cluster cxfs6-8  
  
cmgr> delete cluster cxfs6-8  
  
cmgr> show clusters  
  
cmgr>
```

For example, in prompting mode:

```
cmgr> modify cluster cxfs6-8  
Enter commands, you may enter "done" or "cancel" at any time to exit  
  
Cluster Notify Cmd [optional] ?  
Cluster Notify Address [optional] ?  
Cluster mode <normal|experimental>[optional] ? (normal)
```

```
Cluster ID ? (55)

Current nodes in cluster cxfs6-8:
Node - 1: cxfs6
Node - 2: cxfs7
Node - 3: cxfs8

Add nodes to or remove nodes from cluster cxfs6-8
Enter "done" when completed or "cancel" to abort

cxfs6-8 ? remove node cxfs6
cxfs6-8 ? remove node cxfs7
cxfs6-8 ? remove node cxfs8
cxfs6-8 ? done
Successfully modified cluster cxfs6-8

cmgr> delete cluster cxfs6-8

cmgr> show clusters

cmgr>
```

Display a Cluster with cmgr

To display the clusters and their contents, use the following commands:

```
show clusters
show cluster clustername
```

For example, the following output shows that cluster mycluster has six nodes and two private networks, permitting network failover:

```
cmgr> show cluster mycluster
Cluster Name: mycluster
Cluster Is FailSafe: false
Cluster Is CXFS: true
Cluster ID: 1
Cluster CX mode: normal
```

Cluster mycluster has following 6 machine(s)

```
nodeA
nodeB
nodeC
nodeD
nodeE
nodeF
```

```
CXFS Failover Networks:
network 192.168.0.0, mask 255.255.255.0
network 134.14.54.0, mask 255.255.255.0
```

The multiple networks listed indicates that if the higher priority network should fail, the next priority network will be used. However, the order in which the networks are listed in this output is not an indication of priority. To determine the priority of the networks, you must look at the NIC priorities in the node definition.

Cluster Services Tasks with `cmgr`

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

Start CXFS Services with `cmgr`

To start CXFS services, and set the configuration to automatically restart CXFS services whenever the system is rebooted, use one of the following commands:

```
start cx_services [on node hostname ] for cluster clustername
```

For example, to start CXFS services on all nodes in the cluster:

```
cmgr> start cx_services for cluster cxf6-8
```

Stop CXFS Services with `cmgr`

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

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

```
stop cx_services [on node hostname]for cluster clustername [force]
```

Note: This procedure is only recommended as needed for CXFS administration node because it updates the cluster database and is therefore intrusive to other nodes. When shutting down a CXFS client-only node, do not administratively stop the CXFS services on the node. Rather, let the CXFS services stop by themselves when the client-only node is shut down.

For example:

```
cmgr> stop cx_services on node cxfs6 for cluster cxfs6-8
```

```
CXFS services have been deactivated on node cxfs6 (cluster cxfs6-8)
```

```
cmgr> stop cx_services for cluster cxfs6-8
```

After you have stopped CXFS services in a node, the node is no longer an active member of the cluster.



Caution: If you stop CXFS services, the node will be marked as `INACTIVE` and it will therefore not rejoin the cluster after a reboot. To allow a node to rejoin the cluster, you must restart CXFS services using `cmgr` or the GUI.

Set the Tiebreaker Node with `cmgr`

A *CXFS tiebreaker node* determines whether a CXFS kernel membership quorum is maintained when exactly half of the server-capable nodes can communicate with each other. There is no default CXFS tiebreaker.



Caution: If one of the server-capable nodes is the CXFS tiebreaker in a two server-capable cluster, failure of that node or stopping the CXFS services on that node will result in a cluster-wide forced shutdown. Therefore SGI recommends that you use client-only nodes as tiebreakers so that either server could fail but the cluster would remain operational via the other server.

The reset capability or I/O fencing with switches is **mandatory** to ensure data integrity for all nodes. Clusters should have an odd number of server-capable nodes. If you have an even number of server-capable administration nodes, define a CXFS tiebreaker node. SGI recommends making a client-only node the tiebreaker. (See "CXFS Recovery Issues in a Cluster with Only Two Server-Capable Nodes " on page 526.)

To set the CXFS tiebreaker node, use the modify command as follows:

```
modify cx_parameters
[on node nodename] in cluster clustername
set tie_breaker to hostname
```

To unset the CXFS tiebreaker node, use the following command:

```
set tie_breaker to none
```

For example, in normal mode:

```
cmgr> modify cx_parameters in cluster cxfs6-8
Enter commands, when finished enter either "done" or "cancel"

cxfs6-8 ? set tie_breaker to cxfs8
cxfs6-8 ? done
Successfully modified cx_parameters
```

For example, in prompting mode:

```
cmgr> modify cx_parameters in cluster cxfs6-8

(Enter "cancel" at any time to abort)

Tie Breaker Node ? (cxfs7) cxfs8
Successfully modified cx_parameters

cmgr> show cx_parameters in cluster cxfs6-8
```

```
_CX_TIE_BREAKER=cxfs8
```

Set Log Configuration with `cmgr`

For general information about CXFS logs, see "Set Log Configuration with the GUI" on page 230.

Display Log Group Definitions with `cmgr`

Use the following command to view the log group definitions:

```
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.

Use the following command to see messages logged by a specific daemon on a specific node:

```
show log_group LogGroupName [on node Nodename]
```

To exit from the message display, enter Cntrl-C.

Configure Log Groups with `cmgr`

You can configure a log group with the following command:

```
define log_group log_group on node adminhostname [in cluster clustername]  
  set log_level to log_level  
  add log_file log_file  
  remove log_file log_file
```

Usage notes:

- `log_group` can be one of the following:

```
clconfd  
cli  
crsd  
diags
```

- `log_level` 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
- *log_file*



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

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

```
cmgr> define log_group cli on node cxfs6 in cluster cxfs6-8
```

(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 `cmgr`

Use the following command to modify a log group:

```
modify log_group log_group_name on node hostname [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 cxf6 in cluster cxf6-8
```

(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
```

Revoke Membership of the Local Node with cmgr

To revoke CXFS kernel membership for the local node, such as before the forced CXFS shutdown, enter the following on the local node:

```
admin cxf6_stop
```

This command will be considered as a node failure by the rest of the cluster. The rest of the cluster may then fail due to a loss of CXFS kernel membership quorum, or it may decide to reset the failed node. To avoid the reset, you can modify the node definition to disable the system controller status.

Allow Membership of the Local Node with cmgr

Allowing CXFS kernel membership for the local node permits the node to reapply for CXFS kernel membership. You must actively allow CXFS kernel membership for the local node in the following situations:

- After a manual revocation as in "Revoke Membership of the Local Node with cmgr" on page 610.
- When instructed to by an error message on the console or in /var/adm/SYSLOG.

- After a kernel-triggered revocation. This situation is indicated by the following message in `/var/adm/SYSLOG`:

```
Membership lost - withdrawing from cluster
```

To allow CXFS kernel membership for the local node, use the following command:

```
cmgr> admin cxfs_start
```

See also "Shutdown of the Database and CXFS" on page 332.

CXFS Filesystem Tasks with `cmgr`

This section tells you how to define a filesystem, specify the nodes on which it may or may not be mounted (the *enabled* or *disabled* nodes), and perform mounts and unmounts.

A given filesystem can be mounted on a given node when the following things are true:

- One of the following is true for the node:
 - The default local status is enabled and the node is not in the filesystem's list of explicitly disabled nodes
 - The default local status is disabled and the node is in the filesystem's list of explicitly enabled nodes

See "Define a CXFS Filesystem with `cmgr`" on page 611.

- The global status of the filesystem is enabled. See "Mount a CXFS Filesystem with `cmgr`" on page 618.

Define a CXFS Filesystem with `cmgr`

Use the following commands to define a filesystem and the nodes on which it may be mounted:

```
define cxfs_filesystem logical_filesystem_name [in cluster clustername]  
  set device_name to devicename  
  set mount_point to mountpoint  
  set mount_options to mount_options
```

```
set force to true|false
set dflt_local_status to enabled|disabled
add cxfs_server admin_nodename
    set rank to 0|1|2|...
add enabled_node nodename
add disabled_node nodename
remove cxfs_server admin_nodename
remove enabled_node nodename
remove disabled_node nodename
```

Usage notes:

- Relocation is disabled by default. Recovery and relocation are supported only when using standby nodes. Therefore, you should only define multiple metadata servers for a given filesystem if you are using the standby node model. See "Relocation" on page 20.
- The list of potential metadata servers for any given filesystem must all run the same operating system type.
- `cxfs_filesystem` can be any logical name. Logical names cannot begin with an underscore (`_`) or include any whitespace characters, and can be at most 255 characters.

Note: Within the GUI, the default is to use the last portion of the device name; for example, for a device name of `/dev/cxvm/d761un0s0`, the GUI will automatically supply a logical filesystem name of `d761un0s0`. The GUI will accept other logical names defined with `cmgr` but the GUI will not allow you to modify a logical name; you must use `cmgr` to modify the logical name.

- `device_name` is the device name of an XVM volume that will be shared among all nodes in the CXFS cluster. The name must begin with `/dev/cxvm/`. For more information, see *XVM Volume Manager Administrator's Guide*.
- `mount_point` 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` man page.
- `mount_options` are options that are passed to the `mount` command and are used to control access to the specified XVM volume. For a list of the available options, see the `fstab` man page.

- `force` controls what action CXFS takes if there are processes that have open files or current directories in the filesystem(s) that are to be unmounted. If set to `true`, then the processes will be killed and the unmount will occur. If set to `false`, the processes will not be killed and the filesystem will not be unmounted. The `force` option off (set to `true`) by default.
- `dflt_local_status` defines whether the filesystem can be mounted on all unspecified nodes or cannot be mounted on any unspecified nodes. You can then use the `add enabled_node` or `add disabled_node` commands as necessary to explicitly specify the nodes that differ from the default. There are multiple combinations that can have the same result.

For example, suppose you had a cluster with 10 nodes (`node1` through `node10`). You could use the following methods:

- If you want the filesystem to be mounted on all nodes, and want it to be mounted on any nodes that are later added to the cluster, you would specify:

```
set dflt_local_status to enabled
```

- If you want the filesystem to be mounted on all nodes except `node5`, and want it to be mounted on any nodes that are later added to the cluster, you would specify:

```
set dflt_local_status to enabled
add disabled_node cxf5
```

- If you want the filesystem to be mounted on all nodes except `node5`, and you also do **not** want it to be mounted on any nodes that are later added to the cluster, you would specify:

```
set dflt_local_status to disabled
add enabled_node cxf1
add enabled_node cxf2
add enabled_node cxf3
add enabled_node cxf4
add enabled_node cxf6
add enabled_node cxf7
add enabled_node cxf8
add enabled_node cxf9
add enabled_node cxf10
```

- If you want the filesystem to be mounted on node5 through node10 and on any future nodes, you could specify:

```
set dflt_local_status to enabled
add disabled_node cxfs1
add disabled_node cxfs2
add disabled_node cxfs3
add disabled_node cxfs4
```

To actually mount the filesystem on the enabled nodes, see "Mount a CXFS Filesystem with cmgr" on page 618.

- `cxfs_server` adds or removes the specified CXFS administration node name to the list of potential metadata servers.

Note: After a filesystem has been defined in CXFS, running `mkfs` on it will cause errors to appear in the system log file. To avoid these errors, run `mkfs` before defining the filesystem in CXFS, or delete the CXFS filesystem before running `mkfs`. See "Delete a CXFS Filesystem with cmgr" on page 624.

The following examples shows two potential metadata servers for the `fs1` 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 `cxfs7` (rank 1) will be used. The filesystem is mounted on all nodes.

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

For example, in normal mode:

```
cmgr> define cxfs_filesystem fs1 in cluster cxfs6-8

cxfs_filesystem fs1 ? set device_name to /dev/cxvm/d761un0s0
cxfs_filesystem fs1 ? set mount_point to /mnts/fs1
cxfs_filesystem fs1 ? set force to false
cxfs_filesystem fs1 ? add cxfs_server cxfs6
Enter CXFS server parameters, when finished enter "done" or "cancel"

CXFS server - cxfs6 ? set rank to 0
CXFS server - cxfs6 ? done
```

```
cxfs_filesystem fs1 ? add cxfs_server cxfs7
Enter CXFS server parameters, when finished enter "done" or "cancel"

CXFS server - cxfs7 ? set rank to 1
CXFS server - cxfs7 ? done
cxfs_filesystem fs1 ? set dflt_local_status to enabled
cxfs_filesystem fs1 ? done
Successfully defined cxfs_filesystem fs1

cmgr> define cxfs_filesystem fs2 in cluster cxfs6-8

cxfs_filesystem fs2 ? set device_name to /dev/cxvm/d761un0s1
cxfs_filesystem fs2 ? set mount_point to /mnts/fs2
cxfs_filesystem fs2 ? set force to false
cxfs_filesystem fs2 ? add cxfs_server cxfs8
Enter CXFS server parameters, when finished enter "done" or "cancel"

CXFS server - cxfs8 ? set rank to 0
CXFS server - cxfs8 ? done
cxfs_filesystem fs2 ? set dflt_local_status to enabled
cxfs_filesystem fs2 ? done
Successfully defined cxfs_filesystem fs2
```

For example, in prompting mode:

```
cmgr> define cxfs_filesystem fs1 in cluster cxfs6-8

(Enter "cancel" at any time to abort)

Device ? /dev/cxvm/d761un0s0
Mount Point ? /mnts/fs1
Mount Options[optional] ?
Use Forced Unmount ? <true|false> ? false
Default Local Status <enabled|disabled> ? (enabled)

DEFINE CXFS FILESYSTEM OPTIONS

    0) Modify Server.
    1) Add Server.
    2) Remove Server.
```

G: Reference to cmgr Tasks

- 3) Add Enabled Node.
- 4) Remove Enabled Node.
- 5) Add Disabled Node.
- 6) Remove Disabled Node.
- 7) Show Current Information.
- 8) Cancel. (Aborts command)
- 9) Done. (Exits and runs command)

Enter option:1

No current servers

Server Node ? **cxfs6**
Server Rank ? **0**

- 0) Modify Server.
- 1) Add Server.
- 2) Remove Server.
- 3) Add Enabled Node.
- 4) Remove Enabled Node.
- 5) Add Disabled Node.
- 6) Remove Disabled Node.
- 7) Show Current Information.
- 8) Cancel. (Aborts command)
- 9) Done. (Exits and runs command)

Enter option:1

Server Node ? **cxfs7**
Server Rank ? **1**

- 0) Modify Server.
- 1) Add Server.
- 2) Remove Server.
- 3) Add Enabled Node.
- 4) Remove Enabled Node.
- 5) Add Disabled Node.
- 6) Remove Disabled Node.
- 7) Show Current Information.
- 8) Cancel. (Aborts command)
- 9) Done. (Exits and runs command)

Enter option:9

Successfully defined cxfs_filesystem fs1

cmgr> **define cxfs_filesystem fs2 in cluster cxfs6-8**

(Enter "cancel" at any time to abort)

Device ? **/dev/cxvm/d771un0s1**

Mount Point ? **/mnts/fs2**

Mount Options[optional] ?

Use Forced Unmount ? <true|false> ? **false**

Default Local Status <enabled|disabled> ? (enabled)

DEFINE CXFS FILESYSTEM OPTIONS

- 0) Modify Server.
- 1) Add Server.
- 2) Remove Server.
- 3) Add Enabled Node.
- 4) Remove Enabled Node.
- 5) Add Disabled Node.
- 6) Remove Disabled Node.
- 7) Show Current Information.
- 8) Cancel. (Aborts command)
- 9) Done. (Exits and runs command)

Enter option:1

Server Node ? **cxfs8**

Server Rank ? **0**

- 0) Modify Server.
- 1) Add Server.
- 2) Remove Server.
- 3) Add Enabled Node.
- 4) Remove Enabled Node.
- 5) Add Disabled Node.
- 6) Remove Disabled Node.
- 7) Show Current Information.
- 8) Cancel. (Aborts command)

9) Done. (Exits and runs command)

Enter option:9

Successfully defined cxfs_filesystem fs2

Mount a CXFS Filesystem with cmgr

To mount a filesystem on the enabled nodes, enter the following:

```
admin cxfs_mount cxfs_filesystem logical_filesystem_name [on node nodename] [in cluster clustername]
```

This command enables the *global status* for a filesystem; if you specify the *nodename*, it enables the *local status*. (The global status is only affected if a node name is not specified.) For a filesystem to mount on a given node, both global and local status must be enabled; see "CXFS Filesystem Tasks with cmgr" on page 611.

Nodes must first be enabled by using the `define cxfs_filesystem` and `modify cxfs_filesystem` commands; see "Define a CXFS Filesystem with cmgr" on page 611, and "Modify a CXFS Filesystem with cmgr" on page 619.

For example, to activate the `f1` filesystem by setting the global status to enabled, enter the following:

```
cmgr> admin cxfs_mount cxfs_filesystem fs1 in cluster cxfs6-8
```

The filesystem will then be mounted on all the nodes that have a local status of enabled for this filesystem.

To change the local status to enabled, enter the following:

```
cmgr> admin cxfs_mount cxfs_filesystem fs1 on node cxfs7 in cluster cxfs6-8
```

If the filesystem's global status is disabled, nothing changes. If the filesystem's global status is enabled, the node will mount the filesystem as the result of the change of its local status.

Note: If CXFS 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 services. For more information, see "Start CXFS Services with cmgr" on page 605.

Unmount a CXFS Filesystem with `cmgr`

To unmount a filesystem, enter the following:

```
admin cxfs_unmount cxfs_filesystem filesystemname [on node nodename] [in cluster clustername]
```

Unlike the `modify cxfs_filesystem` command, this command can be run on an active filesystem.

For example, to deactivate the `fs1` filesystem by setting the global status to `disabled`, enter the following:

```
cmgr> admin cxfs_unmount cxfs_filesystem fs1 in cluster cxfs6-8
```

The filesystem will then be unmounted on all the nodes that have a local status of `enabled` for this filesystem.

To change the local status to `disabled`, enter the following:

```
cmgr> admin cxfs_unmount cxfs_filesystem fs1 on node cxfs7 in cluster cxfs6-8
```

If the filesystem's global status is `disabled`, nothing changes. If the filesystem's global status is `enabled`, the node will unmount the filesystem as the result of the change of its local status.

Modify a CXFS Filesystem with `cmgr`

Note: You cannot modify a mounted filesystem.

Use the following commands to modify a filesystem:

```
modify cxfs_filesystem logical_filesystem_name [in cluster clustername]  
  set device_name to devicename  
  set mount_point to mountpoint  
  set mount_options to options  
  set force to true|false  
  set dflt_local_status to enabled|disabled  
  add cxfs_server servername  
    set rank to 0|1|2|...  
  modify cxfs_server servername  
    set rank to 0|1|2|...  
  add enabled_node nodename
```

```
add disabled_node nodename
remove cxfs_server nodename
remove enabled_node nodename
remove disabled_node nodename
```

These are the same commands used to define a filesystem; for more information, see "Define a CXFS Filesystem with cmgr" on page 611.

For example, in normal mode:

```
cmgr> show cxfs_filesystem fs1 in cluster cxfs6-8
```

```
Name: fs1
Device: /dev/cxvm/d761un0s0
Mount Point: /mnts/fs1
Forced Unmount: false
Global Status: disabled
Default Local Status: enabled
```

```
Server Name: cxfs6
Rank: 0
Server Name: cxfs7
Rank: 1
Disabled Client: cxfs8
```

```
cmgr> modify cxfs_filesystem fs1 in cluster cxfs6-8
```

Enter commands, when finished enter either "done" or "cancel"

```
cxfs_filesystem fs3 ? modify cxfs_server cxfs6
```

Enter CXFS server parameters, when finished enter "done" or "cancel"

Current CXFS server cxfs6 parameters:

```
rank :0
CXFS server - cxfs6 ? set rank to 2
CXFS server - cxfs6 ? done
cxfs_filesystem fs1 ? done
```

Successfully modified cxfs_filesystem fs1

```
cmgr> show cxfs_filesystem fs1 in cluster cxfs6-8
```

```
Name: fs1
Device: /dev/cxvm/d761un0s0
```

```
Mount Point: /mnts/fs1
Forced Unmount: false
Global Status: disabled
Default Local Status: enabled
```

```
Server Name: cxfs6
Rank: 2
Server Name: cxfs7
Rank: 1
Disabled Client: cxfs8
```

In prompting mode:

```
cmgr> show cxfs_filesystem fs1 in cluster cxfs6-8
```

```
Name: fs1
Device: /dev/cxvm/d761un0s0
Mount Point: /mnts/fs1
Forced Unmount: false
Global Status: disabled
Default Local Status: enabled
```

```
Server Name: cxfs6
Rank: 0
Server Name: cxfs7
Rank: 1
Disabled Client: cxfs8
```

```
cmgr> modify cxfs_filesystem fs1 in cluster cxfs6-8
```

(Enter "cancel" at any time to abort)

```
Device ? (/dev/cxvm/d761un0s0)
Mount Point ? (/mnts/fs1)
Mount Options[optional] ?
Use Forced Unmount ? <true|false> ? (false)
Default Local Status <enabled|disabled> ? (enabled)
```

MODIFY CXFS FILESYSTEM OPTIONS

- 0) Modify Server.
- 1) Add Server.

- 2) Remove Server.
- 3) Add Enabled Node.
- 4) Remove Enabled Node.
- 5) Add Disabled Node.
- 6) Remove Disabled Node.
- 7) Show Current Information.
- 8) Cancel. (Aborts command)
- 9) Done. (Exits and runs command)

Enter option:0

Current servers:

CXFS Server 1 - Rank: 0	Node: cxfs6
CXFS Server 2 - Rank: 1	Node: cxfs7

Server Node ? **cxfs6**

Server Rank ? (0) **2**

- 0) Modify Server.
- 1) Add Server.
- 2) Remove Server.
- 3) Add Enabled Node.
- 4) Remove Enabled Node.
- 5) Add Disabled Node.
- 6) Remove Disabled Node.
- 7) Show Current Information.
- 8) Cancel. (Aborts command)
- 9) Done. (Exits and runs command)

Enter option:7

Current settings for filesystem (fs1)

CXFS servers:

Rank 2	Node cxfs6
Rank 1	Node cxfs7

Default local status: enabled

No explicitly enabled clients

```

Explicitly disabled clients:
    Disabled Node: cxfs8

    0) Modify Server.
    1) Add Server.
    2) Remove Server.
    3) Add Enabled Node.
    4) Remove Enabled Node.
    5) Add Disabled Node.
    6) Remove Disabled Node.
    7) Show Current Information.
    8) Cancel. (Aborts command)
    9) Done. (Exits and runs command)

Enter option:9
Successfully modified cxfs_filesystem fs3

```

Relocate the Metadata Server for a Filesystem with `cmgr`

If relocation is explicitly enabled in the kernel with the `cxfs_relocation_ok` systune (see "Relocation" on page 20), you can relocate a metadata server to another node using the following command if the filesystem must be mounted on the system that is running `cmgr`:

```
admin cxfs_relocate cxfs_filesystem filesystem_name to node nodename [in cluster clustername]
```

Note: This function is only available on a live system.

To relocate the metadata server from `cxfs6` to `cxfs7` for `fs1` in cluster `cxfs6-8`, enter the following:

```
cmgr> admin cxfs_relocate cxfs_filesystem fs1 to node cxfs7 in cluster cxfs6-8
```

CXFS kernel membership is not affected by relocation. However, users may experience a degradation in filesystem performance while the metadata server is relocating.

For more details, see "Modify a CXFS Filesystem with `cmgr`" on page 619.

Delete a CXFS Filesystem with `cmgr`

Use the following command to delete a filesystem:

```
delete cxfs_filesystem filesystemname [in cluster clustername]
```

For example:

```
cmgr> delete cxfs_filesystem fs2 in cluster cxfs6-8
```

Switches and I/O Fencing Tasks with `cmgr`

The following tasks let you configure switches and I/O fencing. For general information, see "Fence" on page 30.

Note: Nodes without system controllers require I/O fencing to protect data integrity. A switch is mandatory to support I/O fencing; therefore, multiOS CXFS clusters require a switch. See the release notes for supported switches.

Define a Switch with `cmgr`

This section describes how to use the `cmgr` command to define a new Brocade switch to support I/O fencing in a cluster.

Note: To define a switch other than a Brocade switch, such as a QLogic switch, you must use the GUI or the `cxfs_admin` or `hafence(1M)` commands. (You cannot use the `cmgr` command to completely define a switch other than Brocade.) See "Create a Switch with `cxfs_admin`" on page 296 and "Using `hafence` to Manipulate a Switch" on page 318.

To define a new Brocade switch, use the following command:

```
define switch switch_hostname username username password password [mask mask]
```

Usage notes:

- `switch` specifies the hostname of the Fibre Channel switch; this is used to determine the IP address of the switch.

- `username` specifies the user name to use when sending a telnet message to the switch.
- `password` specifies the password for the specified *username*.
- `mask` specifies one of the following:
 - A list of ports in the switch that will never be fenced. The list has the following form, beginning with the # symbol and separating each port number with a comma::

```
#port ,port ,port . . .
```

Each *port* is a decimal integer in the range 0 through 1023. For example, the following indicates that port numbers 2, 4, 5, 6, 7, and 23 will never be fenced:

```
#2 , 4 , 5 , 6 , 7 , 23
```

- A hexadecimal string that represents ports in the switch that will never be fenced. Ports are numbered from 0. If a given bit has a binary value of 0, the port that corresponds to that bit is eligible for fencing operations; if 1, then the port that corresponds to that bit will always be excluded from any fencing operations. For an example, see Figure 11-5 on page 235.

CXFS administration nodes automatically discover the available HBAs and, when fencing is triggered, fence off all of the Fibre Channel HBAs when the `Fence` or `FenceReset` fail action is selected. However, masked HBAs will not be fenced. Masking allows you to prevent the fencing of devices that are attached to the SAN but are not shared with the cluster, to ensure that they remain available regardless of CXFS status. You would want to mask HBAs used for access to tape storage, or HBAs that are only ever used to access local (nonclustered) devices.

For example, using the direct port-number specification method:

```
cmgr> define switch ptg-brocade username admin password password mask #2,4,5,7,65
```

Or, using the hexadecimal bitmask method:

```
cmgr> define switch ptg-brocade username admin password password mask 200000000000000F4
```

Modify a Switch Definition with cmgr

To modify the user name, password, or mask for a Brocade switch, use the following command:

```
modify switch switch_hostname username username password password [mask mask]
```

The arguments are the same as for "Define a Switch with cmgr" on page 624.

Note: To modify the definition of another type of switch, such as QLogic, you must use the GUI, `hafence(1M)`, or `cxfs_admin(1M)` commands. See "Using hafence to Manipulate a Switch" on page 318.

For example, to change the mask for switch `ptg-brocade` from `A4` to `0` (which means that all of the ports on the switch are eligible for fencing), enter the following:

```
cmgr> modify switch ptg-brocade username admin password password mask 0
```

Raise the I/O Fence for a Node with cmgr

Raising an I/O fence isolates the node from the SAN; CXFS sends a message via the telnet protocol to the switch and disables the port. After the node is isolated, it cannot corrupt data in the shared CXFS filesystem. Use the following command:

```
admin fence raise [node nodename]
```

nodename is the name of the node to be isolated.

For example, to isolate the default node, enter the following:

```
cmgr> admin fence raise
```

To isolate node `Node3`, enter the following:

```
cmgr> admin fence raise node Node3
```

Lower the I/O Fence for a Node with `cmgr`

To lower the I/O fence for a given node in order to reenble the port, allowing the node to connect to the SAN and access the shared CXFS filesystem, use the following command:

```
admin fence lower [node nodename]
```

nodename is the name of the node to be reconnected.

For example, to provide access for the default node, enter the following:

```
cmgr> admin fence lower
```

To provide access for node Node3, enter the following:

```
cmgr> admin fence lower node Node3
```

Update Switch Port Information with `cmgr`

To update the mappings in the cluster database between the host bus adapters (HBAs) and switch ports, use the following command:

```
admin fence update
```

You should run this command if you reconfigure any switch or add ports.

Delete a Switch Definition with `cmgr`

To delete a switch, use the following command:

```
delete switch switch_hostname
```

switch_hostname is the hostname of the Fibre Channel switch; this is used to determine the IP address of the switch.

For example:

```
cmgr> delete switch ptg-brocade  
Successfully updated switch config.
```

Show Switches with cmgr

To display the switches in the system, use the following command:

```
show switches
```

To show the switches for a given node, use the following command:

```
show switch hostname
```

For example:

```
cmgr> show switch ptg-brocade
Switch[0]
*Hostname ptg-brocade Username admin Password password Mask 0
Vendor BROCADE Number of ports 8
  0 0000000000000000 Reset
  1 210000e08b0102c6 Reset
  2 210000e08b01fec5 Reset
  3 210000e08b019dc5 Reset
  4 210000e08b0113ce Reset
  5 210000e08b027795 Reset thump
  6 210000e08b019ef0 Reset
  7 210000e08b022242 Reset
```

Query Switch Status with cmgr

To query the status of each port on the switch, use the following command:

```
admin fence query
```

For example:

```
cmgr> admin fence query
Switch[0] "brocade04" has 16 ports
Port 4 type=FABRIC status=enabled hba=210000e08b0042d8 on host o200c
Port 5 type=FABRIC status=enabled hba=210000e08b00908e on host cxfs30
Port 9 type=FABRIC status=enabled hba=2000000173002d3e on host cxfssun3
```

For more verbose display, (which shows all ports on the switch, rather than only those attached to nodes in the default cluster), use the following command:

```
admin fence query verbose
```

For example:

```
cmgr> admin fence query verbose
Switch[0] "brocade04" has 16 ports
Port 0 type=FABRIC status=enabled hba=2000000173003b5f on host UNKNOWN
Port 1 type=FABRIC status=enabled hba=2000000173003adf on host UNKNOWN
Port 2 type=FABRIC status=enabled hba=210000e08b023649 on host UNKNOWN
Port 3 type=FABRIC status=enabled hba=210000e08b021249 on host UNKNOWN
Port 4 type=FABRIC status=enabled hba=210000e08b0042d8 on host o200c
Port 5 type=FABRIC status=enabled hba=210000e08b00908e on host cxfs30
Port 6 type=FABRIC status=enabled hba=2000000173002d2a on host UNKNOWN
Port 7 type=FABRIC status=enabled hba=2000000173003376 on host UNKNOWN
Port 8 type=FABRIC status=enabled hba=2000000173002c0b on host UNKNOWN
Port 9 type=FABRIC status=enabled hba=2000000173002d3e on host cxfssun3
Port 10 type=FABRIC status=enabled hba=2000000173003430 on host UNKNOWN
Port 11 type=FABRIC status=enabled hba=200900a0b80c13c9 on host UNKNOWN
Port 12 type=FABRIC status=disabled hba=0000000000000000 on host UNKNOWN
Port 13 type=FABRIC status=enabled hba=200d00a0b80c2476 on host UNKNOWN
Port 14 type=FABRIC status=enabled hba=1000006069201e5b on host UNKNOWN
Port 15 type=FABRIC status=enabled hba=1000006069201e5b on host UNKNOWN
```

Script Example

The following script defines a three-node cluster of type CXFS. The nodes are of type CXFS.

Note: This example only defines one network interface. The hostname is used here for simplicity; however, you may wish to use the IP address instead to avoid confusion. This example does not address the system controller definitions.

```
#!/usr/cluster/bin/cmgr -if
#
#Script to define a three-node cluster

define node cxfs6
    set hostname to cxfs6
    set is_cxfs to true
    set operating_system to irix
    set node_function to server_admin
    add nic cxfs6
        set heartbeat to true
        set ctrl_msgs to true
        set priority to 1
    done
done

define node cxfs7
    set hostname to cxfs7
    set is_cxfs to true
    set operating_system to irix
    set node_function to server_admin
    add nic cxfs7
        set heartbeat to true
        set ctrl_msgs to true
        set priority to 1
    done
done

define node cxfs8
    set hostname to cxfs8
    set is_cxfs to true
    set operating_system to irix
    set node_function to server_admin
    add nic cxfs8
        set heartbeat to true
        set ctrl_msgs to true
        set priority to 1
    done
done
```

```
define cluster cxfs6-8
    set is_cxfs to true
    set is_failsafe to true
    set clusterid to 20
    add node cxfs6
    add node cxfs7
    add node cxfs8
done
quit
```

After running this script, you would see the following output:

```
Successfully defined node cxfs6
```

```
Successfully defined node cxfs7
```

```
Successfully defined node cxfs8
```

```
Successfully defined cluster cxfs6-8
```

The following script defines two filesystems; fs1 is mounted on all but node cxfs8, and fs2 is mounted on all nodes:

```
#!/usr/cluster/bin/cmgr -if
# Script to define two filesystems
# Define fs1, do not mount on cxfs8
define cxfs_filesystem fs1 in cluster cxfs6-8
set device_name to /dev/cxvm/d76lun0s0
set mount_point to /mnts/fs1
set force to false
add cxfs_server cxfs6
    set rank to 0
done
add cxfs_server cxfs7
    set rank to 1
done
set dflt_local_status to enabled
add disabled_node cxfs8
done
#
# Define fs2, mount everywhere
define cxfs_filesystem fs2 in cluster cxfs6-8
```

```
set device_name to /dev/cxvm/d76lun0s1
set mount_point to /mnts/fs2
set force to false
add cxfs_server cxfs8
set rank to 0
done
set dflt_local_status to enabled
done
```

Creating a cmgr Script Automatically

After you have configured the cluster database, you can use the `build_cmgr_script` command to automatically create a cmgr script based on the contents of the cluster database. The generated script will contain the following:

- Node definitions
- Cluster definition
- Switch definitions
- CXFS filesystem definitions
- Parameter settings
- Any changes made using either the `cmgr` command or the GUI
- FailSafe information (in a coexecution cluster only)

As needed, you can then use the generated script to recreate the cluster database after performing a `cdbreinit`.

Note: You must execute the generated script on the first node that is listed in the script. If you want to execute the generated script on a different node, you must modify the script so that the node is the first one listed.

By default, the generated script is named:

```
/var/cluster/ha/tmp/cmgr_create_cluster_clustername_processID
```

You can specify an alternative pathname by using the `-o` option:

```
build_cmgr_script [-o script_pathname]
```

For more details, see the `build_cmgr_script` man page.

For example:

```
# /var/cluster/cmgr-scripts/build_cmgr_script -o /tmp/newcdb
Building cmgr script for cluster clusterA ...
build_cmgr_script: Generated cmgr script is /tmp/newcdb
```

The example script file contents are as follows; note that because `nodeE` is the first node defined, you must execute the script on `nodeE`:

```
#!/usr/cluster/bin/cmgr -f

# Node nodeE definition
define node nodeE
    set hostname to nodeE.example.com
    set operating_system to IRIX
    set is_failsafe to false
    set is_cxfs to true
    set node_function to server_admin
    set nodeid to 5208
    set reset_type to powerCycle
    add nic nodeE
        set heartbeat to true
        set ctrl_msgs to true
        set priority to 1
    done
done

# Node nodeD definition
define node nodeD
    set hostname to nodeD.example.com
    set operating_system to IRIX
    set is_failsafe to false
    set is_cxfs to true
    set node_function to server_admin
    set nodeid to 5181
    set reset_type to powerCycle
    add nic nodeD
        set heartbeat to true
        set ctrl_msgs to true
        set priority to 1
```

```
done
done

# Node nodeF definition
define node nodeF
    set hostname to nodeF.example.com
    set operating_system to IRIX
    set is_failsafe to false
    set is_cxfs to true
    set node_function to server_admin
    set nodeid to 5401
    set reset_type to powerCycle
    add nic nodeF
        set heartbeat to true
        set ctrl_msgs to true
        set priority to 1
done
done

# Define cluster and add nodes to the cluster
define cluster clusterA
    set is_failsafe to false
    set is_cxfs to true
    set cx_mode to normal
    set clusterid to 35
done

modify cluster clusterA
    add node nodeD
    add node nodeF
    add node nodeE
done

set cluster clusterA

define cxfs_filesystem fs1
    set device_name to /dev/cxvm/fs1
    set mount_point to /fs1
    set force to false
    set dflt_local_status to enabled
    add cxfs_server nodeE
```

```
        set rank to 1
    done
    add cxfs_server nodeD
        set rank to 2
    done
    add cxfs_server nodeF
        set rank to 0
    done
done

define cxfs_filesystem fs2
    set device_name to /dev/cxvm/fs2
    set mount_point to /fs2
    set force to false
    set dflt_local_status to enabled
    add cxfs_server nodeE
        set rank to 1
    done
    add cxfs_server nodeD
        set rank to 2
    done
    add cxfs_server nodeF
        set rank to 0
    done
done

define cxfs_filesystem fs2
    set device_name to /dev/cxvm/fs2
    set mount_point to /fs2
    set force to false
    set dflt_local_status to enabled
    add cxfs_server nodeE
        set rank to 1
    done
    add cxfs_server nodeD
        set rank to 2
    done
    add cxfs_server nodeF
        set rank to 0
    done
done
```

```
# Setting CXFS parameters
modify cx_parameters
    set tie_breaker to none
done

quit
```

Troubleshooting cmgr

You should only use `cmgr` when you are logged in as `root`.

The following message may appear in `/var/cluster/ha/log/cli_hostname` if the underlying command line interface (CLI) was invoked by a login other than `root:CI_IPCERR_AGAIN, ipccnt_connect(): file /var/cluster/ha/comm/clconfd-ipc_cxfs0 lock failed - Permission denied`

Additional cmgr Examples

This section contains the following:

- "Example of Normal CXFS Shutdown Using `cmgr`" on page 637
- "Example of Forced CXFS Shutdown Using `cmgr`" on page 637
- "Example of Rejoining the Cluster after a Stopping CXFS Services Using `cmgr`" on page 638
- "Example of Rejoining the Cluster after a Forced CXFS Shutdown Using `cmgr`" on page 638
- "Example of Configuring Private Network Failover Using `cmgr`" on page 638
- "Example of Configuring a Large Cluster Using `cmgr`" on page 645
- "Example of Performing a Forced CXFS Shutdown Using `cmgr`" on page 646
- "Example of Relocation Error Using `cmgr`" on page 646

- "Example of Migration from an IRIX Cluster to an SGI ProPack Cluster Using `cmgr`" on page 646
- "Example of Querying Node Status Using `cmgr`" on page 654
- "Example of Pinging the System Controller Using `cmgr`" on page 654
- "Example of Monitoring Reset Lines Using `cmgr`" on page 655
- "Example of I/O Fencing Status Using `cmgr`" on page 655
- "Example of Using `build_cmgr_script` to Recreate the Cluster Database" on page 655

Also see:

- "System Reset Connection for CXFS Administration Nodes" on page 178
- "Perform a Power Cycle on a Node" on page 491
- "Reset a Node" on page 491

Example of Normal CXFS Shutdown Using `cmgr`

To perform a normal CXFS shutdown, for example enter the following `cmgr` command:

```
cmgr> stop cx_services on node nodename for cluster clustername
```

This action deactivates CXFS services on **one** node, forming a new CXFS kernel membership after deactivating the node. If you want to stop CXFS services on multiple nodes, you must enter this command multiple times or perform the task using the GUI.

After you stop CXFS services on a node, the node is marked as inactive and is no longer used when calculating the CXFS kernel membership. See "Node Status" on page 406.

Example of Forced CXFS Shutdown Using `cmgr`

To perform an administrative stop, enter the following `cmgr` command to revoke the CXFS kernel membership of the local node:

```
cmgr> admin cxfs_stop
```

This action can also be triggered automatically by the kernel after a loss of CXFS kernel membership quorum.

Example of Rejoining the Cluster after a Stopping CXFS Services Using `cmgr`

The node will not rejoin the cluster after a reboot. The node will rejoin the cluster only when CXFS services are explicitly reactivated with the CXFS GUI (see "Start CXFS Services with the GUI" on page 227) or the following command:

```
cmgr> start cx_services on node nodename for cluster clustername
```

In `cxfs_admin`, you can disable individual nodes with the `disable` command.

Example of Rejoining the Cluster after a Forced CXFS Shutdown Using `cmgr`

After a forced CXFS shutdown, the local node will not resume CXFS kernel membership until the node is rebooted or until you explicitly allow CXFS kernel membership for the local node for example by entering the following `cmgr` command:

```
cmgr> admin cxfs_start
```

Example of Configuring Private Network Failover Using `cmgr`

This section provides an example of modifying a cluster to provide private network failover by using the `cmgr` command.

Suppose your cluster has the following configuration:

```
irix# cxfs-config
Global:
  cluster: mycluster (id 1)
  cluster state: enabled
  tiebreaker: yellow

Networks:

Machines:
...
  node red: node 55    cell 4  enabled  IRIX    server_admin
             hostname: red.example.com
```

```

fail policy: Fence, Shutdown
nic 0: address: 192.168.0.1 priority: 1

node yellow: node 2    cell 3  enabled  IRIX    server_admin
hostname: yellow.example.com
fail policy: Fence, Shutdown
nic 0: address: 192.168.0.2 priority: 1

```

To change the configuration to support private network failover, you would do the following:

1. Ensure that CXFS services are not active.

Note: You cannot add a NIC or a network grouping while CXFS services are active (that is, when `start cx_services` has been executed); doing so can lead to cluster malfunction.

If services have been started, stopped them as follows:

```

[root@linux root]# cmgr -p
Welcome to SGI Cluster Manager Command-Line Interface

```

```

cmgr> stop cx_services for cluster mycluster

```

```

CXFS services have been deactivated in cluster
mycluster

```

2. Add another set of NICs to support a second CXFS network. (The second network will be used as the failover network and can be the public network and thus does not have to be a second CXFS private network.) For example:

```

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

Hostname[optional] ? (red.example.com)
Is this a FailSafe node <true|false> ? (false)
Is this a CXFS node <true|false> ? (true)
Partition ID[optional] ? (0)
Do you wish to modify failure hierarchy[y/n]:n
Reset type <powerCycle|reset|nmi> ? (powerCycle)

```

```
Do you wish to modify system controller
info[y/n]:n
Number of Network Interfaces ? (1) 2
NIC 1 - IP Address ? (192.168.0.1)
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)
NIC 2 - IP Address ? 192.168.1.1
NIC 2 - Heartbeat HB (use network for heartbeats) <true|false> ?
true
NIC 2 - (use network for control messages) <true|false> ? true
NIC 2 - Priority <1,2,...> ? 2
```

Successfully modified node red

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

```
Hostname[optional] ? (yellow.example.com)
Is this a FailSafe node <true|false> ? (false)
Is this a CXFS node <true|false> ? (true)
Partition ID[optional] ? (0)
Do you wish to modify failure hierarchy[y/n]:n
Reset type <powerCycle|reset|nmi> ? (powerCycle)
Do you wish to modify system controller
info[y/n]:n
Number of Network Interfaces ? (1) 2
NIC 1 - IP Address ? (192.168.0.2)
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)
NIC 2 - IP Address ? 192.168.1.2
NIC 2 - Heartbeat HB (use network for heartbeats) <true|false> ?
true
NIC 2 - (use network for control messages) <true|false> ? true
NIC 2 - Priority <1,2,...> ? 2
```

Successfully modified node yellow

Repeat this process for each node. You can use the `cxfs-config` command to display the defined NICs. For example:

```
irix# cxfs-config
...
node red: node 55    cell 4  enabled  IRIX    server_admin
hostname: red.example.com
fail policy: Fence, Shutdown
nic 0: address: 192.168.0.1 priority: 1
nic 1: address: 192.168.1.1 priority: 2

node yellow: node 2    cell 3  enabled  IRIX    server_admin
hostname: yellow.example.com
fail policy: Fence, Shutdown
nic 0: address: 192.168.0.2 priority: 1
nic 1: address: 192.168.1.2 priority: 2
```

3. Configure the NICs into networks. (CXFS will ignore NICs other than priority 1 unless you configure the NICs into networks.)

a. Configure the primary network:

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

Is this a FailSafe cluster <true|false> ? (false)
Is this a CXFS cluster <true|false> ? (true)
Cluster Notify Cmd [optional] ?
Cluster Notify Address [optional] ?
Cluster CXFS mode <normal|experimental>[optional] ? (normal)
Cluster ID ? (1)

Current nodes in cluster mycluster:
Node - 1: green
Node - 2: orange
Node - 3: red
Node - 4: purple
Node - 5: yellow
Node - 6: blue

No networks in cluster mycluster
```

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

```
mycluster ? add net network 192.168.0.0 mask 255.255.255.0  
mycluster ? done  
Successfully modified cluster mycluster
```

At this point, cxfs-config will show the primary network (network 0):

```
irix#  
cxfs-config  
...  
Networks:  
  net 0: type tcpip  192.168.0.0      255.255.255.0  
  
Machines:  
...  
  node red: node 55   cell 4  enabled  IRIX   server_admin  
    hostname: red.example.com  
    fail policy: Fence, Shutdown  
    nic 0: address: 192.168.0.1 priority: 1 network: 0  
    nic 1: address: 192.168.1.1 priority: 2 network: none  
  
  node yellow: node 2   cell 3  enabled  IRIX   server_admin  
    hostname: yellow.example.com  
    fail policy: Fence, Shutdown  
    nic 0: address: 192.168.0.2 priority: 1 network: 0  
    nic 1: address: 192.168.1.2 priority: 2 network: none  
...
```

b. Configure the secondary network:

```
cmgr> modify cluster mycluster  
Enter commands, you may enter "done" or "cancel" at any time to exit  
  
Is this a FailSafe cluster <true|false> ? (false)  
Is this a CXFS cluster <true|false> ? (true)  
Cluster Notify Cmd [optional] ?  
Cluster Notify Address [optional] ?  
Cluster CXFS mode <normal|experimental>[optional] ? (normal)
```

Cluster ID ? (1)

Current nodes in cluster mycluster:

Node - 1: green
Node - 2: orange
Node - 3: red
Node - 4: purple
Node - 5: yellow
Node - 6: blue

No networks in cluster mycluster

Add nodes to or remove nodes/networks from cluster mycluster

Enter "done" when completed or "cancel" to abort

mycluster ? **add net network 192.168.1.0 mask 255.255.255.0**

mycluster ? **done**

Successfully modified cluster mycluster

The `cxfs-config` command will now display the secondary network (network 1):

irix#

cxfs-config

...

Networks:

net 0:	type tcpip	192.168.0.0	255.255.255.0
net 1:	type tcpip	192.168.1.0	255.255.255.0

Machines:

...

node red:	node 55	cell 4	enabled	IRIX	server_admin
hostname: red.example.com					
fail policy: Fence, Shutdown					
nic 0:	address:	192.168.0.1	priority:	1	network: 0
nic 1:	address:	192.168.1.1	priority:	2	network: 1

node yellow:	node 2	cell 3	enabled	IRIX	server_admin
hostname: yellow.example.com					
fail policy: Fence, Shutdown					
nic 0:	address:	192.168.0.2	priority:	1	network: 0

```
nic 1: address: 192.168.1.2 priority: 2 network:
```

1

When you restart `cx_services`, the first membership delivered message will appear:

```
NOTICE: Membership delivered for cells 0x14.  
Cell(age): 3(1) 4(1)
```

To delete a network, do the following:

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

```
Is this a FailSafe cluster <true|false> ? (false)  
Is this a CXFS cluster <true|false> ? (true)  
Cluster Notify Cmd [optional] ?  
Cluster Notify Address [optional] ?  
Cluster CXFS mode <normal|experimental>[optional] ? (normal)  
Cluster ID ? (1)
```

Current nodes in cluster mycluster:

```
Node - 1: green  
Node - 2: orange  
Node - 3: red  
Node - 4: purple  
Node - 5: yellow  
Node - 6: blue
```

Current networks in cluster mycluster:

```
Network 0 - network 192.168.0.0, mask 255.255.255.0  
Network 1 - network 192.168.1.0, mask 255.255.255.0
```

```
cmgr> modify cluster mycluster  
Enter commands, when finished enter either "done" or "cancel"
```

```
mycluster ? remove net network 192.168.1.0  
mycluster ? done  
Successfully modified cluster mycluster
```

While there are networks defined, the cluster will try to use the highest priority network and failover as needed to the lower priority networks as possible. Deleting

all networks will return the cluster to the default mode, in which a network consisting only of the priority 1 NICs is used.

For more information, see "Define a Cluster with `cmgr`" on page 597 and "Modify a Cluster with `cmgr`" on page 601.

Example of Configuring a Large Cluster Using `cmgr`

Following is an example `cmgr` script for configuring a one-node cluster that can be copied and repeated for the number of nodes required:

```
#!/usr/cluster/bin/cmgr -f
# Node nodename definition
define node nodename
    set hostname to nodename
    set operating_system to OS
    set node_function to server_admin|client_admin|client_only
    set is_failsafe to false
    set is_cxfs to true
    set nodeid to nodeID#
    set hierarchy to [system][fence][reset][fencereset][shutdown]
    set reset_type to powerCycle|reset|nmi
    add nic IP address or nodename
        set heartbeat to true
        set ctrl_msgs to true
        set priority to 1
    done
done
# Define cluster and add nodes to the cluster
define cluster clustername
    set is_failsafe to false
    set is_cxfs to true
    set cx_mode to normal
    set clusterid to clusterID#
done
modify cluster clustername
    add node nodename
done
set cluster clustername
define cxfs_filesystem filesystemname
    set device_name to /dev/cxvm/volumename
```

```
        set mount_point to /mountpoint
        set force to false
        set dflt_local_status to enabled
        add cxfs_server server1, server2, etc
            set rank to 0
        done
done
# Setting CXFS parameters
modify cx_parameters
    set tie_breaker to none
done
start cx_services for cluster clustername
quit
```

Example of Performing a Forced CXFS Shutdown Using cmgr

```
# /usr/cluster/bin/cmgr -p
cmgr> admin cxfs_stop
```

Example of Relocation Error Using cmgr

If you try to relocate a filesystem and see an error similar to the following cmgr example, it means that relocation has not been enabled: :

```
CMD(/bin/mount -n -o remount,set_server=node1 /lsan1): exited with status
32 (0x20)

Failed to admin:
    cxfs_relocate

admin command failed
```

To allow the relocation to occur, you must enable relocation as specified in "Relocation" on page 20.

Example of Migration from an IRIX Cluster to an SGI ProPack Cluster Using cmgr

CXFS supports a running cluster with a single type of operating system for administration nodes: either all IRIX or all SGI ProPack. To migrate from an IRIX

cluster to an SGI ProPack cluster, follow the instruction in this chapter. For assistance, contact SGI Managed Services.

You should also be aware the differences between basic IRIX and Linux system administration. See "An Overview of Differences Between IRIX and Linux System Administration" on page 415.

Note: The following procedure assumes that the filesystems in the cluster you want to migrate do not have block sizes greater than the system page size and that they are not real-time filesystems. These types of filesystems are supported on IRIX but not on SGI ProPack.

This example uses `cmgr`, but you could perform a similar procedure using `cxfs_admin` or the GUI. The example begins with a cluster named `performance` having a two IRIX server-capable nodes named `rum` and `snake` and a Solaris client-only node named `ray`:

```
rum # clconf_info
```

```
Event at [2004-02-13 07:57:17]
```

```
Membership since Thu Feb 12 15:15:26 2004
```

Node	NodeID	Status	Age	CellID
snake	1	up	2	1
rum	2	up	2	2
ray	3	up	1	0

```
1 CXFS FileSystems
```

```
/dev/cxvm/V9500 on /cxfs/V9500 enabled server=(snake) 2 client(s)=(ray,rum) status=UP
```

Do the following:

1. Unmount the CXFS filesystems cluster-wide within CXFS. For example:

```
cmgr> admin cxfs_unmount cxfs_filesystem V9500

cxfs_unmount operation successful
```

2. Mount and unmount the filesystems locally, which will ensure that the XFS log plays back cleanly. For example:

```
# mount /dev/cxvm/V9500 /mnt
# umount /mnt
```

3. Stop CXFS services on all nodes. For example on the IRIX node run:

```
cmgr> stop cx_services for cluster performance

CXFS services have been deactivated in cluster performance
```

Note: If you use `cxfs_admin`, you must issue a command for each node.

4. Define the administration node with the SGI ProPack operating system type. For example on the IRIX node run:

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

Hostname[optional] ? bang
Is this a FailSafe node <true|false> ? false
Is this a CXFS node <true|false> ? true
Operating System <IRIX|Linux32|Linux64|AIX|HPUX|Solaris|MacOSX|Windows> ? Linux64
Node Function <server_admin|client_admin|client_only> ? server_admin
Node ID[optional] ? 64
Partition ID[optional] ? (0)
Do you wish to define failure hierarchy[y/n]:y
Hierarchy option 0 <System|FenceReset|Fence|Reset|Shutdown>[optional] ? Fence
Hierarchy option 1 <System|FenceReset|Fence|Reset|Shutdown>[optional] ?
Reset type <powerCycle|reset|nmi> ? (powerCycle)
Do you wish to define system controller info[y/n]:n
Number of Network Interfaces ? (1)
NIC 1 - IP Address ? bang-p
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
```

Successfully defined node bang

5. Add the SGI ProPack administration node to the cluster. For example on the IRIX node rum:

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

```
Is this a FailSafe cluster <true|false> ? (false)  
Is this a CXFS cluster <true|false> ? (true)  
Cluster Notify Cmd [optional] ?  
Cluster Notify Address [optional] ?  
Cluster CXFS mode <normal|experimental>[optional] ? (normal)  
Cluster ID ? (1)
```

Current nodes in cluster performance:

```
Node - 1: ray  
Node - 2: snake  
Node - 3: rum
```

No networks in cluster performance

```
Add nodes to or remove nodes/networks from cluster performance  
Enter "done" when completed or "cancel" to abort
```

```
performance ? add node bang  
performance ? done  
Added node <bang> to cluster <performance>  
Successfully modified cluster performance
```

6. Modify the CXFS filesystems to remove the IRIX administration nodes as metadata servers and add the new SGI ProPack administration node as metadata server. For example, on the IRIX node rum:

```
cmgr> modify cxfs_filesystem V9500  
  
(Enter "cancel" at any time to abort)
```

```
Device ? (/dev/cxvm/V9500)
Mount Point ? (/cxfs/V9500)
Mount Options[optional] ?
Use Forced Unmount ? <true|false> ? (false)
Grio Qualified Bandwidth[optional] ?
Grio managed filesystem ? <true|false>[optional] ?
Default Local Status ? (enabled)
```

MODIFY CXFS FILESYSTEM OPTIONS

- 0) Modify Server.
- 1) Add Server.
- 2) Remove Server.
- 3) Add Enabled Node.
- 4) Remove Enabled Node.
- 5) Add Disabled Node.
- 6) Remove Disabled Node.
- 7) Show Current Information.
- 8) Cancel. (Aborts command)
- 9) Done. (Exits and runs command)

Enter option:2

Current servers:

```
CXFS Server 1 - Rank: 0           Node: rum
CXFS Server 2 - Rank: 1           Node: snake
```

Server Node ? **rum**

- 0) Modify Server.
- 1) Add Server.
- 2) Remove Server.
- 3) Add Enabled Node.
- 4) Remove Enabled Node.
- 5) Add Disabled Node.
- 6) Remove Disabled Node.
- 7) Show Current Information.
- 8) Cancel. (Aborts command)
- 9) Done. (Exits and runs command)

Enter option:2

Current servers:
CXFS Server 1 - Rank: 1 Node: snake

Server Node ? **snake**

- 0) Modify Server.
- 1) Add Server.
- 2) Remove Server.
- 3) Add Enabled Node.
- 4) Remove Enabled Node.
- 5) Add Disabled Node.
- 6) Remove Disabled Node.
- 7) Show Current Information.
- 8) Cancel. (Aborts command)
- 9) Done. (Exits and runs command)

Enter option:1

No current servers

Server Node ? **bang**

Server Rank ? **1**

- 0) Modify Server.
- 1) Add Server.
- 2) Remove Server.
- 3) Add Enabled Node.
- 4) Remove Enabled Node.
- 5) Add Disabled Node.
- 6) Remove Disabled Node.
- 7) Show Current Information.
- 8) Cancel. (Aborts command)
- 9) Done. (Exits and runs command)

Enter option:9

Successfully modified cxfs_filesystem V9500

After you complete this step, the filesystems would show the following information:

```
cmgr> show cxfs_filesystem V9500
```

```
Name: V9500
Device: /dev/cxvm/V9500
Mount Point: /cxfs/V9500
Forced Unmount: false
Global Status: disabled
Default Local Status: enabled
```

```
Server Name: bang
Rank: 1
```

7. Remove the IRIX administration nodes from the cluster. For example, switching to the SGI ProPack node bang:

```
cmgr> modify cluster performance
```

Enter commands, you may enter "done" or "cancel" at any time to exit

```
Is this a FailSafe cluster <true|false> ? (false)
Is this a CXFS cluster <true|false> ? (true)
Cluster Notify Cmd [optional] ?
Cluster Notify Address [optional] ?
Cluster CXFS mode <normal|experimental>[optional] ? (normal)
Cluster ID ? (1)
```

Current nodes in cluster performance:

```
Node - 1: ray
Node - 2: snake
Node - 3: rum
Node - 4: bang
```

Add nodes to or remove nodes/networks from cluster performance

Enter "done" when completed or "cancel" to abort

```
performance ? remove node rum
performance ? remove node snake
performance ? done
Successfully modified cluster performance
```

8. Delete the IRIX administration nodes from the pool. For example, from the SGI ProPack node bang:

```
cmgr> delete node rum
Deleted node (rum).
```

```
cmgr> delete node snake
Deleted node (snake).
```

9. Start CXFS services for all nodes in the cluster. For example, from the SGI ProPack node bang:

```
cmgr> start cx_services for cluster performance
```

CXFS services have been activated in cluster performance

10. Mount the CXFS filesystems. For example, from the SGI ProPack node bang:

```
cmgr> admin cxfs_mount cxfs_filesystem V9500
```

```
cxfs_mount operation successful
```

After completing this procedure, the cluster information is as follows:

```
[root@bang root]# clconf_info
```

```
Event at [2004-02-13 08:44:18]
```

```
Membership since Fri Feb 13 08:44:13 2004
```

Node	NodeID	Status	Age	CellID
ray	3	up	1	0
bang	64	up	1	3

```
1 CXFS FileSystems
```

```
/dev/cxvm/V9500 on /cxfs/V9500 enabled server=(bang) 1 client(s)=(ray) status=UP
```

For more information about using the `cmgr` command to perform this procedure, see the following:

- "Unmount a CXFS Filesystem with `cmgr`" on page 619
- "Stop CXFS Services with `cmgr`" on page 605
- "Define a Node with `cmgr`" on page 577
- "Modify a Cluster with `cmgr`" on page 601
- "Modify a CXFS Filesystem with `cmgr`" on page 619
- "Modify a Cluster with `cmgr`" on page 601
- "Delete a Node with `cmgr`" on page 593
- "Start CXFS Services with `cmgr`" on page 605
- "Mount a CXFS Filesystem with `cmgr`" on page 618

Example of Checking Cluster Status Using `cmgr`

To query node and cluster status, use the following `cmgr` command on a CXFS administration node:

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

Example of Querying Node Status Using `cmgr`

To query node status, use the following `cmgr` command:

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

Example of Pinging the System Controller Using `cmgr`

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

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

Note: This is not required when using `cxfs_admin` because it will attempt to verify that each node is connected as it is added to the cluster.

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.

Example of Monitoring Reset Lines Using `cmgr`

You can use the `cmgr` command to ping the system controller at a node as follows (line break for readability):

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

Note: This is not required when using `cxfs_admin`.

Example of I/O Fencing Status Using `cmgr`

To check the current fencing status, use the `admin fence query` command in `cmgr`

To check current failure action settings, use the `show node nodename` command in `cmgr`.

Example of Using `build_cmgr_script` to Recreate the Cluster Database

Note: Also see "Saving and Recreating the Current Configuration with `cxfs_admin`" on page 300.

You can use the `build_cmgr_script` command from one node in the cluster to create a `cmgr` script that will recreate the node, cluster, switch, and filesystem definitions for all nodes in the cluster database. You can then later run the resulting script to recreate a database with the same contents; this method can be used for missing or corrupted cluster databases.

Note: The `build_cmgr_script` script does not contain local logging information, so it cannot be used as a complete backup/restore tool.

To perform a database backup, use the `build_cmgr_script` script from one node in the cluster, as described in "Creating a `cmgr` Script Automatically" on page 632.



Caution: Do not make configuration changes while you are using the `build_cmgr_script` command.

By default, this creates a `cmgr` script in the following location:

```
/var/cluster/ha/tmp/cmgr_create_cluster_clusname_processID
```

You can specify another filename by using the `-o` option.

To perform a restore on all nodes in the pool, do the following:

1. Stop CXFS services on all nodes in the cluster.
2. Stop the cluster database daemons on each node.
3. Remove all copies of the old database by using the `cdbreinit` command on each node.
4. Execute the `cmgr` script (which was generated by the `build_cmgr_script` script) on the node that is defined first in the script. This will recreate the backed-up database on each node.

Note: If you want to run the generated script on a different node, you must modify the generated script so that the node is the first one listed in the script.

5. Restart cluster database daemons on each node.

For example, to back up the current database, clear the database, and restore the database to all administration nodes, do the following on administration nodes as directed:

On one node:

```
# /var/cluster/cmgr-scripts/build_cmgr_script -o /tmp/newcdb
Building cmgr script for cluster clusterA ...
build_cmgr_script: Generated cmgr script is /tmp/newcdb
```

On one node:

```
cmgr> stop cx_services for cluster clusterA
```

On each node:

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

On each node:

IRIX:

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

SGI ProPack:

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

On each node:

```
# /usr/cluster/bin/cdbreinit
```

On each node:

IRIX:

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

SGI ProPack:

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

*On the *first* node listed in the /tmp/newcdb script:*

```
# /tmp/newcdb
```


Migration from `cmgr` to `cxfs_admin`

If you have scripts that use `cmgr`, you should do the following to use `cxfs_admin`:

1. Do one of the following:
 - Run the `cmgr` script and build the cluster configuration from scratch
 - Start with the cluster database in the desired configuration
2. Run the `cxfs_admin config` command to generate a `cxfs_admin` script. This script should be the `cxfs_admin` equivalent of the `cmgr` script. See "Saving and Recreating the Current Configuration with `cxfs_admin`" on page 300.

Note: The bash shell interprets `cmgr`-generated scripts, but `cxfs_admin` interprets `cxfs_admin`-generated scripts.

Initial Configuration Checklist

Following is a checklist of the steps you must perform when installing and configuring a CXFS system.



Caution: CXFS is a complex product. To ensure that it is installed and configured in an optimal manner, you **must** purchase initial setup services from SGI.

This checklist is not intended to be used directly by the customer, but is provided for reference. It is intended to be used only as an aid; you must be certain to read this entire manual, especially "Configuration Best Practices" on page 139 and Chapter 19, "Troubleshooting" on page 425, before attempting to complete these procedures.

- [] Understand the application use, storage configuration, and I/O patterns at the site. (Not all applications benefit from CXFS; see "Comparison of XFS and CXFS" on page 3.)
- [] Connect the SAN hardware. See the RAID documents.
- [] Is there a private network? This is a requirement.
- [] Is system reset configured for all potential metadata servers, as recommended by SGI? Is system reset or I/O fencing configured for all client-only nodes? One of these solutions is required to ensure data integrity for **all** nodes. See "Isolating Failed Nodes: Failure Policies" on page 28, and "Reset" on page 34.
- [] Are there an odd number of server-capable nodes with CXFS services running? Is there a client-only tiebreaker node? See "Use an Odd Number of Server-Capable Nodes" on page 145 and "Use a Client-Only Tiebreaker" on page 145.
- [] Read this book and any README files and release notes provided with the release. If you have clients running operating systems other than IRIX or SGI ProPack, also read the *CXFS MultiOS Client-Only Guide for SGI InfiniteStorage*.
- [] Verify that the network is usable. See Chapter 6, "IRIX CXFS Installation" on page 101.

- [] Install the CXFS software. See Chapter 6, "IRIX CXFS Installation" on page 101, and Chapter 7, "SGI ProPack CXFS Installation" on page 113. If you have clients running operating systems other than IRIX or SGI ProPack, also see the *CXFS MultiOS Client-Only Guide for SGI InfiniteStorage*.
- [] Modify the configuration files and perform other tasks as needed. See Chapter 8, "Postinstallation Steps" on page 127.
- [] Completely configure and run a **small** cluster (3 nodes). See Chapter 10, "Initial Setup of the Cluster" on page 163.
- [] Look for errors in the daemon log files in the `/var/cluster/ha/logs` directory.
- [] If all is well, add the rest of the nodes. If there are problems, see Chapter 19, "Troubleshooting" on page 425.
- [] Set up the filesystems. See Chapter 10, "Initial Setup of the Cluster" on page 163.

Summary of New Features from Previous Releases

This appendix contains a summary of the new features for each version of this guide.

CXFS Version 1: Original Implementation

CXFS version 1 is the original implementation of CXFS.

IRIX 6.5.6f

Original publication (007-4016-001).

IRIX 6.5.6f

The 007-4016-002 update contains additional troubleshooting information and instructions for unmounting and remounting filesystems with the command line interface. It was reorganized to make the tasks of installation and configuration clearer.

IRIX 6.5.7f

The 007-4016-003 update contains the following:

- Metadata server recovery information
- Administrative shutdown procedures
- Additional troubleshooting information
- Instructions for unmounting and remounting filesystems with the CLI
- Reorganized installation and configuration information

IRIX 6.5.8f

The 007-4016-004 update contains the following:

- Support for hierarchical storage management (HSM) through data management application programming interface (DMAPI), also known as X/Open data storage management specification (XDSM)
- Changes to administrative shutdown, including two new `cmgr` subcommands to stop CXFS services on the local nodes: `admin cxfs_stop` and `admin cxfs_stop`
- Quorum changes without panics

IRIX 6.5.9f

The 007-4016-005 update contains the following:

- Coexecution of CXFS and IRIS FailSafe 2.1, including commands to convert nodes and clusters to apply to both utilities
- Ability to use the `cmgr` command without extra prompting (`-p`), permitting the use of scripts
- New tasks to revoke and allow membership of the local node
- Ability to specify the tie-breaker node, which is used in the process of computing node membership for the cluster when exactly half the nodes in the cluster are up and can communicate with each other
- Clarification that a single subnet should be used

IRIX 6.5.10f

The 007-4016-006 update contains the following:

- Clarifications about CXFS shutdown and database shutdown
- Additional information about CXFS daemons
- Clarifications to the comparison of XFS and CXFS

IRIX 6.5.11f

The 007-4016-007 update contains the following:

- Addition of the Origin 3000 partition ID to node configuration
- Troubleshooting information for a two-node cluster when both nodes go down
- Information about editing the `/etc/hosts` file to use an alternate interface for heartbeat and control.
- Clarification about the use of hardware reset and tie-breaker nodes
- Ability to unset the tie-breaker node
- Use of `fsr`

CXFS Version 2: MultiOS Cluster

CXFS version 2 includes client-only nodes on operating system platforms other than IRIX (*multiOS cluster*, or *heterogeneous clients*).

IRIX 6.5.12f

The 007-4016-008 update contains the following:

- A cluster of at least **three** weighted nodes is recommended for a production environment (that is, one requiring relocation and recovery of the metadata server).

If you use a two-weighted-node cluster for production, you must do one of the following:

- Use reset lines to avoid data corruption and ensure that only one node is running in error conditions (reset lines are recommended for all CXFS clusters and required for use with IRIS FailSafe).
- Weight one node as 1 and the other as 0.
- Set a tie-breaker node.

However, there are issues with a two-weighted-node cluster.

- The new `cluster_status` command, which provides a `curses` interface to display status information gathered by the `cad` daemon (this information is also displayed by the `cxdetail` command).
- Cluster nodes can run adjacent levels of the IRIX operating system (OS); for example, 6.5.11f and 6.5.12f (this applies as of 6.5.12f).
- The ability to execute your own custom scripts around mounting operations.
- Partition ID information.
- Clarification about the following:
 - Hostname resolution rules; it is **critical** that you understand these rules and have files configured properly before attempting to configure a cluster.
 - The difference between *CXFS membership* and *fs2d membership*.
 - Configuration of nodes supported in an IRIS FailSafe and CXFS coexecution environment.
 - Unwritten extent tracking (`unwritten=1|0`) with CXFS.
 - Including the CXFS mount point in the `/etc/exports` file.
 - Number of nodes supported: 16 CXFS nodes, and up to 8 IRIS FailSafe nodes with coexecution.
 - The flow of information in a coexecution cluster.

IRIX 6.5.13f

The 007-4016-009 update contains the following:

- The structure of the CXFS filesystem configuration has changed. CXFS filesystems can now be defined, modified, managed and deleted independently of each other and of the cluster definition. (Previously, the CXFS filesystems were defined as attributes to the cluster definition.)

The new design improves the performance, flexibility and reliability of filesystem configuration. To accommodate clusters mixing nodes running 6.5.12 and 6.5.13, backwards compatibility is enforced by default in 6.5.13. The result is that the performance achievements are not visible; however, if you are 6.5.13 on all nodes in the cluster, you may wish to turn off backwards compatibility. Backwards compatibility will be turned off in the 6.5.14 release.

- Information about locating the `xfsdump` inventory in a shared directory.
- Information about the IRIS FailSafe CXFS resource type that can be used to failover applications that use CXFS filesystems.
- The `-p` option is no longer required when defining filesystems with the `cmgr` command; the scripting capability is therefore provided.
- For certain GUI tasks, the ability to select all nodes at once in addition to specifying nodes individually.
- New `/var/cluster/cmgr-scripts/rotatelog`s script to save log files with day and month name as suffixes.
- The setting to force an unmount of a filesystem using the `umount -k` option is turned off by default in the GUI. There is no default when using the `cmgr` command.
- Clarification of the term *CLI* to mean the underlying set of commands that are used by the `cmgr` cluster manager tool and by the GUI.
- Use of `sgi_apache.sw.server`.
- Correction: real-time filesystems are not currently supported. Changes to reflect this have been made in text.
- New and revised figures.

IRIX 6.5.14f

The 007-4016-011 update contains the following:

- The graphical user interface (GUI) has been improved. The separate cluster view (the `cxdetail` command) and task manager (the `cxtask` command) have been streamlined into one window, the **CXFS Manager**. Both the `cxtask` and `cxdetail` commands are kept for historical purposes; this document refers to just `cxtask` for simplicity.

The new GUI provides the following features:

- Access to tasks through the menu bar or by clicking the right mouse button within the tree view
- Faster filesystem status and cluster status updates

- Access to the `salog(4)` file, which shows every command run from the GUI
- A **Find** textfield helps you find components within the displayed tree-view
- Information about the use of `xfs_repair` and CXFS filesystems.



Caution: Do not use `xfs_repair` on a CXFS filesystem unless you are certain there is a problem.

- Information about using `cmgr(1M)`:
 - Invoking subcommands directly on the command line with the `-c` option
 - Using template scripts provided in the `/var/cluster/cmgr-templates` directory
- Information about MAC labels in a mixed Trusted IRIX and IRIX cluster.
- The structure of the CXFS filesystem configuration was changed with the release of IRIX 6.5.13. Backward compatibility with earlier versions is no longer maintained as of IRIX 6.5.14, because all nodes in the cluster must be running the same or adjacent releases.

If you are upgrading from 6.5.13f entirely to 6.5.14f, there is no further impact.

If you intend to run a mixture of 6.5.13f and 6.5.14f nodes, you must turn off backward compatibility.

If you are upgrading from 6.5.12f or earlier without first installing and running 6.5.13f, then you must perform a one-time manual conversion of your CXFS filesystem definitions.

IRIX 6.5.15f

The 007-4016-012 update contains the following:

Note: Relocation and recovery are deferred in this release.

- Support for clients of other operating systems such as Solaris and Windows NT as defined in the *CXFS MultiOS Client-Only Guide for SGI InfiniteStorage*. These clients will be released asynchronously from the IRIX release. This support will require a

minimum of IRIX 6.5.15f plus appropriate patches. For more information, see your SGI support contact.

- Default scripts are now provided in the `/var/cluster/clconfd-scripts` directory to permit NFS-exporting of CXFS filesystems listed in `/etc/exports`.
- Reset lines are mandatory for two-node and two-weighted node clusters. Larger clusters should have an odd number of weighted nodes, or must have serial reset lines if only two of the nodes are weighted.
- Simplification of Chapter 1. General information about the CXFS Manager GUI and `cmgr` have been moved to their respective reference chapters, coexecution details have been moved into a separate chapter, and the communication flow diagrams and daemon information have been moved into an appendix.
- Information about the error messages that may cause administrators to use `xfstool repair` inappropriately.
- Changes to the `rotatelog` script syntax. The root `crontab` file now has an entry to run the `rotatelog` script weekly. If you run the script twice in one day, it will append the current log file to the previous saved copy, rather than overwriting it.
- A new figure describing some of the various combinations of node and cluster types in a coexecution cluster.

IRIX 6.5.16f

The 007-4016-013 update contains the following:

Note: Relocation and recovery are fully implemented, but the number of associated problems prevents support of these features in CXFS. While data integrity is not compromised, cluster node panics or hangs are likely to occur. These features will be fully supported when these issues are resolved.

- Support for Solaris and Windows NT systems in a multiple operating system (multiOS) cluster, including the following:
 - Information about defining the operating system for a node. For existing clusters that are upgraded to IRIX 6.5.16f, existing nodes will be assigned an operating system type of `IRIX`.

- Information about I/O fencing, which allows a problem node to be isolated from the storage area network (SAN) so that it cannot corrupt data in the shared CXFS filesystem. Solaris and Windows NT nodes require a Brocade switch in order to support I/O fencing for data integrity protection; therefore, the Brocade switch is a required piece of hardware in a cluster running multiple operating systems.
- The new terms *multiOS* and *CXFS client-only node*.
- Support for the L1 controller on SGI Origin 300, SGI Origin 3200C, SGI Onyx 300, and SGI Onyx 3200C systems.
- Information about the CXFS GUI tasks to define and modify a filesystem, which have been split into two pages for ease of use.
- New GUI icons.

IRIX 6.5.17f

The 007-4016-014 update contains the following:

- A new appendix contains an example `/etc/ipfilterd.conf` file that can be used to provide IP filtering for the CXFS private network.
- The `build_cmgr_script` command, which generates a `cmgr` script from the cluster database. The script can be used later to recreate the cluster database after performing a `cdbreinit` command.
- A sample script to `unexport` and locally `umount` an `lofs` filesystem.
- Use of the new command name `cx fsmgr`. The `cx fsmgr` command has the same function as the `cxtask` and `cxdetail` commands, which are kept for historical purposes.
- Clarifications to the following:
 - Starting the CXFS Manager graphical user interface
 - Masking and I/O fencing
 - Terminology such as *cluster*, *node*, and *pool*
 - Terminology used to describe the GUI

IRIX 6.5.18f

The 007-4016-015 update contains the following:

Note: In this release, relocation is disabled by default and recovery is supported only when using standby nodes.

A *standby node* is a metadata server-capable administration node that is configured as a potential metadata server for a given filesystem, but does not currently run any applications that will use that filesystem. To use recovery, you must not run any applications on any of the potential metadata servers for a given filesystem; after the active metadata server has been chosen by the system, you can then run applications that use the filesystem on the active metadata server and client-only nodes.

Relocation and recovery are fully implemented, but the number of associated problems prevents full support of these features in the current release. Although data integrity is not compromised, cluster node panics or hangs are likely to occur. Relocation and recovery will be fully supported in a future release when these issues are resolved.

- IRIX nodes may now be CXFS client-only nodes, meaning that they run a minimal implementation of the CXFS and cluster services, and do not contain a copy of the CXFS cluster database. Client-only nodes are installed with the `cxfs_client` software product.

This change also introduces the term *CXFS administration node*, which is a node that is installed with the `cluster_admin` software product, allowing the node to perform cluster administration tasks and contain a copy of the cluster database. Nodes that you want to run as metadata servers must be installed as CXFS server-capable administration nodes; SGI recommends that all other nodes be installed as client-only nodes.

When you define a node, you no longer need to specify the node weight. This has been replaced by the **Node Function** field, allowing you to choose **Server-capable Admin**, **Client Admin**, or **Client-Only**. (For Solaris and Windows nodes, **Client-Only** is automatically selected for you.) Similar fields are provided for the `cmgr` command.

When upgrading to 6.5.18f, already existing IRIX nodes will by default be assigned as **Server-capable Admin** if they had a weight of 1.

This version also clarifies the terms used for membership: *CXFS kernel membership* and *cluster database membership*.

- New system-tunable parameters:
 - `cxfs_relocation_ok` lets you enable or disable the relocation feature; relocation is disabled by default in this release, and SGI recommends that you **do not** enable it.
 - `cxfsd_min` and `cxfsd_max` let you specify the minimum and maximum number of `cxfsd` threads to run per CXFS filesystem.
- New commands:
 - `cxfs_info` provides status information about the cluster, nodes, and filesystems and is run from a client-only node.
 - `cxfsdump` gathers CXFS configuration information.
- A CXFS cluster is supported with as many as 32 nodes. As many as 16 of those nodes can be CXFS administration nodes and all other nodes can be client-only nodes. You can choose to define a node as a CXFS client administration node, however, SGI strongly recommends that only potential metadata servers be configured as CXFS server-capable administration nodes and that there be an odd number of server-capable nodes for quorum calculation purposes.
- The graphical user interfaces for XVM and CXFS have been combined into one. This guide provides an overview of the XVM-specific tasks provided by the GUI; for details about these tasks, see the *XVM Volume Manager Administrator's Guide*.

The tasks to make, grow, mount/unmount a filesystem are now provided in the GUI.
- Tips about using CXFS and Trusted IRIX.
- Support for Microsoft Windows 2000 systems as client-only nodes. (This guide uses *Windows* to refer to both Microsoft Windows NT and Microsoft Windows 2000 nodes when the information applies equally to both. Information that applies to only one of these types of nodes is identified.)

IRIX 6.5.19f

The 007-4016-016 update contains the following:

- The new rolling annual upgrade policy that permits you to upgrade from 6.5.*n* to the *n*+1 or *n*+4 release, as of 6.5.18f.

- The time required to update and propagate the database across nodes in the cluster has been significantly decreased.
- If you use I/O fencing and `ipfilterd` on a node, the `ipfilterd` configuration must allow communication between the node and the `telnet(1)` port on the switch.
- The following nodes do not contain system controllers and therefore require I/O fencing for data integrity protection:
 - Silicon Graphics Fuel visual workstation
 - Silicon Graphics Octane system
 - Silicon Graphics Octane2 system
- The CXFS Manager graphical user interface (GUI) has added a new icon to represent client-only nodes.
- In preparation for future CXFS MultiOS client releases, the CXFS software now also allows you to specify the Linux, IBM AIX, and Hewlett-Packard HP-UX operating systems when defining a node. For support details, see the *CXFS MultiOS Client-Only Guide for SGI InfiniteStorage* and release notes.
- This version clarifies the various methods to perform cluster database backups and restorations.
- Application programmers should be aware that XFS recently relaxed the requirement that direct I/O be aligned along filesystem block boundaries. As of IRIX 6.5.19f, direct I/O will also be accepted using 512-byte alignment.

This change makes the use of direct I/O on a CXFS partition more consistent with that of other vendor's requirements and thus makes the use of CXFS more transparent. See the description of direct I/O requirements in the `fcntl` man page.

- This version lists the system tunable parameters found in the `/var/sysgen/mtune/cell` file, some of which should not be modified.

IRIX 6.5.20f

The 007–4016–017 update contains the following:

- Changes to the CXFS graphical user interface (GUI):
 - New login connection choices, including support for a remote shell connection, which connects to the server via a user-specified command shell, such as `rsh` or `ssh`.
 - The ability for the `root` user to grant other users permission to execute specific GUI tasks.
 - Use of Java2 for the CXFS GUI, which simplifies the Java installation and co-operation with third-party GUIs. This also enhances the ability to run the GUI through a web browser (via `http://server/CXFSManager/`).
 - Information about using the right mouse button to access tasks appropriate to the selected GUI item.
- Changes to the `cxfsd_min` and `cxfsd_max` defaults, and the `cxfsd_max` legal values.
- More information about memberships, quorums, and tiebreakers.
- A new figure describing standby mode.
- More information about IRIX client-only issues:
 - Client-only node system files
 - Status in log files
 - `cxfs_client` error messages

CXFS Version 3: IRIX or SGI ProPack (Linux 2.4 Kernel) Servers

CXFS version 3 adds support for CXFS metadata servers on SGI Altix systems running Linux (2.4 kernel).

CXFS 3.0

The 007-4016-018 update contains the following:

- Support for SGI ProPack metadata servers on SGI Altix 3000 family of servers and superclusters. A CXFS cluster can contain either SGI ProPack 2.3 server-capable nodes on Altix systems or IRIX server-capable nodes; you cannot mix IRIX and SGI ProPack server-capable nodes within one cluster.

CXFS does not support the relocation or recovery of DMAPi filesystems that are being served by SGI ProPack metadata servers.

Coexecution with FailSafe is not supported on SGI ProPack nodes.

- Due to packaging enhancements, CXFS may now be installed on the M stream or the F stream.

The IRIX CXFS software will no longer be bundled in the IRIX overlay CDs but instead is on a separate *CXFS IRIX Server and Client 3.0 for IRIX 6.5.22* CD. This changes the installation procedure.

Note: If you are upgrading from a previous IRIX release and have CXFS installed, you must upgrade both IRIX and CXFS. If you try to upgrade one without the other, conflicts will occur.

- Information about defining networks for CXFS kernel messaging (in addition to the network used for heartbeat/control). However, use of these networks is **deferred**.
- Support for IRIX real-time filesystems.
- Suggestions for configuring large clusters.
- Information about using `ping` to verify general connectivity and CXFS heartbeat in a multicast environment.
- The GUI has been changed to show a single display for the nodes in the cluster and nodes that are in the pool but not in the cluster. This new selection is **View: Nodes and Cluster**.
- Information about information retaining system core files and the output from the `cxfsdump` utility when reporting problems.
- Information about monitoring heartbeat timeouts for IRIX using Performance Co-Pilot or the `icrash` command.

- The ability to define multiple CXFS filesystems at one time with the GUI.

CXFS 3.1

The 007-4016-019 update contains the following:

- Information about migrating from an IRIX cluster to a SGI ProPack cluster
- Support for a cluster of up to 64 nodes.
- Information about the TP9300 RAID.
- Information about the `cxfs-config` command.
- Clarification that serial hardware reset lines or I/O fencing is **required for all nodes** in order to protect data integrity.
- The ability to define a reset method for a given node to one of the following:
 - `powerCycle` to turn power off and on
 - `reset` to perform a serial reset
 - `nmi` to perform a nonmaskable interrupt

You can define this method using either the `cmgr` command or the GUI. You can manually perform a powercycle or an NMI with the `cmgr` command.

- New appendixes summarizing operating system path differences and new features from previous releases

CXFS 3.2

The 007-4016-021 update contains the following:

- Information about private network failover as defined with the `cmgr` command. (Although the primary network must be private, the backup network may be public.)
- Support for Guaranteed-rate I/O version 2 (GRIOv2) in the IRIX installation procedure.
- Corrections to CXFS and cluster administration path differences between IRIX and SGI ProPack on SGI Altix systems.

- Updated the example for `clconf_info` command. The `clconf_info` command now reports a node as `inactive` rather than `DOWN*` and the unused incarnation number has been removed.
- Support for token obtain optimization. To disable, use the `cxfs_prefetch` system tunable parameter.
- If you have a cluster with an even number of server-capable nodes and no tiebreaker: to avoid a split-brain scenario, you should not use the **Shutdown** setting on any server-capable node.
- Information about multiple Ethernet interfaces on SGI Altix systems and providing persistent device naming.
- Clarification about the `chkconfig` arguments used for IRIX administration nodes, SGI ProPack administration nodes, and client-only nodes.
- Information about the correct options to use for quotas on SGI ProPack clusters (`uquota` and `gquota`).
- Information about serial reset configurations.
- Information about using the `hafence(1M)` command to define a QLogic switch. (You cannot use the GUI or the `cmgr` command to define or modify a switch other than a Brocade switch.)
- If you want to use quotas on a CXFS filesystem, you must install the quota package.
- Information about removing a metadata server from the cluster for maintenance.
- Mount options information.
- Addition of the XVM graphical user interface (GUI) to the CXFS SGI ProPack package.

CXFS 3.3

The 007-4016-022 update contains the following:

- Procedures to remove a single client from the cluster and restore it, and shut down the entire cluster.
- A new chapter about best practices.

- Information about XVM failover.
- Updates to the rolling upgrade policy.
- Information about performing a miniroot install.
- Information about installing the latest Java2 software for use with the CXFS GUI and SGI ProPack.
- Information about modifying the `httpd.conf` file in order to use the CXFS GUI on SGI ProPack.
- Clarifications to terminology.

CXFS 3.4

The 007-4016-023 update contains the following:

- Information about discovering the WWNs
- Support for system reset for SGI Altix systems that use an integrated L2, such as a NUMAlink 4 R-brick, or SGI Altix 3000 Bx2 systems. Configuring a node of this type requires use of `cmgr`.
- Support for SGI ProPack for Linux 4 as a client-only node.
- Support for the `cxfsdump -secure` option for secure remote connections. In cluster mode (the default), `cxfsdump` requires `rsh/ssh` and `rcp/scp` access across all nodes in the cluster.

CXFS Version 4: IRIX or SGI ProPack (Linux 2.6 Kernel) Servers

CXFS version 4 adds support for CXFS metadata servers on SGI Altix systems running Linux 2.6 kernel.

CXFS 4.0

The 007-4016-024 update contains the following:

- Support for SGI ProPack 4 for Linux Service Pack 3.
- Support for IRIX 6.5.29.

- Server-side CXFS client license keys are now supported on server-capable nodes, allowing a client without a node-locked client-side license key to request a license key from the server. Server-side license keys are optional on IRIX metadata servers, but are required on SGI ProPack metadata servers. The licensing software is based on the FLEXlm product from Macrovision Corporation.
- Support for the `cxfs_admin` cluster configuration and administration command, which waits for a command to be completed before continuing, provides the ability to limit which nodes can mount a filesystem, provides an improved interface over `cmgr` (including `<TAB>` completion of commands), and scripting capabilities. In a future release, `cxfs_admin` will replace `cmgr`.
- The availability of dynamic heartbeat monitoring.
- Information about choosing the correct version of XVM failover for your cluster.
- Support for an SGI ProPack node as a server for the Guaranteed Rate I/O version 2 (GRIOv2) feature of CXFS.
- Support for GPT labels.
- Information about using the `md` driver on large SGI Altix systems with CXFS.
- Updates to IRIX CXFS installation procedures.
- Support for specifying a direct list of port numbers to be masked (that is, ports that will never be fenced), rather than using a hexadecimal string that represents the list. Ports that can be masked are now in the range 0 through 1023.
- Ability to specify a vendor when defining a Fibre Channel switch using the GUI.
- Requirement to start the file alteration monitoring (`fam`) service on SGI ProPack nodes in order to use the CXFS graphical user interface (GUI).
- Information about switching between `SGIRDAC` and `SGIAVT` mode for SGI RAID.
- Information about CXFS port usage. CXFS Port Usage.
- Information about the recovery timeout mechanism, in which nodes are polled for progress after a recovery has begun. If recovery for a node is not progressing according to the specified polls, the recovery is considered stalled and the node will then shut down or panic; this prevents the cluster from hanging and keeps filesystems available for the rest of the nodes in the cluster. The following site-changeable system tunable parameters are used:

```
cxfs_recovery_timeout_panic
```

```
cxfs_recovery_timeout_period
cxfs_recovery_timeout_stalled
cxfs_recovery_timeout_start
```

- Information about using the `filestreams` mount option to optimize disk layout.
- Best-practices information about:
 - Using proper storage management procedures
 - Being aware of differences between IRIX and Linux system administration
 - Modifying `updatedb` to avoid unnecessary load
 - Using fast copying for large CXFS files
 - Minimizing the number of switches

Note: The contents of the former “Avoid Problems” section from the Troubleshooting chapter was moved into the Best Practices chapter.

- Addition of SGI RAID information (this information was removed from the release notes).
- Addition of switches information (this information was removed from the release notes).
- A complete list of system-tunable parameters.
- Support for drag-and-drop within the GUI to move nodes between the pool and the cluster.
- Client administration nodes are only supported and appropriate for IRIX nodes running in coexecution with FailSafe. If you are running an earlier release and have an SGI ProPack node that is defined as a client administration node, you must delete the node from the cluster database, uninstall CXFS, reinstall CXFS, and redefine the node as a client-only node. For simplicity, this guide no longer refers to *client-administration node* or *administration node* except in specific instances.
- For clarity, this book now uses the term *SGI ProPack* rather than *Linux* when referring to CXFS nodes running SGI ProPack for Linux.

CXFS 4.1

The 007-4016-025 version includes the following:

- Support for SGI IRIX 6.5.30.
- Support for SGI ProPack 5 running SUSE Linux Enterprise Server 10 (SLES 10).
- Support for SGI License Key (LK) software on SGI ProPack server-capable nodes.

Server-side licensing is required on the following client-only nodes (to determine the Linux architecture type, use the `uname -i` command):

- SGI ProPack 5
- Red Hat Enterprise Linux (RHEL) 4 on x86_64
- SLES 9 on x86_64
- SLES 10 on x86_64 or ia64

Note: For specific release levels, see the release notes.

Other client-only nodes can use either server-side or client-side licensing. However, if one node within a cluster requires server-side licensing, all nodes must use server-side licensing. If no nodes in the cluster require server-side licensing, the nodes can continue to use existing client-side licensing.

Note: Server-side licensing is preferred, and no new client-side licenses will be issued. Customers with support contracts can exchange their existing client-side licenses for new server-side licenses. A future release will not support client-side licensing. For more information, contact SGI customer support.

- Support for the following RAID hardware:
 - SGI InfiniteStorage 10000
 - SGI InfiniteStorage 6700
 - SGI InfiniteStorage 4500
 - SGI InfiniteStorage 4000
- DMAPI is disabled by default on SGI ProPack 5 systems. When you install DMF on a server-capable node, it automatically enables DMAPI. However, if you want

to mount filesystems on an SGI ProPack 5 client-only node with the `dmi` mount option, you must enable DMAPL.

- Information about the appropriate use of `lcrash` for SGI ProPack.
- Information about the procedure required when upgrading from CXFS 3.4.1 or earlier for clusters with three or more server capable nodes.
- Information about removing a metadata server from the cluster while avoiding a reset.
- New `mtcp_rpc_thread` system tunable parameter (for SGI ProPack and Linux third-party nodes) that specifies whether metadata messages are sent from a separate thread in order to save stack space.
- Information about rotating log files on SGI ProPack.
- Information about `SYSLOG credid` warnings.
- The table of mount options is now located in the *CXFS MultiOS Client-Only Guide for SGI InfiniteStorage*.

Glossary

active metadata server

A server-capable administration node chosen from the list of potential metadata servers. There can be only one active metadata server for any one filesystem.

administration node

A node in the pool that is installed with the `cluster_admin.sw.base` software product, allowing the node to perform cluster administration tasks and contain a copy of the cluster database. There are two types of administration nodes: *server-capable administration nodes* and *client administration nodes*.

administrative stop

See *forced CXFS shutdown*

ARP

Address resolution protocol

BMC

Baseboard management controller

cell ID

A number associated with a node that is allocated when a node is added into the cluster definition with the GUI or `cxfs_admin`. The first node in the cluster has cell ID of 0, and each subsequent node added gets the next available (incremental) cell ID. If a node is removed from the cluster definition, its cell ID becomes available. It is not the same thing as the *node ID*.

CLI

Underlying command line interface commands used by the CXFS Manager graphical user interface (GUI).

client

See *CXFS client node*, *CXFS client-only node* and *administration node*.

client administration node

A node that is installed with the `cluster_admin` software product, allowing the node to perform cluster administration tasks and contain a copy of the cluster database, but is not capable of coordinating CXFS metadata. Only supported for IRIX nodes running in coexecution with FailSafe.

client-only node

A node that is installed with the `cxfs_client.sw.base` software product; it does not run cluster administration daemons and is not capable of coordinating CXFS metadata. Any node can be client-only node. See also *server-capable administration node*

cluster

A *cluster* is the set of systems (nodes) configured to work together as a single computing resource. A cluster is identified by a simple name and a cluster ID. A cluster running multiple operating systems is known as a *multiOS cluster*.

There is only one cluster that may be formed from a given pool of nodes.

Disks or logical units (LUNs) are assigned to clusters by recording the name of the cluster on the disk (or LUN). Thus, if any disk is accessible (via a Fibre Channel connection) from machines in multiple clusters, then those clusters must have unique names. When members of a cluster send messages to each other, they identify their cluster via the cluster ID. Cluster names must be unique.

Because of the above restrictions on cluster names and cluster IDs, and because cluster names and cluster IDs cannot be changed once the cluster is created (without deleting the cluster and recreating it), SGI advises that you choose unique names and cluster IDs for each of the clusters within your organization.

cluster administration daemons

The set of daemons on a server-capable administration node that provide the cluster infrastructure: `fs2d`, `cad`, `cmond`, `crsd`. See "CXFS Control Daemon" on page 26.

cluster administrator

The person responsible for managing and maintaining a cluster.

cluster database

Contains configuration information about all nodes and the cluster. The database is managed by the cluster administration daemons.

cluster domain

XVM concept in which a filesystem applies to the entire cluster, not just to the local node. See also *local domain*.

cluster database membership

The group of administration nodes in the **pool** that are accessible to cluster administration daemons and therefore are able to receive cluster database updates; this may be a subset of the nodes defined in the pool. The cluster administration daemons manage the distribution of the cluster database (CDB) across the administration nodes in the pool. (Also known as *user-space membership* and *fs2d database membership*.)

cluster ID

A unique number within your network in the range 1 through 255. The cluster ID is used by the operating system 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 must not change a cluster ID after the cluster has been defined. Clusters IDs must be unique.

cluster mode

One of two methods of CXFS cluster operation, `Normal` or `Experimental`. In `Normal` mode, CXFS resets any node for which it detects heartbeat failure; in `Experimental` mode, CXFS ignores heartbeat failure. `Experimental` mode allows you to use the kernel debugger (which stops heartbeat) without causing node failures. You should only use `Experimental` mode during debugging.

control messages

Messages that cluster software sends between the cluster nodes to request operations on or distribute information about cluster nodes. Control messages and heartbeat messages are sent through a node's network interfaces that have been attached to a control network.

cluster node

A node that is defined as part of the cluster. See also *node*.

coexecution

The ability to run CXFS and IRIS FailSafe together. For more information, see "Overview of FailSafe Coexecution" on page 41.

control network

The network that connects nodes through their network interfaces (typically Ethernet) such that CXFS can send heartbeat messages and control messages through the network to the attached nodes. CXFS uses the highest priority network interface on the control network; it uses a network interface with lower priority when all higher-priority network interfaces on the control network fail.

CXFS client daemon

The daemon (`cxfs_client`) that controls CXFS services on a client-only node. See "CXFS Client Daemon" on page 26.

CXFS control daemon

The daemon (`clconfd`) that controls CXFS services on an administration node. See "Cluster Administration Daemons" on page 24.

CXFS database

See *cluster database*.

CXFS kernel membership

The group of CXFS nodes that can share filesystems in the cluster, which may be a subset of the nodes defined in a cluster. During the boot process, a node applies for CXFS kernel membership. Once accepted, the node can share the filesystems of the cluster. (Also known as *kernel-space membership*.) CXFS kernel membership differs from *cluster database membership* and FailSafe membership. For more information about FailSafe, see *FailSafe Administrator's Guide for SGI InfiniteStorage*.

CXFS services

The enabling/disabling of a node, which changes a flag in the cluster database. This disabling/enabling does not affect the daemons involved. The daemons that control

CXFS services are `clconfd` on an administration node and `cxfs_client` on a client-only node. See "CXFS Services" on page 25.

CXFS services start

To enable a node, which changes a flag in the cluster database, by using an administrative task in the CXFS GUI or the `cxfs_admin enable` command.

CXFS services stop

To disable a node, which changes a flag in the cluster database, by using the CXFS GUI or the `cxfs_admin disable` command. See also *forced CXFS shutdown*.

CXFS shutdown

See *forced CXFS shutdown* and *shutdown*

CXFS tiebreaker node

A node identified as a tiebreaker for CXFS to use in the process of computing CXFS kernel membership for the cluster, when exactly half the nodes in the cluster are up and can communicate with each other. There is no default CXFS tiebreaker. SGI recommends that the tiebreaker node be a client-only node. The CXFS tiebreaker differs from the FailSafe tiebreaker; see *FailSafe Administrator's Guide for SGI InfiniteStorage*.

database

See *cluster database*.

database membership

See *cluster database membership*.

details area

The portion of the GUI window that displays details about a selected component in the view area. See also *view area*.

domain

See *cluster domain* and *local domain*.

dynamic heartbeat monitoring

Starts monitoring only when an operation is pending. Once monitoring initiates, it monitors at 1-second intervals and declares a timeout after 5 consecutive missed seconds, just like *static heartbeat monitoring*.

FailSafe Membership

The group of nodes that are actively sharing resources in the cluster, which may be a subset of the nodes defined in a cluster. FailSafe membership differs from *CXFS kernel membership* and *cluster database membership*. For more information about FailSafe, see *FailSafe Administrator's Guide for SGI InfiniteStorage*.

failure action hierarchy

See *failpolicy methods*

failpolicy methods

The set of instructions that determine what happens to a failed node; the second instruction will be followed only if the first instruction fails; the third instruction will be followed only if the first and second fail. The available actions are: *fence*, *fencereset*, *reset*, and *shutdown*. Also known as *failure action hierarchy*

fence

The failure policy method that isolates a problem node so that it cannot access I/O devices, and therefore cannot corrupt data in the shared CXFS filesystem. I/O fencing can be applied to any node in the cluster (CXFS clients and metadata servers). The rest of the cluster can begin immediate recovery.

fencereset

The failure policy method that fences the node and then, if the node is successfully fenced, performs an asynchronous system reset; recovery begins without waiting for reset acknowledgment. If used, this fail policy method should be specified first. If the fencing action fails, the reset is not performed; therefore, *reset* alone is also highly recommended for all server-capable nodes (unless there is a single server-capable node in the cluster).

fencing recovery

The process of recovery from fencing, in which the affected node automatically withdraws from the CXFS kernel membership, unmounts all file systems that are using an I/O path via fenced HBA(s), and then rejoins the cluster.

forced CXFS shutdown

The withdrawal of a node from the CXFS kernel membership, either due to the fact that the node has failed somehow or by issuing an `admin cxfs_stop` command. This disables filesystem and cluster volume access for the node. The node remains enabled in the cluster database. See also *CXFS services stop* and *shutdown*.

fs2d database membership

See *cluster database membership*.

gratuitous ARP

ARP that broadcasts the MAC address to IP address mappings on a specified interface

heartbeat messages

Messages that cluster software sends between the nodes that indicate a node is up and running. Heartbeat messages and *control messages* are sent through the node's network interfaces that have been attached to a control network.

heartbeat interval

The time between heartbeat messages. The node timeout value must be at least 10 times the heartbeat interval for proper CXFS operation. The higher the number of heartbeats (smaller heartbeat interval), the greater the potential for slowing down the network. See also *dynamic heartbeat monitoring* and *static heartbeat monitoring*.

I/O fencing

See *fence*.

IPMI

Intelligent Platform Management Interface

kernel-space membership

See *CXFS kernel membership*.

local domain

XVM concept in which a filesystem applies only to the local node, not to the cluster. See also *cluster domain*.

log configuration

A log configuration has two parts: a *log level* and a *log file*, both associated with a *log group*. The cluster administrator can customize the location and amount of log output, and can specify a log configuration for all nodes or for only one node. For example, the *crsd* log group can be configured to log detailed level-10 messages to the *crsd-foo* log only on the node *foo* and to write only minimal level-1 messages to the *crsd* log on all other nodes.

log file

A file containing notifications for a particular *log group*. A log file is part of the *log configuration* for a log group.

log group

A set of one or more CXFS processes that use the same log configuration. A log group usually corresponds to one daemon, such as *gcd*.

log level

A number controlling the number of log messages that CXFS will write into an associated log group's log file. A log level is part of the log configuration for a log group.

membership

See *cluster database membership* and *CXFS kernel membership*.

membership version

A number associated with a node's cell ID that indicates the number of times the CXFS kernel membership has changed since a node joined the membership.

metadata

Information that describes a file, such as the file's name, size, location, and permissions.

metadata server

The administration node that coordinates updating of metadata on behalf of all nodes in a cluster. There can be multiple potential metadata servers, but only one is chosen to be the active metadata server for any one filesystem.

metadata server recovery

The process by which the metadata server moves from one node to another due to an interruption in CXFS services on the first node. See also *recovery*

multiOS

A cluster that is running multiple operating systems, such as IRIX and Solaris.

multiport serial adapter cable

A device that provides four DB9 serial ports from a 36-pin connector.

node

A *node* is an operating system (OS) image, usually an individual computer. (This use of the term *node* does not have the same meaning as a node in an SGI Origin 3000 or SGI 2000 system and is different from the NUMA definition for a brick/blade on the end of a NUMALink cable.)

A given node can be a member of only one pool (and therefore) only one cluster.

See also *administration node*, *client-only node*, *server-capable administration node*, and *standby node*

node ID

An integer in the range 1 through 32767 that is unique among the nodes defined in the pool. You must not change the node ID number after the node has been defined. It is not the same thing as the *cell ID*.

node membership

The list of nodes that are active (have CXFS kernel membership) in a cluster.

node timeout

If no heartbeat is received from a node in this period of time, the node is considered to be dead. The node timeout value must be at least 10 times the heartbeat interval for proper CXFS operation.

notification command

The command used to notify the cluster administrator of changes or failures in the cluster and nodes. The command must exist on every node in the cluster.

owner host

A system that can control a node remotely, such as power-cycling the node. At run time, the owner host must be defined as a node in the pool.

owner TTY name

The device file name of the terminal port (TTY) on the *owner host* to which the system controller is connected. The other end of the cable connects to the node with the system controller port, so the node can be controlled remotely by the owner host.

pool

The *pool* is the set of nodes from which a particular cluster may be formed. Only one cluster may be configured from a given pool, and it need not contain all of the available nodes. (Other pools may exist, but each is disjoint from the other. They share no node or cluster definitions.)

A pool is formed when you connect to a given node and define that node in the cluster database using the CXFS GUI. You can then add other nodes to the pool by defining them while still connected to the first node, or to any other node that is already in the pool. (If you were to connect to another node and then define it, you would be creating a second pool).

port password

The password for the system controller port, usually set once in firmware or by setting jumper wires. (This is not the same as the node's root password.)

potential metadata server

A server-capable administration node that is listed in the metadata server list when defining a filesystem; only one node in the list will be chosen as the active metadata server.

quorum

The number of nodes required to form a cluster, which differs according to membership:

- For CXFS kernel membership:
 - A majority (>50%) of the server-capable nodes in the cluster are required to **form** an initial membership
 - Half (50%) of the server-capable nodes in the cluster are required to **maintain** an existing membership
- For cluster database membership, **50%** of the **nodes in the pool** are required to form and maintain a cluster.

recovery

The process by which a node is removed from the CXFS kernel membership due to an interruption in CXFS services. It is during this process that the remaining nodes in the CXFS kernel membership resolve their state for cluster resources owned or shared with the removed node. See also *metadata server recovery*

relocation

The process by which the metadata server moves from one node to another due to an administrative action; other services on the first node are not interrupted.

reset

The failure policy method that performs a system reset via a serial line connected to the system controller. The reset may be a powercycle, serial reset, or NMI (nonmaskable interrupt).

server-capable administration node

A node that is installed with the `cluster_admin` product and is also capable of coordinating CXFS metadata.

server-side licensing

Licensing that uses license keys on the CXFS server-capable nodes; it does not require node-locked license keys on CXFS client-only nodes. The license keys are node-locked to each server-capable node and specify the number and size of client-only nodes that may join the cluster membership. All nodes require server-side licensing.

shutdown

The fail action hierarchy selection that tells the other nodes in the cluster to wait before reforming the CXFS kernel membership. The surviving cluster delays the beginning of recovery to allow the node time to complete the shutdown. See also *forced CXFS shutdown*.

snooping

A security breach involving illicit viewing.

split-brain syndrome

A situation in which multiple clusters are formed due to a network partition and the lack of reset and/or CXFS tiebreaker capability.

spoofing

A security breach in which one machine on the network masquerades as another.

standby node

A server-capable administration node that is configured as a potential metadata server for a given filesystem, but does not currently run any applications that will use that filesystem.

static heartbeat monitoring

Monitors constantly at 1-second intervals and declares a timeout after 5 consecutive missed seconds (default). See also *dynamic heartbeat monitoring*.

storage area network (SAN)

A dedicated, high-speed, scalable network of servers and storage devices designed to enhance the storage, retrieval, and management of data

system controller port

A port sitting on a node that provides a way to power-cycle the node remotely. Enabling or disabling a system controller port in the cluster database tells CXFS whether it can perform operations on the system controller port.

system log file

Log files in which system messages are stored

tiebreaker node

See *CXFS tiebreaker node*.

user-space membership

See *cluster database membership*.

view area

The portion of the GUI window that displays components graphically. See also *details area*.

Index

6.5.12f and earlier filesystem conversion, 135
64-bit scalability, 5
100baseT, 39

A

access control lists, 5
ACLs, 5
activate CXFS services
 cmgr, 605
 cxf_admin, 279
 GUI, 227
ACTIVE cluster status, 402
active metadata server, 12
add a node
 cmgr, 577
 cxf_admin, 270
 GUI, 216
add nic, 577
address resolution protocol (ARP), 542
admin command (cmgr), 563
administration, 305
administration best practices, 152
administration daemon, 128
administration membership, 19
administration node, 14
administration software installation
 SGI ProPack, 116
administration tools, 43
advisory record locks, 9
affinity and XVM failover v2, 48
age, 404
allocation of space, 5
allow CXFS kernel membership
 cmgr, 610
 GUI, 232

alternate logins, 187
alternate root, 104, 109, 110
analyze I/O performance, 198
apache server, 186
ARP, 542
asynchronous buffering techniques, 5
atime, 504
AutoLoad boot parameter, 135
automatic restart of nodes, 135
AVT, 376

B

B-trees, 5
backups, 342
backups and client nodes, 153
bandwidth, 5, 8
baseboard management controller
 See "BMC", 539
best practices, 139
block size, 149, 537
blue text, 194
BMC, 34, 539, 582, 583
Brocade switch, 67
Brocade switch GUI, 142
Brocade Web Tools V2.0, 142
BSD interfaces, 9
buffer cache activity, 433, 442
buffer coherency, 36
buffering disks, 8
bufview, 433, 442
build a cmgr script automatically, 632
build_cmgr_script command, 633
bulkstat, 328

C

cad

- messages, 477
 - options file, 128, 129
 - process, 24, 164, 496
 - processes, 129
 - verify it is running, 165
- cad daemon and ports, 320
- cad.options file, 129
- can't run remotely, 481
- capacity of the system, 158
- CDB, 12
- cdb-exitop, 481
- cdbreinit, 158, 429, 485, 490
- cdbutil, 442
- cell ID, 383, 404, 447
- cell_tkm_feature_disable, 359
- cfs, 434
- change a filesystem
- cmgr, 619
 - GUI, 244
- change a node
- cmgr, 587
 - GUI, 217
- change an existing node into a client-only node, 324
- Channel traffic, 459
- checklist, 661
- chkconfig, 166
- chkconfig arguments, 529
- chkconfig settings, 321
- clconf_info, 331, 438, 439, 441, 570
- clconfd
- errors, 468
 - options file, 133
 - process, 496
- clconfd-scripts directory, 325
- clconfd.options file, 133
- clearing the database, 485
- cli
- errors, 468
 - messages, 479

- CLI (underlying command line interface), 43
- client administration node, 14, 274, 578
- client removal, 371
- client restoration, 373
- client software installation
- SGI ProPack, 120
- client vnodes, 434
- client-only node, 13, 273, 578
- gathering status from, 405
- client-only node preference, 145
- client-only node tiebreaker, 145
- client.options file, 134
- client_timeout, 449
- cluster (terminology), 10
- cluster administration daemon
- options file, 129
- cluster conversion
- cmgr, 602
 - GUI, 226
- cluster daemon restart, 489
- cluster database
- automatic cmgr script generation, 632
 - backup/restore, 379
 - definition of term, 4
 - membership quorum stability, 142
 - quorum, 142
 - re-creation, 632
 - shutdown, 333
 - terminology, 12
- cluster database membership, 19
- cluster database membership and quorum, 511
- cluster database membership log, 515
- cluster definition
- cmgr, 598
 - GUI, 224
- cluster deletion
- cmgr, 603
 - cxfs_admin, 285
 - GUI, 227
- cluster display
- cmgr, 604

- cxfs_admin, 285
 - GUI, 227
- cluster domain, 342
- cluster ID
 - changing, 158
 - specification
 - cmgr, 598
 - cxfs_admin, 284
 - GUI, 226
- cluster ID changes, 158
- cluster manager tools, 43
- cluster membership, 19
- cluster migration, 415, 646
- cluster mirror status error, 474
- cluster mode, 225
- cluster modification
 - cmgr, 601
 - GUI, 225
- cluster status
 - tools for troubleshooting, 431
 - verification, 402
- cluster tasks
 - cmgr, 597
 - cxfs_admin, 282
 - GUI, 224
- cluster_admin subsystem, 24, 496
- cluster_control subsystem, 24, 496
- cluster_mgr, 43
 - See "cmgr", 558
- cluster_services subsystem, 496
- cluster_status
 - See "clconf_info and cxfs_info", 439
- cmgr
 - automatic script generation, 632
 - c option, 563
 - command line execution, 563
 - exiting, 561
 - help, 559
 - initial configuration and, 564
 - overview, 558
 - script mode, 629
 - scripts and, 561
 - See "configuration tasks", 557
 - shell and, 563
 - template files, 564
- cmgr-create-cluster, 564
- cmgr-create-node, 564
- cmgr-templates directory, 564
- cmond
 - errors, 473
 - process, 164, 496
 - verify it is running, 164
- cms_dead(), 436
- cms_declare_membership(), 436
- cms_fence_timeout, 360
- cms_fence_timeout_action, 360
- cms_follower(), 436
- cms_leader(), 436
- cms_local_fail_action, 352
- cms_nascent(), 436
- cms_reset_error_override, 361
- cms_reset_timeout, 361
- cms_reset_timeout_action, 361
- cms_shutdown(), 436
- cms_stable(), 436
- cms_trace_enable, 362
- cmsd, 496
- cmsd kernel thread, 440
- coexecution
 - communication paths, 503
 - See "FailSafe coexecution", 38
- colors and states, 201
- command buttons, 194
- command line interface, 43
 - See also "cmgr", 558
- commands installed, 56
- communication paths, 498
- comparison
 - network and CXFS filesystems, 7
 - XFS and CXFS, 3
- compute power, 40
- concat creation, 195
- concatenated slice limit, 537

- concatenation, 341
- concepts for CXFS, 10
- configuration best practices, 139
- configuration checker (cxfs-config), 382
- configuration checklist, 661
- configuration consistency, 143
- configuration overview, 56, 61
- configuration tasks
 - cmgr, 557
 - cluster definition, 598
 - cluster deletion, 603
 - cluster display, 604
 - cluster tasks, 597
 - connectivity test, 596
 - convert a FailSafe cluster, 602
 - CXFS services, 605
 - cxfs tiebreaker, 606
 - defaults, 576
 - define a switch, 624
 - delete a switch, 627
 - fail action hierarchy, 580
 - fence lowering, 627
 - fence raising, 626
 - fencing, 580
 - filesystem deletion, 624
 - filesystem modification, 619
 - filesystem mount/unmount, 618
 - log configuration, 608
 - membership allow/revoke, 610
 - metadata server definition, 611
 - metadata server relocation, 623
 - node deletion, 593
 - node display, 595
 - node modification, 587
 - node NMI, 591
 - node power cycle, 491
 - node reset, 491
 - node tasks, 577
 - notify administrator of cluster changes, 598
 - switch modification, 626
 - update switch port information, 627
 - cxfs_admin
 - cluster deletion, 285
 - cluster display, 285
 - cluster tasks, 282
 - CXFS services, 279
 - cxfs tiebreaker, 283
 - define a switch, 296
 - delete a switch, 297
 - fail action hierarchy, 274
 - fencing, 274
 - filesystem deletion, 294
 - filesystem mount/unmount, 292
 - metadata server relocation, 293
 - node deletion, 279
 - node display, 280
 - node tasks, 270
- GUI
 - cluster definition, 224
 - cluster deletion, 227
 - cluster display, 227
 - cluster modification, 225
 - connectivity test, 223
 - convert a FailSafe cluster, 226
 - convert a FailSafe node, 222
 - CXFS services start/stop, 227
 - cxfs tiebreaker, 229
 - delete a switch, 236
 - display a node, 223
 - fence lowering, 237, 577
 - fence raising, 237
 - filesystem deletion, 247
 - filesystem modification, 244
 - filesystem mount/unmount, 245
 - log configuration, 230
 - membership allow/revoke, 232
 - metadata server definition, 242
 - metadata server relocation, 248
 - node addition/removal, 217
 - node definition, 208
 - node deletion, 223
 - node modification, 217
 - node resets, 217

- set up a new filesystem, 171
 - setting up a new cluster, 170
 - switch definition, 233
 - switch modification, 236
 - update switch port information, 236
- configuration tools, 141
- configuration verification, 151
- configure for automatic restart, 135
- configure network interfaces, 94
- configure system files, 127
- configure the serial ports, 99
- configure with the GUI, 168
- configure XVM volumes, 429
- connectivity test
 - cmgr, 596
 - GUI, 223
 - serial, 180
- contacting SGI with problems
 - SGI ProPack, 493
- contiguous allocation of space, 5
- control network, 19, 210
- convert a FailSafe cluster
 - cmgr, 602
 - GUI, 205, 226
- convert a FailSafe node
 - cmgr, 592
 - GUI, 205, 222
- convert filesystem definitions for 6.5.12f and earlier, 135
- convert from FailSafe task, 205
- corpseleader process, 496
- cpuinfo, 493
- create a cluster
 - cmgr, 597
 - cxfs_admin, 282
 - GUI, 224
- create a filesystem metadata server
 - cmgr, 611
 - GUI, 242
- create a node
 - cmgr, 577
 - cxfs_admin, 270
 - GUI, 208
- cron jobs, 153
- crontab, 153, 154, 328
- crsd
 - process, 24, 164, 496
 - verify it is running, 165
- crsd daemon and ports, 320
- crsd errors, 472, 480
- ctime, 504
- cxdetail (cxfsmgr), 189
- CXFS client node, 12
- CXFS client-only node, 14
- CXFS kernel and ports, 320
- CXFS kernel membership
 - allow/revoke
 - cmgr, 610
 - current information, 441
- CXFS kernel membership allow/revoke
 - GUI, 232
- CXFS kernel membership, quorum, and tiebreaker, 512
- CXFS Manager
 - See "configuration tasks, GUI", 170
 - See "GUI", 191
- CXFS membership
 - problems, 440
 - state determination, 435
- CXFS services
 - start
 - cmgr, 605
 - cxfs_admin, 279
 - GUI, 227
 - stop, 179
 - cmgr, 605
 - cxfs_admin, 279
 - GUI, 228
- CXFS shutdown forced, 610
 - cmgr, 606
 - GUI, 232
- CXFS shutdown normal, 335
 - GUI, 228

- CXFS tiebreaker node
 - cmgr, 606
 - cxfs_admin, 283
 - GUI, 229
 - cxfs*mount.sh scripts, 325
 - cxfs*umount.sh scripts, 325
 - cxfs-config, 382
 - cxfs-reprobe, 428
 - cxfs_admin
 - access, 255
 - access permissions, 268
 - advanced mode, 260
 - attributes, 256
 - basic mode, 260
 - c option, 264
 - class, 256
 - command history, 263
 - command line execution, 264
 - defaults, 266
 - editor style, 266
 - error behavior, 266
 - exiting, 269
 - exiting from, 269
 - filesystem tasks, 287
 - help, 258
 - interrupting a command, 264
 - line wrapping, 266
 - lock, 259
 - mode, 266
 - modes, 260
 - network failover tasks, 295
 - object, 256
 - overview, 255
 - prompting mode, 262
 - safety, 259
 - scripts and, 265
 - set command, 266
 - show a filesystem, 294
 - steal, 259
 - switch tasks, 295
 - syntax, 256
 - <TAB> completion, 261
 - cxfs_client
 - process, 496
 - cxfs_client subsystem, 496
 - cxfs_client.options, 134
 - cxfs_client_push_period, 352
 - cxfs_cluster, 166
 - cxfs_dcvn_timeout, 353
 - cxfs_extents_block_size, 357
 - cxfs_extents_delta, 353
 - cxfs_extents_delta_depth, 357
 - cxfs_info, 405
 - cxfs_punch_hole_restrict, 353
 - cxfs_recovery_slowdown, 362
 - cxfs_recovery_timeout_panic, 151, 363
 - cxfs_recovery_timeout_period, 151, 363
 - cxfs_recovery_timeout_stalled, 151, 363
 - cxfs_recovery_timeout_start, 151, 364
 - cxfs_relocation_ok, 354
 - cxfs_server_push_period, 354
 - cxfs_shutdown_time, 358
 - cxfs_token_fault_tolerant, 364
 - cxfs_token_track, 365
 - cxfs_validate_objid, 366
 - cxfs_verify_existence_token, 366
- cxfsdp, 157
 - cxfsd process, 497
 - cxfsd_max, 355
 - cxfsd_min, 355
 - cxfsd_sync_force, 367
 - cxfsdump, 437, 493
 - cxfsfilesystemUpgrade, 135
 - cxfslicense -d, 85
 - cxfsmgr, 189, 429
 - cxtask (cxfsmgr), 189
- ## D
- daemon
 - list of, 495
 - restart, 489

- verify, 164
 - data flow, 504
 - data integrity protection, 146
 - data management API, 328
 - Data Migration Facility, 41, 328
 - data security, 28
 - database
 - clearing, 485
 - dump, 442
 - membership, 19
 - shutdown, 333
 - database membership
 - See "cluster database membership", 19
 - DB9 serial port, 545
 - dcshake process, 497
 - dcvn, 434
 - dcvnlist, 434
 - deactivate CXFS services
 - cmgr, 605
 - cxfs_admin, 279
 - GUI, 227, 228
 - define a cluster
 - cmgr, 598
 - GUI, 224
 - defragmenter, 5
 - defragmenter software, 155
 - delete a cluster
 - cmgr, 603
 - cxfs_admin, 285
 - GUI, 227
 - delete a filesystem
 - cmgr, 624
 - cxfs_admin, 294
 - GUI, 247
 - delete a node
 - cxfs_admin, 279
 - delete a node from the cluster
 - cmgr, 601
 - GUI, 216
 - delete a node from the pool
 - cmgr, 593
 - GUI, 222
 - detach volume elements, 196
 - details area, 192
 - /dev/cxvm directory, 439
 - device block size, 537
 - df, 432
 - diagnostic tasks
 - cmgr, 596
 - differences between IRIX and SGI ProPack
 - paths, 533
 - direct-to-disk I/O, 9
 - disk blocks, free, 432
 - disk buffering, 8
 - disk configuration, 200
 - disk management, 342
 - disk striping, 341
 - disk unlabing, 196
 - display a cluster
 - cmgr, 604
 - cxfs_admin, 285
 - GUI, 227
 - display nodes
 - cmgr, 595
 - cxfs_admin, 280
 - GUI, 223
 - DMAPI, 328
 - DME, 41, 328
 - DNS, 96, 166
 - domain, 342
 - DOWN node state, 407
 - drag and drop, 198
 - drag-and-drop, 200
 - dsvn, 434
 - dump from metadata server, 346
 - dump of the database, 442
 - dynamic heartbeat monitoring, 27
- E**
- eag:mac-ip=label, 398
 - Edit menu, 193

- enterprise client license key, 77
- entitlement ID, 82, 83
- ERROR cluster status, 402
 - /etc/chkconfig, 166
 - /etc/config/netif.options, 97
 - /etc/config/routed.options, 97
 - /etc/exports, 128
 - /etc/flexlm/license.dat, 83
 - /etc/fstab in XFS (not CXFS), 3
 - /etc/hosts file, 95
 - /etc/init.d, 489
 - /etc/init.d/cluster, 315
 - /etc/init.d/cxfs_client, 315
 - /etc/init.d/cxfs_cluster, 315
 - /etc/inittab, 99
 - /etc/ipfilterd.conf, 529
 - /etc/lk/keys.dat, 84
 - /etc/mtab, 4
 - /etc/nsswitch.conf file, 94
 - /etc/services file, 128
 - /etc/sysconfig/locate, 154
- Ethernet network, 36
- examples
 - admin ping, 180
 - administrative communication within one node, 498
 - administrative communication within one node under coexecution, 503
 - bufview, 442
 - clconf_info, 570
 - cluster definition, 568
 - cluster display, 569
 - cmgr, 563
 - cmgr scripts, 562
 - communication between nodes in the pool , 500
 - communication for a node not in a cluster, 501
 - communication paths, 498
 - CXFS Manager GUI, 169
 - CXFS software installation
 - SGI ProPack, 121
 - cxfs_admin, 264
 - cxfs_admin scripts, 265
 - daemon communication within one
 - administration node, 499
 - daemon communication within one
 - administration node under coexecution, 504
 - define filesystem, 571
 - /etc/config/netif.options, 97
 - /etc/config/routed.options, 97
 - /etc/hosts file, 96
 - /etc/inet/hosts file, 96
 - /etc/inittab, 99
 - fs2d logging and tracing, 132
 - fs2d options, 132
 - GUI screens, 190
 - ifconfig, 98, 178
 - IP filtering, 530
 - metadata flow, 505
 - metadata server distribution, 14
 - name services, 96
 - node configuration, 565
 - ping output, 178
 - ping output for SGI ProPack, 98
 - pool and cluster concepts, 11
 - potential metadata servers, 571
 - private network interface test, 178
 - private network interface test for SGI ProPack, 98
 - quit cmgr, 576
 - relocation versus recovery, 23
 - reset, 35
 - .rhosts, 112
 - sar, 441
 - serial port configuration, 99
 - show clusters, 264, 563
 - show nodes in pool, 567
 - show status, 570
 - start CXFS services, 569
 - stop CXFS services, 179
 - test serial reset connection, 179
 - unmount LOFS filesystems, 327
 - verify cluster daemons are running, 165
- exclusive write tokens, 160

exit operations and alternate root, 104, 109, 110
 exitops and alternate root, 104, 109, 110
 exporting a CFS filesystem, 37
 exports file , 128
 extent lists systune, 353
 extent tracking, 159

F

failover2.conf, 48
 FailSafe coexecution
 overview, 38, 390
 terminology, 41
 FailSafe conversion for a cluster
 cmgr, 602
 GUI, 205, 226
 FailSafe conversion for a node
 cmgr, 592
 GUI, 222
 FailSafe membership, 19
 FailSafe membership, quorum, and Tiebreaker, 514
 fam, 156
 fast copy, 157
 FC cable connections and Brocade, 71
 fencing
 lower (enable access)
 cmgr, 627
 GUI, 237
 raise (disable access)
 cmgr, 626
 GUI, 237
 fencing and hardware changes, 367
 fencing and security of data, 31
 fencing port requirements, 320
 file locking, 36
 File menu, 189, 193
 file size/offset maximum, 537
 filestreams, 161
 filesystem block size, 537
 filesystem buffer cache activity monitor, 442
 filesystem comparison, network and CXFS, 7
 filesystem configuration, 149
 filesystem creation, 3
 filesystem defragmenter software, 155
 filesystem deletion
 cmgr, 624
 cxfs_admin, 294
 GUI, 247
 filesystem features of XFS supported by CXFS, 4
 filesystem format changes between 6.5.12f and 6.5.13f, 135
 filesystem growth, 346
 filesystem is inaccessible, 446
 filesystem local-host access, 6
 filesystem maintenance, 344
 Filesystem Manager, 3
 filesystem metadata server definition
 cmgr, 611
 GUI, 242
 filesystem modification
 cmgr, 619
 GUI, 244
 filesystem mounting
 cmgr, 601, 618
 cxfs_admin, 292
 GUI, 245
 overview, 345
 XFS differences, 3
 filesystem mounting/unmounting
 GUI, 246, 247
 filesystem network access, 6
 filesystem reorganizer, 5
 filesystem repair, 154
 filesystem response times, 5
 filesystem restriction of access, 5
 filesystem specifications, 537
 filesystem structure, 3
 filesystem tasks
 cmgr, 611
 GUI, 237
 guided configuration, 171
 filesystem unmounting

- cmgr, 619
- cxfs_admin, 293
- GUI, 245
- filesystem view, 2
- filesystems and FailSafe, 395
- filtering IP traffic, 529
- find, 153
- find and crontab, 154, 328
- Find text field, 197
- firmware for switches, 72
 - Brocade, 67
- forced CXFS shutdown, 337
 - GUI, 228
- forced CXFS shutdown and restart, 159, 338
- forced shutdown, 149
- forced unmount
 - cmgr, 613
 - GUI, 246
 - recommended, 150
- free disk blocks, 432
- fs2d
 - database membership, 19
 - errors, 475, 480
 - options file, 130
 - process, 24, 164, 496
 - verify it is running, 164
- fs2d daemon and ports, 320
- fs2d membership
 - See "CXFS kernel membership", 19
- fs2d.options file, 130
- fs2d_log, 142, 438
- fsr, 155, 328
- fsr_xfs, 5
- fstab in XFS (not CXFS), 3
- function of a node, 209, 273, 578
- fx, 175, 571

G

- gather cluster configuration information, 437
- gigabit ethernet, 39

- GPT labels, 54
- gratuitous ARP, 542
- GRIO, 9
- GRIOV2 license, 77
- growing filesystems, 346
- guaranteed-rate I/O, 9
- GUI
 - configuring with, 168
 - multiple instances, 189
 - overview, 43, 191
 - starting, 188
 - tasks
 - See "configuration tasks", 183
 - web-based version, 186
- GUI and xvm command differences, 200
- GUI will not run, 454
- guided configuration tasks, 192

H

- hardware installed, 493
- hardware inventory, 426
- hardware requirements, 38
- heartbeat considerations, 526
- heartbeat monitor systune, 350
- heartbeat monitoring, 27
- heartbeat network, 19, 210
- heartbeat systune, 349
- heartbeat timing, 20
- help
 - for cmgr, 559, 563
 - for cxfs_admin, 258
 - for GUI, 170
- Help button, 43
- help menu, 194
- hierarchical storage management, 328
- hinv, 426
- host bus adapter, 31
- hostname control network, 210
- hostname resolution, 93

hostname/IP-address pairings, 170
 hosts file, 95
 hsm, 328
 hub, 39
 hung system, 447
 hyperthread and multi-core processors, 75

I

I/O device configuration, 429
 I/O fencing, 147
 I/O fencing and integrity of data, 31
 I/O monitor for XVM, 196
 I/O operations, 5
 I/O overhead, 40
 I/O performance analysis, 198
 icons and states, 201
 icrash, 160, 434, 439, 440
 idbg, 437
 ifconfig, 178
 ifconfig command, 98
 INACTIVE
 cluster status, 402
 node state, 407
 incarnation, 404
 initial cluster configuration
 cmgr, 564
 GUI, 168
 initial configuration, 164
 initial configuration checklist, 661
 inittab, 99
 input instructions, 200
 installation overview, 56, 61
 Intelligent Platform Management Interface
 See "IPMI", 539
 intelplus, 456
 internode communication, 94
 inventory file, 347
 ioconfig, 429
 IP address and control network, 210
 IP address error, 465

IP address, changing, 93
 IP filtering, 529
 IP-address/hostname pairings, 170
 ipfilterd chkconfig, 529
 ipfilterd_inactive_behavior, 529
 IPMI, 539, 583
 IPMI issues, 455
 ipmitool, 456, 539
 IRIX
 FLEXlm license key verification, 107
 IRIX and Linux system administration
 differences, 415
 irix installation, 102
 is_* commands, 577
 item view
 See "details view", 192
 IX brick, 545

J

java files and release level, 187
 Java-enabled web browser access to GUI, 44
 java_plugin, 186
 joining pools, 517
 jumbo frames, 39

K

kernel membership state determination, 435
 kernel running on SGI ProPack, 493
 kernel status tools, 434
 kernel threads, 495
 keys.dat, 84

L

L1, 277, 582
 L1 system controller, 552

- L2, 277, 582
- L2 system controller , 545
- label disks, 195
- LAN, 7
- lanplus, 456, 544
- large cluster, 144
- large clusters, 170
- lcrash, 161
- Legato NetWorker, 342
- library for record locking, 9
- license
 - Brocade, 69
 - GUI and, 454
- license key
 - error, 464
 - obtaining, 82, 84
 - SGI webpage, 92
 - verification on IRIX, 107
 - verification on SGI ProPack, 120
- license.dat, 83
- licensing, 76
- Linux and IRIX system administration
 - differences, 415
- local area networks (LANs), 7
- local domain, 342
- local node, 336
- local node failure systune, 352
- locate, 153
- locks, 9
- lofs filesystems, 327
- log configuration
 - cmgr, 608
 - GUI, 230
- log files, 437
 - consume too much disk space, 446
 - errors, 476
 - list of, 400
 - management, 339
 - monitoring, 400
 - names and sizes, 157
- log in problems, 448
- log-based filesystem, 4

- logging fs2d options, 131
- logical volume creation, 175, 571
- logical volume reconfiguration, 199
- ls, 153
- lsmmod, 493
- lspci, 493
- LUN limit, 537
- LUN numbers supported, 65

M

- MAC label and Trusted IRIX, 398
- maintenance, 305
- mandatory access control label and Trusted IRIX, 398
- mandatory record locks, 9
- manual CXFS startup/shutdown
 - SGI ProPack, 315
- mask and switch definition, 234
- md driver and SGI Altix systems, 116
- membership
 - See "cluster database membership or CXFS kernel membership", 19
- membership delivered, 447
- membership log, 515
- membership quorum stability, 142
- membership weight, 584
- memberships, 509
- memory requirements, 141
- mesg_delay_time, 358
- mesglist, 434
- mesgtcpaccept process, 497
- mesgtcpdiscovery, 497
- mesgtcpmulticast, 497
- mesgtcprcv process, 497
- message trace, 434
- metadata
 - logging, 36
 - terminology, 3, 12
 - transaction examples, 6

metadata flow, 504
 metadata server
 chosen from potential list, 37
 definition
 cmgr, 611
 GUI, 242
 discovery, 329
 model, 36
 order of, 571
 recovery, 332
 relocation
 cmgr, 623
 cxf_admin, 293
 GUI, 248
 terminology, 3, 12
 metadata server dedication, 144
 metadata server removal, 369
 metadata server restoration, 371
 metadata server-capable administration node, 12
 migrating from an IRIX to an SGI ProPack
 cluster, 415, 646
 miniroot installation, 481
 mirror creation, 195
 mirror status error, 474
 mirroring, 341
 missing XVM volumes, 201
 mixed operating system metadata servers, 143
 mkfs, 3, 172, 175, 571
 mkpart, 214, 276, 579
 MMSC, 277, 582
 MMSC serial port, 556
 mode of cluster, 225
 modify a filesystem
 cmgr, 587, 619
 GUI, 244
 modify a node
 cmgr, 587
 GUI, 217
 modules loaded on SGI ProPack, 493
 monitoring tools, 433
 mount
 command, 3, 439

 filesystems, 345
 cmgr, 618
 cxf_admin, 292
 GUI, 245
 options, 242
 points, 242
 see mounted filesystems, 432
 mount filesystems, 571
 mount scripts, 325
 mounted filesystems, showing, 432
 MSC, 277, 582
 MSC system controller, 554
 msenlow, minthigh, 399
 mtcp_hb_period, 349
 mtcp_hb_watchdog, 350
 mtcp_mesg_validate, 356
 mtcp_nodelay, 350
 mtcp_reserve_size, 358
 mtcp_rpc_thread, 350
 mtime, 504
 multiOS cluster, 31
 multiplexer, 35

N

name restrictions, 93
 named pipes, 9
 NetBackup, 342
 netif.options, 97
 Netscape and the Brocade switch GUI, 142
 Netscape server, 186
 netstat, 431
 network
 interface configuration, 94
 network connectivity
 cmgr, 596
 GUI, 223
 network filesystem comparison, 7
 network issues, 140
 network partition, 148

- network requirements, 39
 - network routing, 97
 - network status, 431
 - network switch, 39
 - NetWorker, 342
 - networks, 19
 - new features summary, 663
 - NFS, 7, 343
 - NFS and CXFS, 6
 - NFS exporting, 37, 128, 325
 - NFS serving, 141
 - NIS, 95, 166
 - NMI a node
 - cmgr, 591
 - NMI reset method, 34
 - node
 - isolation, 28
 - state, 406
 - status, 406
 - tasks
 - See "configuration tasks", 208
 - terminology, 10
 - node addition
 - cmgr, 577
 - GUI, 216
 - node conversion
 - cmgr, 592
 - GUI, 222
 - node definition
 - cmgr, 577
 - cxfs_admin, 270
 - GUI, 208
 - node deletion
 - cmgr, 593
 - cxfs_admin, 279
 - GUI, 222
 - node display
 - cmgr, 595
 - cxfs_admin, 280
 - GUI, 223
 - node function, 209, 273, 578
 - node ID changes, 158
 - node ID, changing, 158
 - node modification
 - cmgr, 587
 - GUI, 217
 - node NMI
 - cmgr, 591
 - node power cycle
 - cmgr, 491
 - node removal, 156
 - node reset
 - cmgr, 491
 - GUI, 217
 - node shutdown, 155
 - node status
 - database shutdown, 334
 - forced CXFS shutdown, 338
 - normal CXFS shutdown, 336
 - node status tools, 431
 - node tasks
 - cmgr, 577
 - cxfs_admin, 270
 - GUI, 208
 - node weight, 584
 - node-locked license keys, 76
 - normal CXFS shutdown, 228
 - notify administrator of cluster changes
 - cmgr, 598
 - GUI, 225
 - nsd, 166
 - nss_enterprise.sw.server, 186
 - nsswitch.conf file, 94
 - NT nodes, 343
 - number of nodes supported, 41
 - NVRAM variables, 135
- O**
- old cluster, 440
 - operating system activity data, 433
 - operating system differences, 533

osview, 158, 433
 output to gather
 SGI ProPack, 493
 overview, 61

P

packages installed
 administration, 57
 client only, 56
 packages running on SGI ProPack, 493
 partition, 276, 579
 partition of network and reset of hardware, 149
 path differences, 533
 pcp_eoe, 408
 pcp_eoe.sw.xvm, 408
 peer-to-disk model, 8
 perform tasks, 197
 Performance Co-Pilot, 198, 434
 XVM statistics, 104, 110, 408
 Performance Co-Pilot XVM I/O monitor, 196
 performance monitoring tools, 433
 physical LUN limit, 537
 physical storage tools, 426
 physical volumes, showing, 429
 physvol, 463
 ping, 98, 178
 ping and BMC, 455
 pipes (named), 9
 pmgsvm, 196, 409
 pmie, 434
 pmieconf, 434
 pool, 11
 port configuration, 99
 port usage, 320
 POSIX pathname, 8
 potential metadata server, 12
 power cycle a node
 cmgr, 491
 power cycle reset method, 34
 preinstallation steps, 93

preliminary steps, 164
 private network, 19, 94, 140
 heartbeat and control, 94
 interface test , 98
 private network failover configuration, 368
 private network interface test, 178
 private network required, 39
 Privilege Manager, 188
 problem location, 441
 problem reporting
 SGI ProPack, 493
 processor type on SGI ProPack, 493
 prompt mode for cmgr, 559
 PRUNEFSS, 154
 PRUNEPATHS, 154

Q

QLogic FC switch, 72
 quit cmgr, 576
 quorum, 142
 quorum leader and ports, 320
 quorum stability, 142
 quorums, 509
 quotas, 5, 343

R

race to reset, 149
 RAID, 63
 raise a fence
 cmgr, 626
 GUI, 237
 RDAC, 376
 read and metadata
 flow, 504
 reboot, 486
 rebooting without rejoining the cluster, 488
 reconfigure a cluster, 440

- record locks, 9
- record-locking library, 9
- recovering a two-node cluster, 486
- recovery, 155
 - features, 342
 - filesystem and XFS, 4
 - of the metadata server, 332
 - terminology, 22
- recovery issues with two server-capable nodes, 526
- Recovery process, 496
- recovery timeout mechanism, 151
- recreating the cluster database, 490
- redirecting the console, 552
- reinitialize the database, 429
- rejoin the cluster
 - forced CXFS shutdown, 338, 638
 - normal CXFS shutdown, 336, 638
- release history, 663
- reliability of the filesystem, 4
- relocate a metadata server
 - cmgr, 623
 - cxfs_admin, 293
 - GUI, 248
- relocation, 20, 155
- remove a cluster
 - cmgr, 603
 - cxfs_admin, 285
 - GUI, 227
- remove a filesystem
 - cmgr, 624
 - cxfs_admin, 294
 - GUI, 247
- remove a nod
 - cxfs_admin, 279
- remove a node from the cluster
 - cmgr, 601
 - GUI, 216
- remove a node from the pool
 - cmgr, 593
 - GUI, 222
- remove nic, 577
- removing and restoring cluster members, 369
- reorganizer, 5
- reporting problems
 - SGI ProPack, 493
- requirements, 38
- reset, 40, 514
 - removing, 159
- reset a database, 485
- reset a node
 - cmgr, 491
 - GUI, 217
- reset configurations, 539
- reset lines, 159
- reset methods, 34
- reset race, 149
- reset services daemon, 128
- residual cluster, 440
- restart (avoiding), 339
- restart cluster daemons, 489
- restarting CXFS, 374
- restarting CXFS services, 485
- restarting nodes automatically, 135
- restore, 346
- restricted system tunable parameters, 356
- restrictions with CXFS filesystems, 9
- revoke CXFS kernel membership
 - cmgr, 610
 - GUI, 232
- rfind, 153
- rhel pd_min, 351
- rhost database entry and Trusted IRIX, 398
- /.rhosts file, 111
- root filesystem and CXFS, 9
- rotatelogs, 340
- rotating log files, 339
- routed, 97
- routed.options, 97
- Router Discovery Advertisements, 98
- routing, 97
- RPC port mapper, 320
- rpm, 493
- rsh, 111

S

- Samba, 7, 343
- Samba serving, 141
- SAN
 - documentation, 661
 - use of, 2
- sar, 433, 441
- scalability (64 bit0), 5
- script generation for cmgr, 632
- script mode and cmgr, 629
- scripts
 - cmgr, 561
 - cxfs_admin, 265
 - pre/post-mount scripts, 325
 - templates for cmgr, 564
- selecting items to view or modify, 196, 249, 252
- separate pools, 517
- serial connections
 - cmgr, 596
 - GUI, 223
- serial hardware reset
 - See "reset", 34
- serial port configuration, 99
- serial port multiplexer, 35
- serial port server, 103, 108
- serial reset connection test, 178
- server vnodes, 434
- server-capable administration node, 274, 578
- server-capable node, 12
- server-capable node numbers, 145
- server-side licensing, 76
- services file, 128
- services for CXFS
 - stop, 179
- set commands, 577
- set up a new cluster, 170
- SGI ProPack
 - administration software installation, 116
 - client software installation, 120
 - FLEXlm license key verification, 120
 - manual CXFS startup/shutdown, 315
 - software
 - maintenance, 134
 - space requirements, 121
 - sgi-cad, 128
 - sgi-cmsd, 128
 - sgi-crsd, 128
 - sgi-gcd, 128
 - sgi_apache.sw.server, 186
 - SGIAVT, 376
 - SGIAVT mode, 53
 - SGIRDAC, 376
 - SGIRDAC mode, 53
 - shell use with cmgr, 563
 - shortcuts, 194
 - show a cluster
 - cmgr, 604
 - cxfs_admin, 285
 - GUI, 227
 - show a node
 - cmgr, 595
 - cxfs_admin, 280
 - GUI, 223
 - show clients/servers, 435
 - show command (cmgr), 563
 - show switches, 297, 628
 - shut down nodes unobtrusively, 155
 - shutdown
 - cluster database, 333
 - forced CXFS cluster database, 337
 - normal CXFS
 - cxfs_admin, 275, 581
 - GUI, 228
 - normal CXFS database, 335
 - restart and, 159, 338
 - sinfo, 435
 - single filesystem view, 2
 - single-system view, 38
 - site-changeable system tunable parameters, 348
 - size of the cluster, 41
 - slice disk, 195
 - slocate, 154

- software maintenance
 - SGI ProPack, 134
- software mix, 143
- software requirements, 38
- software upgrades, 157
- space
 - allocation, 5
 - requirements, 102, 108
- space requirements
 - SGI ProPack, 121
- speed, 5
- split-brain syndrome, 146
- standby node, 155
- start
 - CXFS processes
 - SGI ProPack, 315
- start command (cmgr), 563
- start CXFS, 159
- start CXFS services, 336, 638
 - cmgr, 605
 - cxfs_admin, 279
 - GUI, 227
- static heartbeat monitoring, 27
- statistics for an XVM volume, 433
- status
 - cluster, 402
 - node, 406
 - system controller, 655
 - system, overview, 399
- sthreads, 435, 496
- stop command (cmgr), 563
- stop CXFS processes on Solaris, 315
- stop CXFS services
 - cmgr, 605
 - cxfs_admin, 279
 - GUI, 228
- stopping CXFS, 374
- storage area network, 2
- storage management, 151
- storage tools, 426
- stripe creation, 195
- striping, 341
- subnet, 19, 141
- subvolume creation, 195
- summary of new features, 663
- swap to a file, 9
- switch, 31
 - Brocade, 67
 - definition
 - cmgr, 624
 - cxfs_admin, 296
 - GUI, 233
 - display, 297, 628
 - QLogic, 72
 - query, 628
 - recommendations, 39
- switch definition
 - GUI, 233
- switch deletion
 - cmgr, 627
 - cxfs_admin, 297
 - GUI, 236
- switch modification
 - cmgr, 626
 - GUI, 236
- switch port information
 - cmgr, 627
 - GUI, 236
- switch use, 149
- switchshow
 - Solaris, 451
- sysadm_cxfs.sw.desktop, 185
- sysadmdesktop, 187, 188
- sysctrl* commands, 577
- SYSLOG, 193, 346, 438, 439, 441, 447, 466
- SYSLOG errors, 466
- system
 - files, 127
 - software communication paths, 498
 - status, 399
 - view, 38
- system activity, 433
- system activity data, 433

- system administration differences between IRIX and Linux, 415
- system buffer cache activity, 433
- system capacity, 158
- system controller status, 655
- system controller types, 277, 582
- system hang, 447
- system log file, 193
- system reset, 147
- system reset configurations, 539
- system tunable parameters
 - restricted dynamic, 359
 - restricted static, 357
 - site-changeable dynamic, 352
 - site-changeable static, 349
- System V interfaces, 9

T

- Tape Management Facility (TMF), 41
- tasks, 197
- Tasks menu, 193
- TCP port, 320
- TCP/IP network, 36
- TCP/IP network requirements, 39
- tcp_channels, 436
- tcpmux, 455
- tcpmux/sgi_sysadm, 455
- telnet
 - limit simultaneous sessions for Brocade, 70
- telnet and fencing, 72
- telnet and I/O fencing, 147
- telnet port and I/O fencing, 31
- template files, 564
- temporary names, 195
- term definitions, 200
- terminology, 10
- test command (cmgr), 563
- test connectivity
 - cmgr, 596
 - GUI, 223

- Test Connectivity task, 441
- test the system, 178
- tgconfig, 39
- threads, 495
- tiebreaker node
 - cmgr, 606
 - cxfs_admin, 283
 - GUI, 229
- tiebreakers, 145, 512
- tigon, 39
- TMF, 41
- tokens, 160, 504
- Toolchest, 185
- tools, 43
- trace messages, 434
- transaction rate, 5
- transaction rates, 7
- transform an existing node into a client-only node, 324
- tree view
 - See "view area", 192
- TRIX (Trusted IRIX), 41
- troubleshooting, 425
 - cad messages, 477
 - cannot access filesystem, 446
 - cell ID count and "Membership delivered" messages, 447
 - clconfd daemon death, 462
 - clconfd errors, 468
 - clearing the database, 485
 - cli errors, 468
 - cli messages, 479
 - cluster configuration tools, 429
 - cluster control tools, 430
 - cluster daemon restart, 489
 - cluster database membership quorum stability, 142
 - cluster/node status tools, 431
 - cmond errors, 473
 - common problems, 442
 - corrective actions, 484

- crsd errors, 472, 480
- eliminating a residual cluster, 440
- exclusive write tokens, 160
- fs2d errors, 475, 480
- general messages, 476
- GUI use, 142
- GUI will not run, 454
- identify the cluster status, 438
- IP address error, 465
- kernel status tools, 434
- license key error, 464
- locate the problem, 441
- log file errors, 476
- log files, 437
- log files consume too much disk space, 446
- lost CXFS membership, 464
- messages, 458
- networking tools, 431
- no cluster name ID error, 463
- node is detected but never joins
 - membership, 447
- normal messages, 458
- out of logical swap space, 462
- performance monitoring tools, 433
- physical storage tools, 426
- reboot, 486
- reboot before changing node ID or cluster ID, 158
- rebooting without rejoining the cluster, 488
- recovering a two-node cluster, 486
- recreating the cluster database, 490
- removing reset lines, 159
- restart CXFS after a forced CXFS shutdown, 159
- restarting CXFS services, 485
- strategy, 425
- SYSLOG errors, 466
- system is hung, 447
- tools, 426
- unable to define a node, 447
- unwritten extent tracking, 160
- xfs_repair appropriate use, 154
- you cannot log in, 448

- trusted IRIX, 41, 157

U

- UDP port, 320
- uname, 102, 108, 493
- UNKNOWN cluster status, 402
- UNKNOWN node state, 407
- unlabel disks, 196
- unmount (forced)
 - cmgr, 613
 - GUI, 246
 - recommended, 150
- unmount a filesystem
 - cmgr, 619
 - cxfs_admin, 293
 - GUI, 246, 247
- unmount LOFS filesystems, 327
- unmount scripts, 325
- unmounting, 345
- unwritten extent tracking, 159
- UP node state, 406
- updatedb, 154
- UPDATEDB_PRUNEPATHS, 154
- upgrading
 - transform an existing node into a client-only node, 324
- upgrading from 6.5.12f or earlier, 135
- user jobs and metadata server, 153
- user-space membership, 19
- /usr/bin/icrash, 434
- /usr/cluster/bin/clconf_info, 570

V

- /var/adm/SYSLOG, 193, 438, 439, 441, 447, 466
- /var/cluster/clconfd-scripts directory, 325
- /var/cluster/cmgr-scripts/rotatelogs, 340
- /var/cluster/cmgr-templates, 564

/var/cluster/ha/log/fs2d_log, 438
 /var/cluster/ha/tmp/cmgr_create_cluster_clustername_process, 632
 /var/flexlm/license.dat, 83
 /var/log/messages, 193
 /var/xfsdump/inventory, 347
 verify cluster daemons are running, 164
 verify that chkconfig arguments are on, 166
 VERITAS NetBackup, 342
 view area, 192
 view cluster components, 196
 view component details, 197
 VLAN, 141
 vnodes, 434
 volume creation, 175, 195, 571
 volume element detach, 196
 volume header creation, 175, 571
 volume management, 341
 volume manager, 40
 volume topologies, 199
 volume-element deletion, 196

W

web browser access to GUI, 44
 web-based version of the GUI, 186
 weight, 584
 weighted nodes (server-capable administration nodes), 13
 what is CXFS?, 2
 when to use CXFS, 5
 Windows nodes, 343
 workstation client license key, 77
 worldwide node name, 31
 worldwide port name
 SGI ProPack, 450
 write and metadata flow, 504
 write tokens, 160
 WWN discovery, 375

WWNN, 31
 WWNN, 632
 SGI ProPack, 450

X

X/Open Data Storage Management Specification, 328
 XFS
 comparison to CXFS, 3
 features supported, 4
 xfs
 quotas, 343
 xfs_fsr, 5, 155
 xfs_repair, 155
 xfs_repair appropriate use, 154
 xfsd, 495
 xfsdump and xfsrestore, 346
 XSDM, 328
 xthreads, 497
 XVM, 341
 logical volume creation, 175, 571
 requirement, 40
 statistics, 408
 xvm, 429, 432, 433, 439, 440
 XVM cluster mirror status error, 474
 XVM failover, 45
 XVM I/O monitor, 196
 XVM license, 77
 XVM shortcuts, 194
 XVM volume mapping to storage targets, 375
 XVM volumes, 429

Z

zoning, 31