

CXFS™ MultiOS Client-Only Nodes: Installation and Configuration Guide

007-4507-004

CONTRIBUTORS

Written by Lori Johnson

Edited by Susan Wilkening

Illustrated by Chrystie Danzer

Production by Glen Traefald

Engineering contributions to the book by Neil Bannister, Dale Brantly, David Chatterton, Mark Cruciani, Tad Dolphay, Dave Ellis, Eric Eppe, Andrew Gildfind, Dennis Kender, Aaron Mantel, Troy McCorkell, Ken McDonell, Terry Merth, Daniel Moore, Max Matveev, Barry Naujok, Tim Sirianni, Michael Umansky, Madan Valluri

COPYRIGHT

© 2002 Silicon Graphics, Inc. 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 Silicon Graphics, Inc.

LIMITED RIGHTS LEGEND

The electronic (software) version of this document was developed at private expense; if acquired under an agreement with the USA government or any contractor thereto, it is acquired as "commercial computer software" subject to the provisions of its applicable license agreement, as specified in (a) 48 CFR 12.212 of the FAR; or, if acquired for Department of Defense units, (b) 48 CFR 227-7202 of the DoD FAR Supplement; or sections succeeding thereto. Contractor/manufacturer is Silicon Graphics, Inc., 1600 Amphitheatre Pkwy 2E, Mountain View, CA 94043-1351.

TRADEMARKS AND ATTRIBUTIONS

Silicon Graphics, SGI, the SGI logo, IRIX, and Origin are registered trademarks and CXFS, FailSafe, IRIS FailSafe, O2, Trusted IRIX, and XFS are trademarks of Silicon Graphics, Inc.

AIX is a registered trademark of IBM Corp. Brocade and Silkworm are trademarks of Brocade Communication Systems, Inc. Disk Manager is a registered trademark of ONTRACK Data International, Inc. FibreStar and JNI are registered trademarks and EZ Fibre is a trademark of JNI Corporation. GLOBEtrotter and FLEXIm are registered trademarks of GLOBEtrotter Software and Macrovision Corporation. HP-UX is a trademark of Hewlett-Packard Company. InstallShield is a registered trademark of InstallShield Software Corporation in the United States and/or other countries. Intel and Pentium are registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries. Legato NetWorker is a registered trademark of Legato Systems, Inc. Linux is a registered trademark of Linus Torvalds. Microsoft, Windows, and Windows NT are registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. SANsurfer and QLogic are registered trademarks of QLogic Corporation. Sun, SunOS, Solaris, and Ultra Enterprise are trademarks or registered trademarks of Sun Microsystems, Inc. UNIX and the X device are registered trademarks of The Open Group in the United States and other countries.

Screen snaps of the EZ Fibre product are provided by express permission of JNI Corporation.

Cover design by Sarah Bolles, Sarah Bolles Design, and Dany Galgani, SGI Technical Publications.

New Features in this Guide

This guide contains the following new features:

- Support for Microsoft Windows 2000 nodes in a CXFS MultiOS cluster. This guide uses *Windows* to refer to both Microsoft Windows NT and Microsoft Windows 2000 systems. See Chapter 6, "Windows NT and Windows 2000 Platforms", page 53.
- Support for SGI TP9100s; see Chapter 2, "SGI RAID Firmware", page 13. For additional details, see the release notes.
- A new section: "Configuring Two HBAs for Failover Operation", page 63.
- Support for the JNI 5.1.1 and later driver on Solaris clients, which simplifies the installation steps; see "JNI Fibre Channel Host Bus Adapter Installation and Configuration", page 26.
- DMAPI support for all platforms.
- Removal of the Solaris limitation requiring more kernel threads.

Record of Revision

Version	Description
001	March 2002 Original publication with the CXFS MultiOS Clients 2.0 release for IRIX 6.5.16f.
002	May 2002 Revised to support the CXFS MultiOS Clients 2.1 release for IRIX 6.5.16f. This release supports the Sun Microsystems Solaris and Microsoft Windows NT platforms.
003	June 2002 Revised to support the CXFS MultiOS 2.1.1 release for IRIX 6.5.16f. This release supports the Sun Microsystems Solaris and Microsoft Windows NT platforms.
004	August 2002 Revised to support the CXFS MultiOS 2.2 release for IRIX 6.5.17f. This release supports the Sun Microsystems Solaris, Microsoft Windows NT, and Microsoft Windows 2000 platforms.

Contents

About This Guide	xv
Prerequisites	xv
Related Publications	xv
Obtaining Publications	xvii
Conventions	xvii
Reader Comments	xviii
1. Introduction	1
When to Use CXFS	2
CXFS on Client-Only Nodes	2
CXFS Processes	3
Licenses	3
Cluster Administration	3
User Administration for CXFS	4
Performance Considerations	5
Requirements	5
Recommendations	8
Overview of the Installation and Configuration Steps	10
Solaris Overview	10
Windows Overview	11
2. SGI RAID Firmware	13
Required SGI RAID Firmware	13
Required SGI TP9400 RAID Firmware	13
Required SGI TP9100 RAID Firmware	13

RAID Firmware Verification	13
For More Information	14
3. Brocade Fibre Channel Switch Verification	15
Required Brocade Fibre Channel Switch Firmware and License	15
Verifying the Brocade License	16
Verifying the Brocade Switch Firmware Version	16
Changing to the Brocade FC Cable Connections	16
4. Obtaining CXFS and XVM FLEXlm Licenses	19
Obtain the Host Information Required for the License	19
Solaris Host Information	19
Windows Host Information	20
Obtaining and Install the Licenses	20
For More Information	20
5. Solaris Platform	21
CXFS on Solaris	21
Requirements Specific to Solaris	21
CXFS Commands Installed on Solaris	22
Solaris Log Files	23
Solaris Limitations and Considerations	23
Access Control Lists and Solaris	24
FLEXlm License Verification for Solaris	25
JNI Fibre Channel Host Bus Adapter Installation and Configuration	26
Installing the JNI HBA	26
Installing and Running the EZ Fibre Configuration GUI	28
Verifying the JNI HBA Installation	37
Preinstallation Steps for Solaris	38

Hostname Resolution and Network Configuration Rules for Solaris	39
Adding a Private Network for Solaris Nodes	40
Verifying the Private and Public Networks for Solaris	45
Client Software Installation Steps for Solaris	45
Installation Overview	46
Verifying the Solaris Installation	47
Postinstallation Steps for Solaris: Creating the I/O Fencing File	48
Manual CXFS Startup/Shutdown for Solaris	50
Software Maintenance for Solaris	50
Upgrading the CXFS for Solaris Software	51
Modifying the CXFS for Solaris Software	51
6. Windows NT and Windows 2000 Platforms	53
CXFS on Windows	53
Requirements Specific to Windows	54
CXFS Commands Installed on Windows	54
Windows Log Files	55
Functional and Performance Limitations Specific to Windows	55
UNIX Perspective of CXFS on a Windows Node	55
Windows Perspective of CXFS on a Windows Node	57
Performance Considerations on a CXFS Windows Node	59
Forced Unmount on a Windows Node	60
QLogic Fibre Channel Host Bus Adapter Installation for Windows	60
Confirming the QLogic HBA Installation	61
Support for More than 8 LUNs under Windows NT	61
Configuring Two HBAs for Failover Operation	63
Preinstallation Steps for Windows	64
Hostname Resolution and Network Configuration Rules for Windows	65

Adding a Private Network for Windows Nodes	65
Adding a Private Network for Windows NT Nodes	66
Adding a Private Network for Windows 2000 Nodes	67
Verifying the Private and Public Networks for Windows	68
Client Software Installation Steps for Windows	69
Postinstallation Steps for Windows	73
Configuring the FLEXlm License for Windows	74
Creating the Windows I/O Fencing File	75
Performing User Configuration	77
Checking Permissions on the Password, Group, and Fencing Files	78
Creating a New Hardware Profile	78
Manual CXFS Startup/Shutdown for Windows	81
Software Maintenance for Windows	82
Modifying the CXFS for Windows Software	82
Upgrading the CXFS for Windows Software	84
Removing the CXFS for Windows Software	84
Downgrading the CXFS for Windows Software	85
7. Cluster Configuration	87
Defining the Client-Only Nodes	87
Adding the Client-Only Nodes to the Cluster	89
Defining the Switch for I/O Fencing	90
Starting CXFS Services on the Client-Only Nodes	90
Mounting Filesystems on the Client-Only Nodes	91
Restarting the Windows Node	91
Verifying the Cluster	92
Forced Unmount of CXFS Filesystems	93

8. Troubleshooting	95
Identifying Problems on Solaris Nodes	95
Is the Solaris Node in the Cluster?	95
Are there Error Messages for the Solaris Node?	95
Identifying Problems on Windows Nodes	96
Is the CXFS Software Running Correctly on the Windows Node?	96
Windows NT CXFS Software Verification	96
Windows 2000 CXFS Software Verification	97
Is the Windows Node in the Cluster?	98
Are There Error Messages for the Windows Node?	98
Windows Error Message Explanations	99
Common Problems and Solutions	101
Incorrect Configuration	101
Determining If A Client-Only Node Is Fenced	101
Common Solaris Problems: JNI Problems	101
Common Windows Problems	102
Windows QLogic Problems	102
Filesystems Are Not Displayed on a Windows Node	103
Large Log Files on Windows	103
Windows Failure on Restart	104
Memory Configuration of the Windows Node	104
Reporting Problems	104
Reporting Solaris Problems	104
Reporting Windows Problems	105
Retain Windows NT Information	106
Retain Windows 2000 Information	107
Save Crash Dumps for Windows	108

Contents

Glossary 109

Index 115

Figures

Figure 5-1	Example: Second Window: EZ Fibre Configuration Utility - Standalone . . .	30
Figure 5-2	Location of icon (+) to Display the HBA	31
Figure 5-3	Example: After Clicking + to Display the HBA	32
Figure 5-4	Location of the Icon to Display the Adapter Parameters	33
Figure 5-5	Example: After Clicking the HBA Icon to Show the Adapter Parameters . . .	34
Figure 5-6	After Clicking the Adapter Information Tab	35
Figure 5-7	After Clicking on LUN-Level Zoning	36
Figure 5-8	Example: After Mapping the LUNs and Committing the Changes	37
Figure 6-1	Select All Software Components	70
Figure 6-2	Enter CXFS Details	71
Figure 6-3	Review the Settings	72
Figure 6-4	Start CXFS Driver	73
Figure 6-5	Modify the CXFS for Windows	83

About This Guide

This publication documents the CXFS MultiOS 2.2 release for IRIX 6.5.17f. This release supports Sun Microsystems Solaris nodes, Microsoft Windows NT nodes, and Microsoft Windows 2000 nodes.

Prerequisites

This guide assumes the following:

- The reader is familiar with the information presented in the *CXFS Version 2 Software Installation and Administration Guide* and the operating system documentation.
- The IRIX CXFS cluster is installed and operational.
- The CXFS client-only nodes have the appropriate platform-specific operating system software installed.

Related Publications

The following documents contain additional information (if you are viewing this document online, you can click on [TPL Link](#) below to link to the book on the SGI TechPubs library):

- CXFS documentation:
 - Platform-specific release notes
 - *CXFS Version 2 Software Installation and Administration Guide* ([TPL link](#))
- SGI TP9400 documentation :
 - *TPSSM7 RAID Software Concepts Guide for TP9400* ([TPL link](#))
 - *SGI TP9400 RAID Owner's Guide* ([TPL link](#))
 - *SGI TP9400 RAID Administration Guide* ([TPL link](#))

The SGI TP9400 documentation is available on the release CD in the following files:

- `tp9400_sw_concepts_guide.pdf`

- `tp9400_owners_guide.pdf`
 - `tp9400_admin_guide.pdf`
 - SGI TP9100:
 - *TPM Installation Instructions and User's Guide for TP9100*
 - JNI host bus adapter (HBA) card and driver documentation:
 - *Installation Guide, FCE-6460 and FCE2-6460 PCI-to-Fibre Channel Host Bus Adapters (Solaris, Windows NT/2000, Novell, AIX, HP-UX, MAC-OS, Linux) JNI FibreStar*
 - *Quick Installation Guide, Solaris, AIX and Windows JNI EZ Fibre*
- Also see the JNI website at:
- <http://www.jni.com>
- QLogic HBA card and driver documentation:
 - *Hardware Installation Guide for the QLA2xxx Board Family*
 - *Software Installation Guide for the QLA2xxx Board Family*
- See the QLogic website at:
- <http://www.qlogic.com>
- Solaris documentation:
 - *Solaris 8 Installation Guide*
 - *Solaris 8 System Administration Collection*
 - *Solaris 8 Advanced Installation Guide*
- See the Sun Microsystems website at:
- <http://www.sun.com>
- Sun Microsystems hardware documentation:
 - *Ultra Enterprise 450 Systems Manual*
 - *Ultra Enterprise 6000/5000/4000 Systems Manual*
 - *Ultra Enterprise 6000/5000/4000 Systems Installation Guide*

- *Ultra Enterprise 10000 SSP 3.0 User's Guide*
- Windows software documentation:
 - *Start Here Microsoft Windows NT Workstation: Basics and Installation*
 - *Microsoft Windows 2000 Quick Start Guide*
- Hardware documentation for the Intel platform
- *Flexible License Manager End User Manual* from GLOBEtrotter Software, Inc.

Obtaining Publications

To obtain SGI documentation, go to the SGI Technical Publications Library at:

<http://techpubs.sgi.com>.

Conventions

Note: 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.

The following conventions are used throughout this document:

Convention	Meaning
<code>command</code>	This fixed-space font denotes literal items such as commands, files, routines, path names, signals, messages, and programming language structures.
<i>variable</i>	Italic typeface denotes variable entries and words or concepts being defined.
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.

GUI	This font denotes the names of graphical user interface (GUI) elements such as windows, screens, dialog boxes, menus, toolbars, icons, buttons, boxes, fields, and lists.
[]	Brackets enclose optional portions of a command or directive line.
...	Ellipses indicate that a preceding element can be repeated.

Reader Comments

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

You can contact us in any of the following ways:

- Send e-mail to the following address:

`techpubs@sgi.com`

- Use the Feedback option on the Technical Publications Library World Wide Web page:

`http://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:

Technical Publications
SGI
1600 Amphitheatre Pkwy., M/S 535
Mountain View, California 94043-1351

- Send a fax to the attention of "Technical Publications" at +1 650 932 0801.

We value your comments and will respond to them promptly.

Introduction

This guide provides an overview of the installation and configuration procedures for CXFS client-only nodes running SGI CXFS clustered filesystems. A *CXFS client-only node* runs a subset of CXFS processes and services.

This release supports client-only nodes running the following operating systems:

- Solaris 8
- Windows NT 4.0 Service Pack 6
- Windows 2000 Service Pack 2

Note: 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.

A cluster running multiple operating systems is known as a *multiOS cluster*.

Many of the procedures mentioned in this guide will be performed by SGI personnel or other qualified service personnel. Details for these procedures are provided in other documents.



Caution: CXFS is a complex product. To ensure that CXFS is installed and configured in an optimal manner, it is **mandatory** that you purchase SGI installation services developed for CXFS. Contact your local SGI sales representative for details.

For general information about CXFS terminology, concepts, and configuration, see the *CXFS Version 2 Software Installation and Administration Guide*.

This chapter discusses the following:

- "When to Use CXFS", page 2
- "CXFS on Client-Only Nodes", page 2
- "Overview of the Installation and Configuration Steps", page 10

When to Use CXFS

You should use CXFS when you have multiple hosts 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.
- All processes that perform reads/writes for a given file reside on the same host.
- Multiple processes on multiple hosts read the same file.
- Direct-access I/O is used for reads/writes for multiple processes on multiple hosts.
- Large files and file accesses are being used.

For most filesystem loads, the preceding scenarios 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. (*Metadata* is information that describes a file, such as the file's name, size, location, and permissions.) CXFS is faster than NFS because the data does not go through the network.

CXFS on Client-Only Nodes

This section contains the following:

- "CXFS Processes", page 3
- "Licenses", page 3
- "Cluster Administration", page 3
- "User Administration for CXFS", page 4
- "Performance Considerations", page 5
- "Requirements", page 5
- "Recommendations", page 8

CXFS Processes

When CXFS is started on a client-only node, a user-space daemon/service is started that provides the required processes. This is a subset of the processes needed on an IRIX node.

Licenses

You must have the following licenses:

- Brocade license. See "Required Brocade Fibre Channel Switch Firmware and License", page 15.
- CXFS FLEXlm license installed on every node in the cluster; see Chapter 4, "Obtaining CXFS and XVM FLEXlm Licenses", page 19.

Note: XVM provides a mirroring feature. If you want to access a mirrored volume from a given node in the cluster, you must install the "XVM Volume Plexing" FLEXlm license on that node. Only those nodes that will access the mirrored volume must be licensed. For information about purchasing this license, see your SGI sales representative.

Cluster Administration

There must be at least one IRIX node in the cluster that is responsible for updating that filesystem's metadata. This node is referred to as the *CXFS metadata server*. Only IRIX nodes can be metadata servers; client-only nodes cannot be metadata servers. The CXFS cluster database is not stored on client-only nodes; only IRIX nodes contain the cluster database.

An IRIX node is required to perform administrative tasks, using either the `cmgr(1M)` command or the CXFS graphical user interface (GUI). For more information about using these tools, see the *CXFS Version 2 Software Installation and Administration Guide*.

Note: The premount and postmount scripts are supported on IRIX and Solaris nodes; they are not supported on Windows nodes. The scripts behave the same on IRIX and Solaris nodes, but the pathnames on Solaris are as follows:

```
/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
```

For information about using these scripts on IRIX nodes or Solaris nodes, see the *CXFS Version 2 Software Installation and Administration Guide*.

User Administration for CXFS

A CXFS cluster requires a consistent user identification scheme across all hosts in the cluster so that one person using different cluster nodes has the same access to the files on the cluster.

The following must be observed to achieve this consistency:

- Users must have the same usernames on all nodes in the cluster. An individual user identifier (UID) should not be used by two different people anywhere in the cluster. Ideally, group names should also be the same on all nodes in the cluster.
- The `/etc/passwd` and `/etc/group` files from the CXFS metadata server must be installed on the client-only node. These files are used to determine the equivalent UNIX UID and group identifiers (GIDs) of each user.

If these files are modified on the CXFS metadata server, the files should be reinstalled on the client-only node.

Note: Under Windows, the CXFS software will detect that these files have changed and will apply the updated contents when mapping Windows users to UNIX UIDs. The Administrator user on the Windows node will be mapped to the root user on the UNIX nodes for the purpose of file access controls. Therefore, access to the Administrator user account on Windows nodes should be treated with the same caution as the root user on UNIX nodes.

Performance Considerations

CXFS may not give optimal performance under the following circumstances:

- When you are using NFS to export a CXFS filesystem from a CXFS client. Performance will be much better when the export is performed from a CXFS metadata server than when it is performed from a CXFS client-only node.
- 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 when using 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.

- When distributed applications write to shared files that are memory mapped.

Also see "Functional and Performance Limitations Specific to Windows", page 55.

Requirements

Using a client-only node in a CXFS cluster requires the following:

- A supported SAN hardware configuration.

Note: For details about supported hardware, see the Entitlement Sheet that accompanies the base CXFS release materials. Using unsupported hardware constitutes a breach of the CXFS license. CXFS does **not** support the Silicon Graphics O2 workstation as a CXFS node nor does it support JBOD.

- At least one IRIX node to act as the metadata server and from which to perform cluster administration tasks. CXFS should be installed on the IRIX node before CXFS is installed on the client-only node.
- A private 100baseT TCP/IP network connected to each node, to be dedicated to the CXFS private heartbeat and control network. All nodes must be configured to use the same subnet.
- IRIX 6.5.16f or later, plus appropriate patches, installed on the IRIX nodes in a multiOS cluster. In a multiOS cluster, all IRIX nodes must be at a minimum of IRIX 6.5.16f or later, plus appropriate patches to allow a CXFS Solaris or Windows node to be present. For more details, see the platform-specific CXFS MultiOS release notes.
- A FLEXlm license key for CXFS and optionally XVM. The CXFS license is required for all nodes in the pool; an XLV license is required for each node from which you want to access a mirrored XVM volume.
- An SGI-sold and supported Brocade Fibre Channel 2400, 2800, or 3800 switch, which is required to protect data integrity.

IRIX nodes use serial reset lines or I/O fencing to protect the integrity of the data stored in the cluster. (One of these methods is mandatory for the IRIX nodes in a cluster with only two IRIX nodes or with only two weighted IRIX nodes. Larger clusters should have an odd number of weighted IRIX nodes.)

The *I/O fencing* feature isolates a problem node so that it cannot access I/O devices and therefore cannot corrupt data in the shared CXFS filesystem. This feature can only be used with a Brocade Fibre Channel switch; therefore, the Brocade switch is a required piece of hardware in a multiOS cluster.

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 only access 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. Instead, SGI chose to implement a simpler form of fencing, enabling/disabling a host's Brocade ports.

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.



Caution: Do not perform a `telnet` to the switch and leave the session connected.

- A cluster of no more than 16 nodes. At least one IRIX node must have weight in order to be a potential metadata server; other nodes can be CXFS clients. All Solaris nodes and Windows nodes are CXFS client-only nodes.

A cluster in which both CXFS and IRIS FailSafe 2.1 or later are run (known as *coexecution*) is supported with a maximum of 16 nodes, as many as 8 of which can run FailSafe. However, FailSafe cannot run on Solaris nodes or Windows nodes.

- No nodes within the cluster running Trusted IRIX. A multiOS cluster cannot contain Trusted IRIX nodes.
- Ensure that the appropriate IRIX software is installed on the potential metadata server nodes. For example, if you want to use `quotas` and access control lists (ACLs) on any cluster node, the `oe.sw.quotas`, `nfs.sw.acl_nfs`, and `oe.sw.acl` subsystems must be installed on the IRIX nodes listed as potential metadata servers. For more information, see *IRIX Admin: Disks and Filesystems*, *IRIX Admin: Backup, Security and Accounting*, and your site's IRIX system administrator.

Also see "Requirements Specific to Solaris", page 21, and "Requirements Specific to Windows", page 54.

- Use an odd number of weighted IRIX nodes.
- Weight potential metadata servers only (unless weighting a stable IRIX client gives you an odd number).
- Shut down cluster services before maintenance.
- Avoid relocation and recovery in the current release.
- Use the following good practices:
 - Unmount the filesystems from the metadata server, shut down the node, and remount the filesystem when possible.
 - If relocation is necessary, limit activity on a filesystem.
 - Do the following before shutting down a node:
 - Unmount filesystems.
 - Shut down cluster services.
- Do not run power management software, which may interfere with the CXFS cluster.
- Enable the *forced unmount* feature for CXFS filesystems, which is turned off by default. Normally, an unmount operation will fail if any process has an open file on the filesystem. However, a forced unmount allows the unmount to proceed regardless of whether the filesystem is still in use.

Many sites have found that enabling this feature improves the stability of their CXFS cluster, particularly in situations where the filesystem must be unmounted.

The method used to implement this feature is platform-specific:

- 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 Solaris nodes, a similar function is performed with the `fuser -k` command and a `umount -f` command.

- On Windows nodes, a forced unmount does not terminate user processes by default. For details, see "Forced Unmount on a Windows Node", page 60.

This feature is available on an IRIX node by the following CXFS GUI menu:

Tasks

- > **Filesystems**
- > **Unmount a Filesystem**

In the CXFS GUI, click the **Force** toggle in the **Unmount Filesystem** task.

You can also specify this feature using the `cmgr(1M)` commands to define the filesystem. For more information, see "Forced Unmount of CXFS Filesystems", page 93.

You must use `cmgr` from an IRIX node, and the GUI must be connected to an IRIX node.

For more information, see the *CXFS Version 2 Software Installation and Administration Guide*, the Solaris `fuser(1M)` man page, and the Solaris and IRIX `umount(1M)` man pages.

Overview of the Installation and Configuration Steps

This section provides an overview of the installation, verification, and configuration steps for each platform type:

- "Solaris Overview"
- "Windows Overview", page 11

Solaris Overview

Note: For additional details, see the CXFS MultiOS release notes for Solaris. If you run into problems, see Chapter 8, "Troubleshooting", page 95.

Following is the order of installation and configuration steps for a CXFS Solaris node:

1. Install the Solaris 8 operating system according to the directions in the Solaris documentation (if not already done).

2. Install and verify the SGI TP9400 or SGI TP9100 RAID. See Chapter 2, "SGI RAID Firmware", page 13.
3. Install and verify the Brocade Fibre Channel switch. See Chapter 3, "Brocade Fibre Channel Switch Verification", page 15.
4. Obtain and install the CXFS and (optional) XLV licenses. See Chapter 4, "Obtaining CXFS and XVM FLEXlm Licenses", page 19.
5. Install and verify the JNI host bus adapter (HBA). You will install the JNIC146x package, which provides software for the Fibre Channel card and driver. See "JNI Fibre Channel Host Bus Adapter Installation and Configuration", page 26.
6. Prepare the Solaris node, including adding a private network. See "Preinstallation Steps for Solaris", page 38.
7. Use the Solaris `pkgadd(1M)`, `pkgrm(1M)`, and `pkginfo(1)` commands as directed to install the `SGIcxf`s package, which provides the following:
 - The `/etc/init.d/cxf`s_cluster command and associated scripts in the `/etc/rc` directory for automatic startup and shutdown
 - Drivers required for CXFS (`xvm` and `cell`)
 - The CXFS module
 - The `cxf`s_client commandSee "Client Software Installation Steps for Solaris", page 45.
8. Create the I/O fencing file. See "Postinstallation Steps for Solaris: Creating the I/O Fencing File", page 48.
9. Configure the cluster to add the new Solaris node and mount filesystems. See Chapter 7, "Cluster Configuration", page 87.

Windows Overview

This information applies to both Windows NT and Windows 2000 nodes unless otherwise noted.

Note: For additional details, see the CXFS MultiOS release notes for Windows. If you run into problems, see Chapter 8, "Troubleshooting", page 95.

Following is the order of installation and configuration steps for a CXFS Windows node:

1. Install the Windows operating system according to the directions in the Windows documentation (if not already done).
2. Install Windows NT Service Pack 6 or Windows 2000 Service Pack 2 according to the directions in the Windows documentation (if not already done).
3. Install and verify the SGI TP9400 or SGI TP9100 RAID. See Chapter 2, "SGI RAID Firmware", page 13.
4. Install and verify the Brocade Fibre Channel switch. See Chapter 3, "Brocade Fibre Channel Switch Verification", page 15.
5. Obtain the CXFS and (optional) XLV licenses. See Chapter 4, "Obtaining CXFS and XVM FLEXlm Licenses", page 19.
6. Install and verify the QLogic host bus adapter (HBA) and driver. See "QLogic Fibre Channel Host Bus Adapter Installation for Windows", page 60.
7. Prepare the Windows node, including adding a private network. See "Preinstallation Steps for Windows", page 64.
8. Install the CXFS software. See "Client Software Installation Steps for Windows", page 69.
9. Perform post-installation configuration steps:
 - "Configuring the FLEXlm License for Windows", page 74
 - "Creating the Windows I/O Fencing File", page 75
 - "Performing User Configuration", page 77
 - "Creating a New Hardware Profile", page 78
10. Configure the cluster to add the new Windows node, restart the node, and mount filesystems. See Chapter 7, "Cluster Configuration", page 87.
11. Start CXFS services on the Windows node to see the mounted filesystems under the configured drive letter.

SGI RAID Firmware

The SGI TP9400 or SGI TP9100 RAID will be initially installed and configured by SGI personnel.

Required SGI RAID Firmware

This section describes the required RAID firmware for the SGI TP9400 and SGI TP9100.

Required SGI TP9400 RAID Firmware

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 Mojave code must be installed according to FCO 1056.
- If you have a 4884 unit, the Mojave code is installed by default.

Note: By default, the TP9400 supports 32 logical units (LUNs). If additional LUNs are required, you must obtain a separate software enabling key; this key will support a maximum of 128 LUNs. Contact your SGI sales representative for the SGI software partitioning key.

Required SGI TP9100 RAID Firmware

The TP9100 4.0 CD contains the required version 7.75 controller firmware and NVSRAM files for the 1-Gbit TP9100. The TP9100 5.0 CD contains the required version 8.29 firmware and NVSRAM files for the 2-Gbit TP9100.

RAID Firmware 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.

For More Information

The following documentation is used to install and verify the RAID:

- SGI TP9400:
 - *TPSSM7 RAID Software Concepts Guide for TP9400*
 - *SGI TP9400 RAID Owner's Guide*
 - *SGI TP9400 RAID Administration Guide*
- SGI TP9100:
 - *TPM Installation Instructions and User's Guide for TP9100*

Brocade Fibre Channel Switch Verification

In order to protect data integrity, Solaris nodes and Windows nodes must use the *I/O fencing* feature, which isolates a problem node so that it cannot access I/O devices and therefore cannot corrupt data in the shared CXFS filesystem. This feature can only be used with an SGI-sold and supported Brocade Fibre Channel switch; therefore, the Brocade switch is a required piece of hardware in a multiOS cluster.

The Brocade Fibre Channel switches will be initially installed and configured by SGI personnel. You can use the information in this chapter to verify the installation.

Required Brocade Fibre Channel Switch Firmware and License

This release supports SGI-sold and supported Brocade Silkorm 2400 (8-port), 2800 (16-port), 3200 (8-port, 2-Gbit), 3800 (16-port, 2-Gbit) switches. See the release notes for further information.

All Brocade switches contained within the SAN fabric must have the appropriate Brocade license key installed and must be upgraded to the firmware level v2.2.1c or later. (This release was tested with v2.2.1c.) If the current firmware level of the switches must be upgraded, please contact your local SGI service representative or customer support center.

For the Brocade 3800 switch, Firmware level 3.0 must be installed. The Brocade switch must be configured so that its Ethernet interface is accessible from all IRIX cluster nodes using `telnet`. Either the public or private heartbeat network can be used to do this.



Caution: The `telnet` port must be kept free in order for I/O fencing to succeed.

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
```

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:

```
brocade:admin> version
Kernel:      5.3.1
Fabric OS:   v2.2.1c                               <== Firmware Revision
Made on:     Mon Dec 18 11:39:26 PST 2000
Flash:       Mon Dec 18 11:40:01 PST 2000
BootProm:    Thu Jun 17 15:20:39 PDT 1999
```

Changing to 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 using either the CXFS GUI or the `cmgr(1M)` command.
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 the `cmgr(1M)` command to add/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 IRIX servers can correctly identify the Brocade ports used by all clients.
6. If connected to a different RAID device, restart the Windows NT node.

Consult the *CXFS Version 2 Software Installation and Administration Guide* for instructions to configure I/O fencing.

Obtaining CXFS and XVM FLEXlm Licenses

The software licensing used by CXFS is based on the FLEXlm product from GLOBETrotter Software, Inc. For all supported platforms, a FLEXlm license is required to use CXFS. Perform the procedures in this chapter to satisfy this requirement.

XVM provides a mirroring feature. If you want to access a mirrored volume from a given node in the cluster, you must install the "XVM Volume Plexing" FLEXlm license on that node. Only those nodes that will access the mirrored volume must be licensed. For information about purchasing this license, see your SGI sales representative.

Obtain the Host Information Required for the License

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. The method used to obtain this information is platform-specific.

Solaris Host Information

To obtain the host identifier and hostname of the system on which you will run CXFS, execute the following Solaris commands:

```
/usr/bin/hostid  
/usr/bin/hostname
```

For example:

```
# /usr/bin/hostid  
830dad77  
# /usr/bin/hostname  
cxfssun2
```

When you are asked for the license manager host identifier, provide this information. You must have a separate license for each host on which CXFS is installed.

Windows Host Information

FLEXlm requires that you supply the Ethernet (MAC) address in order to generate the FLEXlm license. This address is known as the *Physical Address* in Windows. You can obtain this information in one of the following ways:

- View the network adapter properties in the **Windows Control Panel**
- Open a command prompt window and run the following command:

```
C:\>ipconfig /all
```

If the machine has more than one network interface, you should use the Physical Address of the private network interface.

Note: Windows NT licenses cannot be used under Windows 2000 and vice versa. If you are upgrading a Windows node to Windows 2000, you must obtain a new license.

Obtaining and Install the Licenses

Along with your entitlement number, you will receive a URL to a key generation page. To obtain your permanent CXFS and XVM licenses, follow the instructions on the key generation page. After the required information is provided, a key will be generated and displayed on the webpage along with installation instructions.

See also "FLEXlm License Verification for Solaris", page 25, and "Configuring the FLEXlm License for Windows", page 74.

For More Information

For more information about licensing, see the following webpage:

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

If you cannot use the web key generation page, you can contact the SGI order desk at 800 800 4SGI (800 800 4744).

For more information on FLEXlm, you may order the *Flexible License Manager End User Manual* from GLOBEtrouter Software, Inc.

Solaris Platform

This chapter contains the following:

- "CXFS on Solaris"
- "FLEXIm License Verification for Solaris", page 25
- "JNI Fibre Channel Host Bus Adapter Installation and Configuration", page 26
- "Preinstallation Steps for Solaris", page 38
- "Client Software Installation Steps for Solaris", page 45
- "Postinstallation Steps for Solaris: Creating the I/O Fencing File", page 48
- "Manual CXFS Startup/Shutdown for Solaris", page 50
- "Software Maintenance for Solaris", page 50

CXFS on Solaris

This section contains the following:

- "Requirements Specific to Solaris"
- "CXFS Commands Installed on Solaris", page 22
- "Solaris Log Files", page 23
- "Solaris Limitations and Considerations", page 23
- "Access Control Lists and Solaris", page 24

Requirements Specific to Solaris

In addition to the items listed in "Requirements", page 5, using a Solaris node to support CXFS requires the following:

- Solaris 8 operating system.

- One to four JNI FibreStar FCE-6460-N (PCI) 2-Gbit Fibre Channel host bus adapters (HBAs).

Note: 1-Gbit HBAs and Sbus HBAs are not supported.

- One or more of the following Sun Microsystems hardware platform series:
 - Ultra Enterprise 250
 - Ultra Enterprise 450
 - Ultra Enterprise 4000
 - Ultra Enterprise 3000
 - Ultra Enterprise 5000
 - Ultra Enterprise 6000
 - Ultra Enterprise 10000

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 a UFS or NFS filesystem on top of a Solaris CXFS filesystem, this is not recommended.

CXFS Commands Installed on Solaris

The following commands are shipped as part of the CXFS Solaris package:

- `/usr/cxfs_cluster/bin/cxfs_client` (the CXFS client service)
- `/usr/cxfs_cluster/bin/cxfslicense`
- `/usr/cxfs_cluster/bin/xvmprobe`

These commands provide all of the services needed to include a Solaris node in a CXFS cluster. The `pkgadd(1M)` output lists all software added; see "Installation Overview", page 46.

For more information, see the `cxfs_client(1M)` and `xvmprobe(1M)` man pages.

Solaris Log Files

The `cxfs_client` command creates a `/var/log/cxfs_client` log file. (There is no `/var/cluster` log on Solaris nodes.) This log file is not rotated or truncated.

For information about the log files created on IRIX nodes, see the *CXFS Version 2 Software Installation and Administration Guide*.

Solaris Limitations and Considerations

CXFS for Solaris has the following limitations and considerations:

- For optimal performance, you should set the value of the Solaris system tunable parameter `maxphys` in the `/etc/system` file. Do the following:
 1. Make a backup copy of the `/etc/system` file.

Note: Exercise extreme caution in changing `/etc/system` and always make a backup copy.

2. Change the value of `maxphys` to `0x800000` (hexadecimal) in the `/etc/system` file.
3. Reboot the Solaris node. This causes the change to take effect.
4. Verify that the new value for `maxphys` is in effect by running the following command:

```
# echo "maxphys/X" | adb -k
      physmem 1f03f
      maxphys:
      maxphys:          800000
```

- CXFS Solaris nodes cannot view or edit user and group quotas because CXFS administration must be performed from an IRIX node. However, user and group quotas are enforced correctly by the metadata server.

To view or edit your quota information, you must log in to an IRIX cluster node and make any necessary changes. If you would like to provide a viewing

command such as `repquota`, you could construct a Solaris shell script similar to the following:

```
#!/bin/sh
#
# Where repquota lives on IRIX
repquota=/usr/etc/repquota
# The name of an IRIX node in the cluster
irixnode=cain
rsh $irixnode "$repquota $*"
exit
```

- The minimum block size supported is 2 KB, determined by a maximum of 4 extents per 8-KB page. (XFS uses a default block size of 4 KB unless overridden by an administrator to a different blocksize value, for example 2 KB or 8 KB.)
- All disk devices attached to JNI controllers must be for use only by CXFS disks; do not attach non-disk devices to any JNI controllers that are configured for CXFS use. This restriction is required because all disk devices on JNI controllers configured for CXFS make use of the whole disk volume, which must be conveyed to Solaris via modification in the JNI driver to the value returned by the `READ_CAPACITY SCSI` command.

Access Control Lists and Solaris

All CXFS files have UNIX mode bits (read, write, and execute) and optionally an access control list (ACL). For more information, see the `chmod(1)` and `setfacl(1)` man pages.

If you restore a CXFS file that had an ACL containing only owner-ACL entries (that is, owner/group/other/mask) from a Solaris node, upon restoration one of the following will happen:

- **When using `tar(1)`, `cpio(1)`, and Legato Networker:** The ACL will be lost because these tools behave "intelligently" by not calling `acl(2)` to set an ACL if the file has only owner/group/other/mask entries. These tools will only set the file mode. However, this does not present a change in functionality because an access permissions check on the mode and the ACL containing only owner entries will give the same result.

- **When using other backup/restore utilities:** A mask will be added to the ACL if the application calls `acl(2)` for every file.

A backup/restore utility that calls `acl(2)` to set an ACL for every file will result in a file being restored with four ACL entries (that is, owner/group/other/mask), even though it may have originally had only three (that is, owner/group/other). This is due to a requirement in `getfacl(1)` that it receive four ACL entries for the `GETACL` command to `acl(2)`. (If fewer than four entries are returned, `getfacl` will report an error).

Note: Solaris ACLs can have up to 1024 entries for a file. A directory can have 1024 entries as well as an additional 1024 entries for the default ACL. For CXFS running on IRIX, ACLs are limited to 25 entries for a file and a total of 50 for a directory (that is, the directory ACL plus the default ACL).

FLEXlm License Verification for Solaris

Use the `cxfslicense` command with the `-d` option to verify that the FLEXlm licenses have been installed properly.

If the CXFS license is properly installed, you will see the following:

```
# /usr/cxfs_cluster/bin/cxfslicense -d
CXFS license granted.
```

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

```
Cluster services: CXFS not properly licensed for this host. Run
    '/usr/cxfs_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 `SYSLOG` 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 CLCONF failed the CXFS license check. Use the
Mar  4 12:58:05 5B:typhoon-q32    '/usr/cxfs_cluster/bin/cxfslicense -d'
Mar  4 12:58:05 5B:typhoon-q32 command to diagnose the license problem.
```

JNI Fibre Channel Host Bus Adapter Installation and Configuration

This section provides an overview of the JNI HBA installation and verification for Solaris nodes:

- "Installing the JNI HBA"
- "Installing and Running the EZ Fibre Configuration GUI", page 28
- "Verifying the JNI HBA Installation", page 37

These procedures may be performed by you or by a qualified Sun service representative. You must be logged in as `root` to perform the steps listed in this section.

Installing the JNI HBA

You can use one to four JNI HBAs for CXFS per Sun machine. (Other HBAs may be present that are not shared with the CXFS cluster.)

To install the JNI HBA, perform the following steps. Additional details are provided in various chapters/sections of the *Installation Guide, FCE-6460 and FCE2-6460 PCI-to-Fibre Channel Host Bus Adapters (Solaris, Windows NT/2000, Novell, AIX, HP-UX, Mac OS, Linux) JNI FibreStar*, as noted.

1. Install the JNI host bus adapter (HBA) into the Solaris system. Perform the steps in the following chapter:
 - "Hardware Installation"
2. Bring the system back up using the steps listed in the following "Verifying" sections (the following represents the location of these sections in the manual):
 - "Unix Server DriverSuite"
 - "Solaris Driver"
 - "Verifying Hardware in OpenBoot PROM"
 - "Verifying Hardware in Solaris"

You will be required to perform a Solaris `boot -r` after installing hardware.



Caution: If you do not see the expected results, do not attempt to continue. Instead, diagnose why the card is not being seen by the hardware.

3. Install the latest JNI HBA driver software (JNIC146x) and Storage Networking Industry Association (SNIA) application programming interface package (JNISnia), according to the instructions in the following “Installing” section:

- “Unix Server DriverSuite”
 - “Solaris Driver”
 - “Installing the Software”

You can retrieve the driver and SNIA package from the following JNI website:

<http://www.jni.com/Drivers>

- a. Under **Locate Driver by Product**, click on **FCE-6460**.
- b. Under the **Solaris** section, left click **JNIC146x.pkg** and save as the following pathname:

```
/var/tmp/JNIC146x.pkg
```

Verify that the driver attached correctly to the HBA and that the package installed correctly by following the verification steps at the end of the section. Do not proceed until the verification succeeds.

4. Set the HBA to fabric mode:

- a. In the `/kernel/drv/jnic146x.conf` file, change the following lines:

```
# FcLoopEnabled=1;  
# FcFabricEnabled=0;
```

Delete the # character at the beginning of each line to uncomment it and change the values so that loop is disabled (0) and fabric is enabled (1). When you are done, the lines will appear as follows:

```
FcLoopEnabled=0;  
FcFabricEnabled=1;
```

- b. Reboot the Solaris node by entering the following command:

```
# init 6
```

Installing and Running the EZ Fibre Configuration GUI

After you have verified the installation of the HBA and the driver's attachment to it, you are ready to install and run the JNI EZ Fibre program. This graphical user interface (GUI) will modify the driver's configuration file, `/kernel/drv/jnic146x.conf`, so that it lists the worldwide node name (WWNN) and worldwide port name (WWPN) for the devices on your Fibre Channel.

For general installation information, see *Quick Installation Guide, Solaris, AIX and Windows JNI EZ Fibre*.

Do the following:

1. Install the GUI and change to the appropriate directory:
 - a. Download the latest EZ Fibre GUI from the JNI website:
 - i. Go to the following website:

```
http://www.jni.com/Drivers
```
 - ii. Under **Locate Driver by Product**, click on **FCE-6460**
 - iii. Under the **Solaris** section, left click **EZF_22j.tar** or later and save as the following pathname (for example):

```
/var/tmp/EZF_22j.tar
```

- b. Extract the saved file using the `tar(1)` command. For example:

```
# tar xvf /var/tmp/EZF_22j.tar
```

- c. Change to the directory where the extracted GUI command is located:

```
# cd /var/tmp/EZF_22
```

- d. Run the `install.sh` script:

```
# ./install.sh
```

```
Checking for required and recommended patches...
```

```
checkpatches.sh: Note - The following OS vendor recommended patches are  
not installed or have been superseded -- please consult the EZ-Fibre  
read me:
```

```
108434-01 108435-01
<press enter to continue>
```

```
InstallAnywhere is preparing to install...
Installer using temporary disk space in '/tmp' ($TMPDIR not set).
```

You should install the GUI into the default location. When you see **Congratulations!**, click on **Done**.

2. Change to the following directory and read any README files you find there:

```
# cd /opt/jni/ezfibre/standalone
```

3. Invoke the EZ Fibre GUI from within the same directory by entering the following command:

```
# ./ezf
```

Two windows will appear. The first, titled **Refresh SAN information**, will say **Discovering LUNs for HBA#0**.

After a short while, this window will be replaced by a larger window, as shown in the example in Figure 5-1. (The example screen snaps in this document may not exactly match the display you see.)

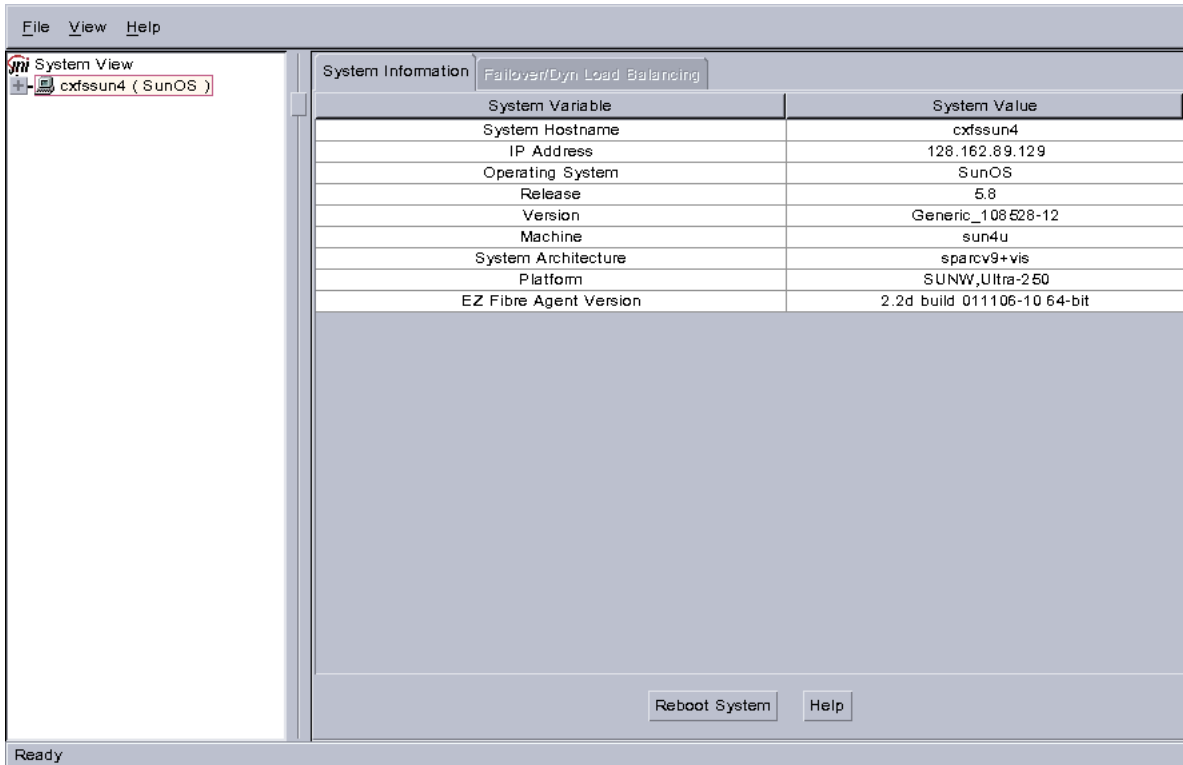


Figure 5-1 Example: Second Window: EZ Fibre Configuration Utility - Standalone

The left-hand pane of this window displays a listing of systems. Find the system you are configuring and click on the + sign next to it; this action expands the display so that it shows the installed JNI HBA on the system. Figure 5-2 highlights the + sign.

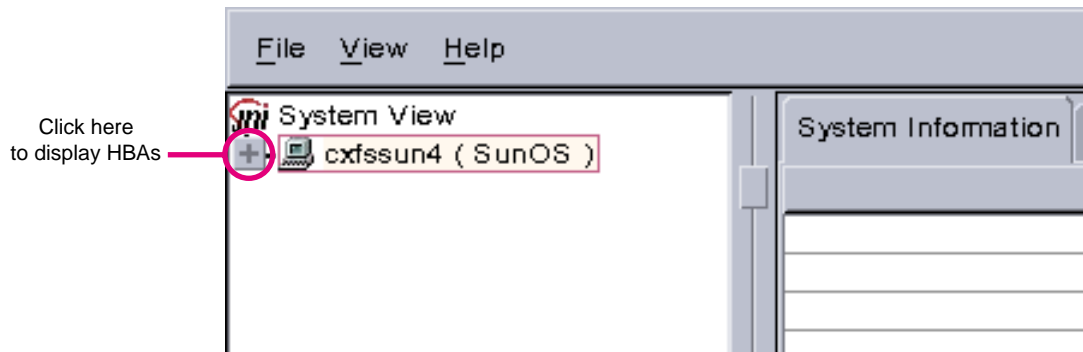


Figure 5-2 Location of icon (+) to Display the HBA

Figure 5-3 shows an example of the display after clicking on the + sign for `cxfsun4`, which shows the JNI HBA.




Figure 5-3 Example: After Clicking + to Display the HBA

4. Click on the icon to the right (not the + sign to the left). Figure 5-4 shows the icon.

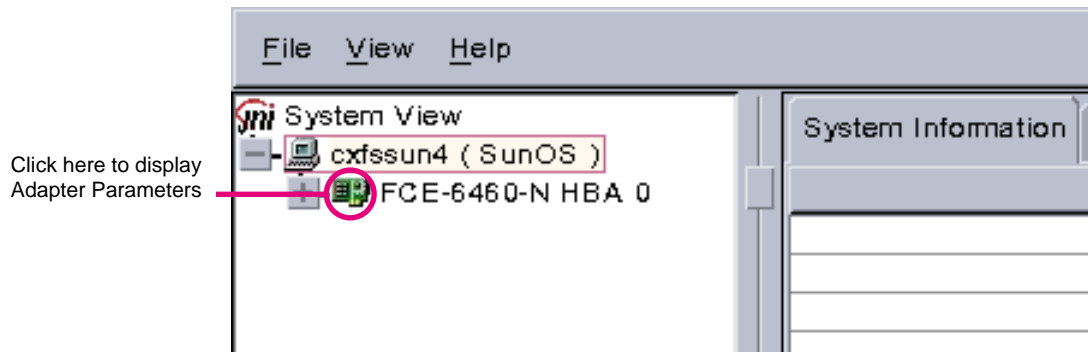


Figure 5-4 Location of the Icon to Display the Adapter Parameters

The right-hand pane will change to show **Adapter Parameters** for the selected HBA, as shown in Figure 5-5.

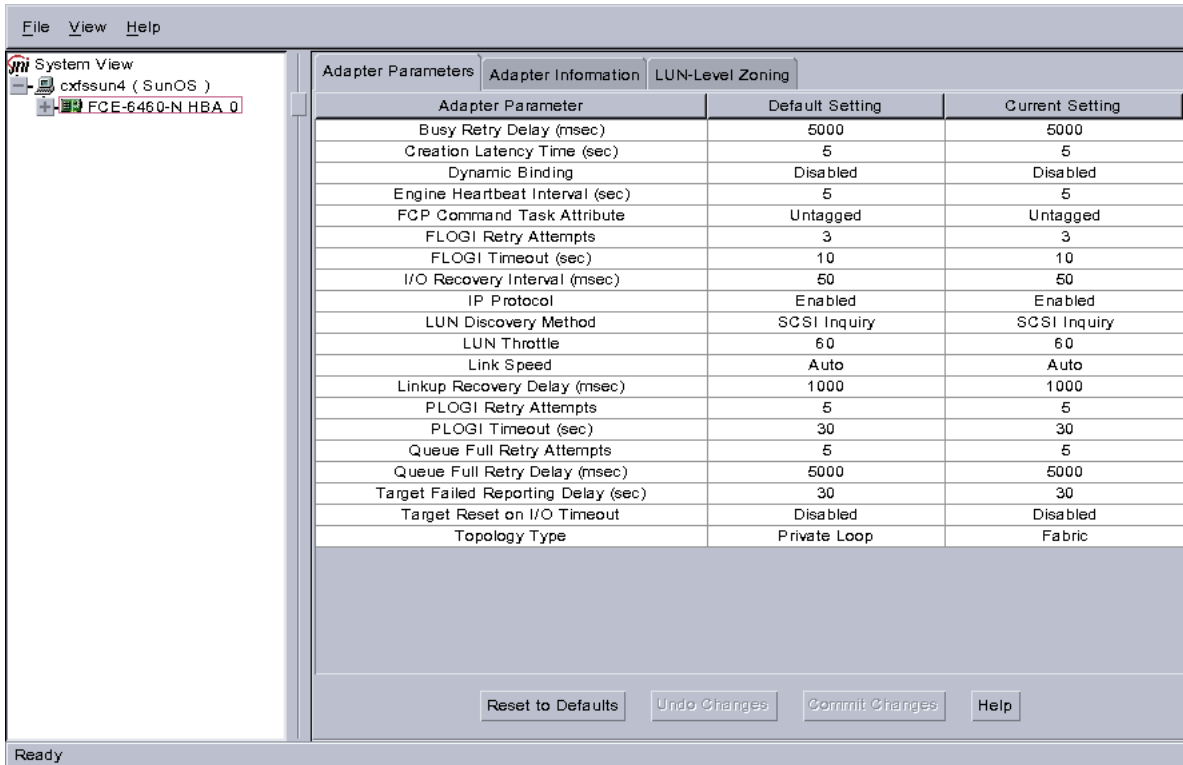


Figure 5-5 Example: After Clicking the HBA Icon to Show the Adapter Parameters

- a. Click on the **Adapter Information** tab to see the information in Figure 5-6.

The last two lines show the WWNN and WWPN of the JNI HBA. You will need the WWPN numbers when you configure the `/etc/fencing.conf` file, so take note of them now; for more information about the `/etc/fencing.conf` file, see "Postinstallation Steps for Solaris: Creating the I/O Fencing File", page 48.

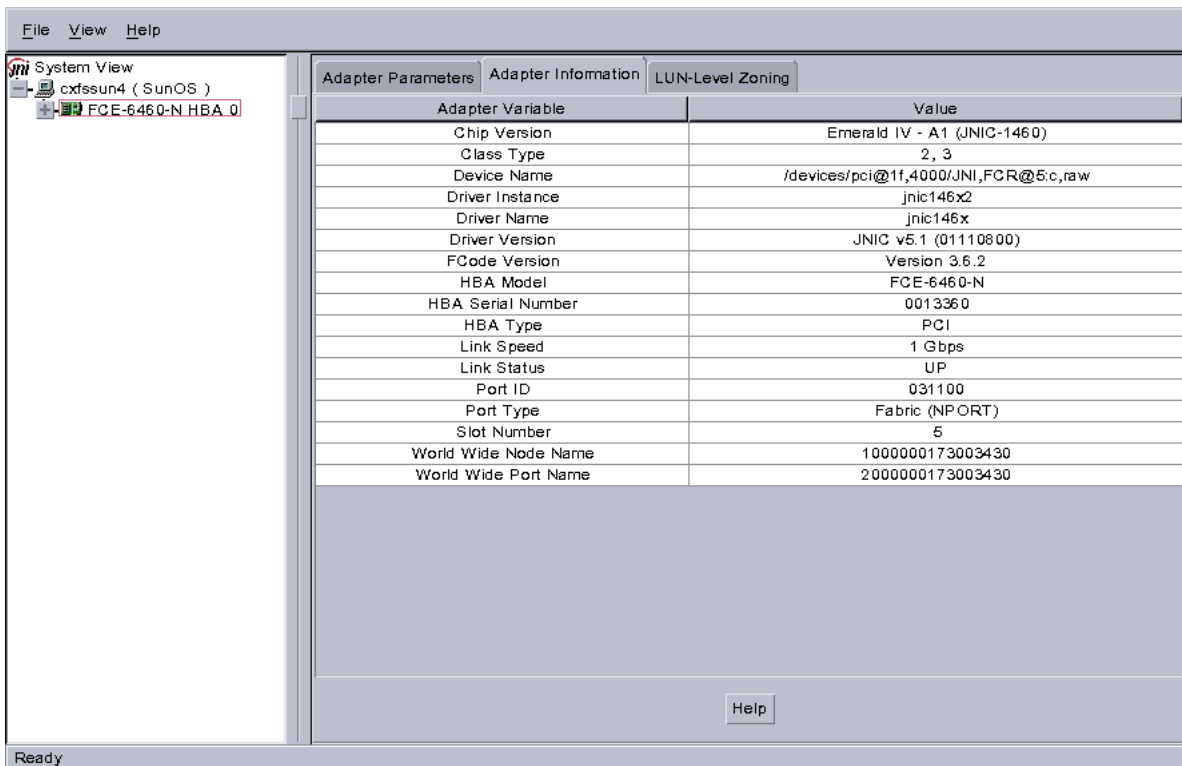


Figure 5-6 After Clicking the Adapter Information Tab

- b. Click on the **LUN-Level Zoning** tab in the left-hand pane to display a list of all the known devices on the selected HBA, as shown in Figure 5-7.




Figure 5-7 After Clicking on LUN-Level Zoning

5. Select the devices that should be accessed through the HBA.

For each device you want to access, click on the corresponding box in the **Mapped** column to make a check mark appear, as shown in Figure 5-8. After you have selected all the desired devices for the HBA, click on **Commit Changes**. The LUNs you map will depend upon your own site’s needs.



Caution: In this example, LUN 31 is used for administration by the TP9400. This LUN must not be used for other purposes; do not map it or use it for XVM volumes.

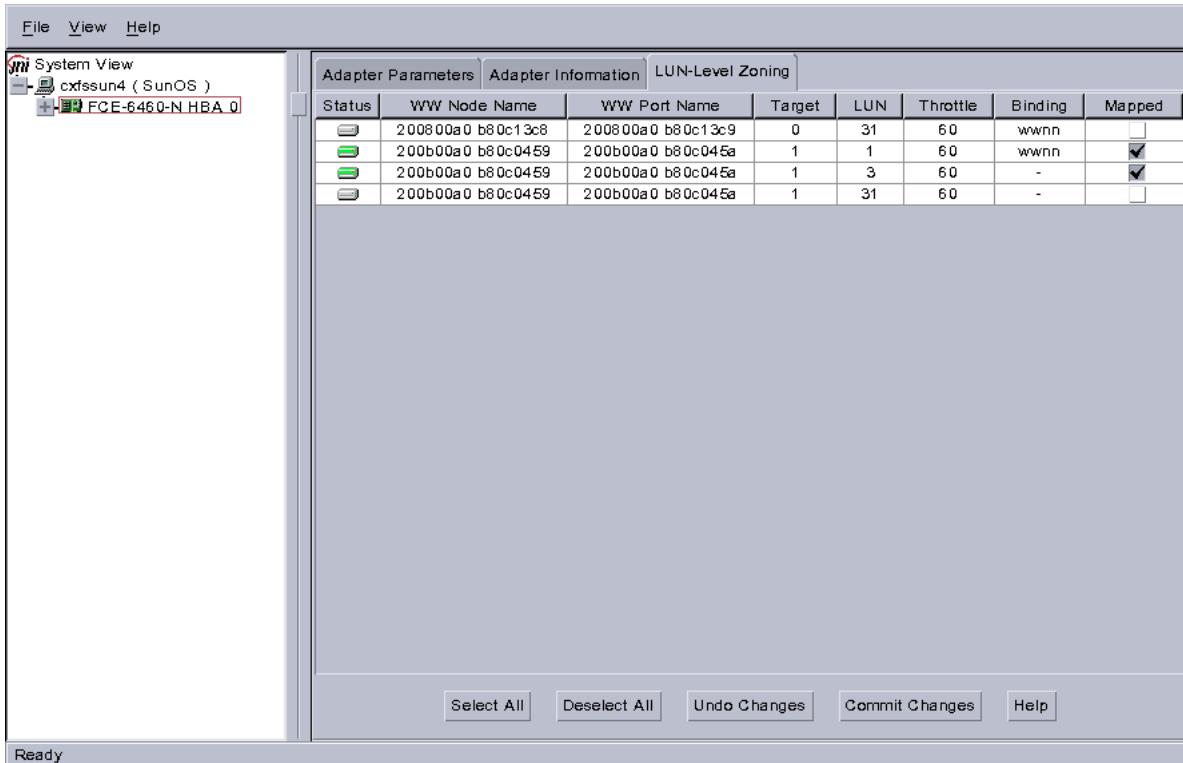


Figure 5-8 Example: After Mapping the LUNs and Committing the Changes

6. Reboot the system to make the changes take effect:

```
# init 6
```

Verifying the JNI HBA Installation

After the system reboots, you should verify that the devices were correctly configured by running the Solaris `format(1M)` command. You should see a list of each device you selected. For example:

```
# format
Searching for disks... done

c4t1d1: configured with capacity of 133.99GB
```

c4t1d3: configured with capacity of 133.99GB

AVAILABLE DISK SELECTIONS:

- 0. c0t0d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
/pci@1f,4000/scsi@3/sd@0,0
- 1. c4t1d1 <SGI-TP9400-0401 cyl 65533 alt 2 hd 64 sec 67>
/pci@1f,4000/JNI,FCR@5/sd@1,1
- 2. c4t1d3 <GI-TP9400-0401 cyl 65533 alt 2 hd 64 sec 67>
/pci@1f,4000/JNI,FCR@5/sd@1,3

Specify disk (enter its number):

In this example, disks 1 and 2 are being addressed by the JNI driver, as indicated by the presence of JNI,FCR in the pathname.

The system log and console display may display warning messages because the disks have IRIX labels on them. For example:

```
Mar 5 14:17:33 cxfssun4 scsi: WARNING: /pci@1f,4000/JNI,FCR@5/sd@1,1 (sd154):
Mar 5 14:17:33 cxfssun4          corrupt label - wrong magic number
Mar 5 14:17:33 cxfssun4 scsi:      Vendor 'SGI', product 'TP9400', 284203008 512 byte blocks
Mar 5 14:17:33 cxfssun4 scsi:      WARNING: /pci@1f,4000/JNI,FCR@5/sd@1,3 (sd155):
Mar 5 14:17:33 cxfssun4          corrupt label - wrong magic number
Mar 5 14:17:33 cxfssun4 scsi:      Vendor 'SGI', product 'TP9400', 284203008 512 byte blocks
```

This situation will be corrected automatically by CXFS after it is installed.

Note: You should not be alarmed by the preceding messages, nor should you try to relabel the disks with the `format` command. At this point, you are only trying to achieve connectivity to the devices, and the content is not important.

If you are having trouble with the verification steps, see "Common Solaris Problems: JNI Problems", page 101.

Preinstallation Steps for Solaris

When you install the CXFS software on the client-only node, 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 or a qualified Sun service representative will perform on your Solaris nodes prior to installing the CXFS software. It contains the following sections:

- "Hostname Resolution and Network Configuration Rules for Solaris"
- "Adding a Private Network for Solaris Nodes", page 40
- "Verifying the Private and Public Networks for Solaris", page 45

Hostname Resolution and Network Configuration Rules for Solaris



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

The following hostname resolution rules and recommendations apply to client-only nodes:

- Hostnames cannot begin with an underscore (_) or include any whitespace characters.
- The private network IP addresses on a running node in the cluster cannot be changed while cluster 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 for that network.

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

Adding a Private Network for Solaris Nodes

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

Note: A private network is **required** for use with CXFS. Only the private network is used by CXFS for heartbeat/control messages.

You may skip some steps, depending upon the starting conditions at your site. For details about any of these steps, see the Solaris documentation.

1. If your system is already operational and on the network, skip to step 2.

If your Solaris system has **never** been set up, bring the system to single-user mode. For example, go to the PROM prompt and boot the Solaris node into single-user mode:

```
> boot -s
```

As a last resort, you can reach the PROM prompt by pressing the L1-A (or Stop-A) key sequence.

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

The `/etc/inet/hosts` (or `/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/inet/hosts` (or `/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/inet/hosts` (or `/etc/hosts`) file.

The decision to use fully qualified domain names is usually a matter of how the clients (such as NFS) are going to resolve names for their client server programs, 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/inet/hosts` (or `/etc/hosts`) on all nodes. For example:

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

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

For more information, see the `hosts(4)`, `named(1M)`, and `nis(1)` 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.

For more information, see "Hostname Resolution and Network Configuration Rules for Solaris", page 39.

3. 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.)

4. Determine the name of the private interface by using the `ifconfig(1M)` command as follows:

```
# ifconfig -a
```

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.

For example, on an Ultra Enterprise 250, the integrated Ethernet is hme0; this is the public network. The following `ifconfig` output shows that only the public interface exists:

```
# ifconfig -a
lo0: flags=1000849<UP,LOOPBACK,RUNNING,MULTICAST,IPv4> mtu 8232 index 1
    inet 127.0.0.1 netmask ff000000
hme0: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 index 2
    inet 128.162.2.91 netmask ffffffff broadcast 128.162.2.255
    ether 8:0:20:d2:29:c5
```

If the second network does not appear, do the following:

- a. If you do not have the PCI card installed, install it. Refer to your PCI documentation for instructions.

If your card is already installed, skip to step b.

- b. Use the output from the `dmesg(1M)` command to determine the interface name for the private network; look for the network interface that immediately follows the public network; you may wish to search for `Found`. For example:

```
# dmesg
Feb  6 09:38:36 ue250 last message repeated 42 times
Feb  6 11:38:40 ue250 pseudo: [ID 129642 kern.info] pseudo-device: devinfo0
Feb  6 11:38:40 ue250 genunix: [ID 936769 kern.info] devinfo0 is /pseudo/devinfo@0
Feb  6 11:38:41 ue250 hme: [ID 517527 kern.info] SUNW,hme0 : PCI IO 2.0 (Rev Id = c1) Found
Feb  6 11:38:41 ue250 genunix: [ID 936769 kern.info] hme0 is /pci@1f,4000/network@1,1
Feb  6 11:38:41 ue250 hme: [ID 517527 kern.info] SUNW,hme1 : PCI IO 2.0 (Rev Id = c1) Found
Feb  6 11:38:41 ue250 hme: [ID 517527 kern.info] SUNW,hme1 : Local Ethernet address = 8:0:20:cc:43:48
Feb  6 11:38:41 ue250 pcipsy: [ID 370704 kern.info] PCI-device: SUNW,hme@1,1, hme1
Feb  6 11:38:41 ue250 genunix: [ID 936769 kern.info] hme1 is /pci@1f,2000/SUNW,hme@1,1
```

The second network is hme1; this is the private network, and is displayed after hme0 in the `dmesg` output. In this example, hme1 is the value needed in step c and in step 5 below.

- c. Initialize the private network's interface by using the `ifconfig(1M)` command as follows, where *interface* is the value determined in step b:

```
ifconfig interface plumb
```

For example:

```
# ifconfig hme1 plumb
```

After performing the `plumb`, the `hme1` interface will appear in the `ifconfig` output, although it will not contain the appropriate information (the correct information will be discovered after the system is rebooted later in step 8). For example, at this stage you would see the following:

```
ifconfig -a
lo0: flags=1000849<UP,LOOPBACK,RUNNING,MULTICAST,IPv4> mtu 8232 index 1
    inet 127.0.0.1 netmask ff000000
hme0: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 index 2
    inet 128.162.2.91 netmask ffffffff broadcast 128.162.2.255
    ether 8:0:20:d2:29:c5
hme1: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 index 3
    inet 0.0.0.0 netmask ff000000 broadcast 255.0.0.0
    ether 8:0:20:d2:29:c5
```

5. Create a file named `/etc/hostname.interface`, where *interface* is the value determined in step 4. This file must contain the name of the **private** network. For example:

```
# cat /etc/hostname.hme1
cxfssun3-priv
```

Note: In this scenario, `/etc/hostname.hme0` must contain the same value as the `/etc/nodename` file. For example:

```
# cat /etc/hostname.hme0
cxfssun3
# cat /etc/nodename
cxfssun3
```

The Solaris `/etc/nodename` file is analogous to the IRIX `/etc/sys_id` file.

6. Edit the `/etc/netmasks` file to include the appropriate entries.
7. (Optional) Edit the `/.rhosts` file if you want to use remote access or if you want to use the connectivity diagnostics provided with CXFS.

Ensure that the `/.rhosts` file on each Solaris node allows all of the nodes in the cluster to have access to each other. The connectivity tests execute a `ping(1)`

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(1) as user root.

For example, suppose you have a cluster with three nodes: irix0, sun1, and sun2. The /.rhosts files could be as follows (the prompt denotes the node name):

```
irix0# cat /.rhosts
sun1 root
sun1-priv root
sun2 root
sun2-priv root
```

```
sun1# cat /.rhosts
irix0 root
irix0-priv root
sun2 root
sun2-priv root
```

```
sun2# cat /.rhosts
irix0 root
irix0-priv root
sun1 root
sun1-priv root
```

8. Reboot the Solaris system:

```
# init 6
```

At this point, ifconfig will show the correct information for the private network.

For example:

```
ifconfig -a
lo0: flags=1000849<UP,LOOPBACK,RUNNING,MULTICAST,IPv4> mtu 8232 index 1
    inet 127.0.0.1 netmask ff000000
hme0: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 index 2
    inet 128.162.2.91 netmask ffffffff broadcast 128.162.2.255
    ether 8:0:20:d2:29:c5
hme1: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 index 3
    inet 10.1.1.36 netmask ffffffff broadcast 10.1.1.255
    ether 8:0:20:d2:29:c5
```

Verifying the Private and Public Networks for Solaris

For each private network on each Solaris node in the pool, verify access with the Solaris `ping(1)` command. Enter the following, where *nodeIPAddress* is the IP address of the node:

```
# /usr/sbin/ping -s -c 3 nodeIPAddress
```

For example:

```
# /usr/sbin/ping -s -c 3 128.162.89.128
PING 128.162.89.128: 56 data bytes
64 bytes from cxfssun3.americas.sgi.com (128.162.89.128): icmp_seq=0. time=0. ms
64 bytes from cxfssun3.americas.sgi.com (128.162.89.128): icmp_seq=1. time=0. ms
64 bytes from cxfssun3.americas.sgi.com (128.162.89.128): icmp_seq=2. time=0. ms
64 bytes from cxfssun3.americas.sgi.com (128.162.89.128): icmp_seq=3. time=0. 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:

```
# /usr/sbin/ifconfig eri0
eri0: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 index 2
    inet 128.162.89.127 netmask ffffffff broadcast 128.162.89.255
    ether 0:3:ba:d:ad:77
```

In the first 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.

Client Software Installation Steps for Solaris

The CXFS software will be initially installed and configured by SGI personnel. This section provides an overview of those procedures. You can use the information in this section to verify the installation.

Installation Overview

Installing the CXFS client CD for Solaris requires approximately 20 MB of space.

To install the required software on a Solaris node, SGI personnel will do the following:

1. Verify that the node has been upgraded to Solaris 8 (also known as SunOS 5.8) according to the *Solaris 8 Advanced Installation Guide*. Use the following command to display the currently installed system:

```
# uname -r
```

This command should return a value of 5.8.

2. Do the following:
 - a. Insert the CXFS MultiOS CD-ROM.
 - b. Read the already inserted CD-ROM as follows:

```
# pkgadd -d /cdrom/cdrom01/solaris/SGIcxfs.pkg
```

For example, where ellipses (. . .) indicate truncated output:

```
# pkgadd -d /cdrom/cdrom01/solaris/SGIcxfs.pkg
```

The following packages are available:

```
1  SGIcxfs      SGI CXFS client software
      (sparc) release 2.2
```

Select package(s) you wish to process (or 'all' to process all packages). (default: all) [?,??,q]:

Processing package instance <SGIcxfs> from </cdrom/solaris/SGIcxfs.pkg>

. . .

This package contains scripts which will be executed with super-user permission during the process of installing this package.

Do you want to continue with the installation of <SGIcxfs> [y,n,?] **y**

Installing SGI CXFS client software as <SGIcxfs>

```
## Installing part 1 of 1.
```



```

/etc/init.d/cxfs_cluster
/etc/rc0.d/K77cxfs_cluster <symbolic link>
/etc/rc2.d/S77cxfs_cluster <symbolic link>
/usr/cxfs_cluster/bin/cxfs_client
/usr/cxfs_cluster/bin/cxfs_soclosed
/usr/cxfs_cluster/bin/cxfslicense
/usr/cxfs_cluster/bin/xvmprobe
/usr/cxfs_cluster/doc/relnotes_2.2
/usr/kernel/drv/cell.conf
/usr/kernel/drv/sparcv9/cell
/usr/kernel/drv/sparcv9/xvm
/usr/kernel/drv/xvm.conf
/usr/kernel/fs/sparcv9/cxfs
/usr/sbin/clmount
/usr/share/man/man1/xvmprobe.1
/usr/share/man/man1m/cxfs_client.1m
[ verifying class <none> ]
## Executing postinstall script.
Starting CXFS services...

```

Installation of <SGIcxfs> successful.

- c. Verify that the CXFS license key has been installed. See Chapter 4, "Obtaining CXFS and XVM FLEXlm Licenses", page 19.

For example:

```

# /usr/cxfs_cluster/bin/cxfslicense -d
CXFS license granted.

```

Verifying the Solaris Installation

To verify that the CXFS software has been installed properly, use the `pkginfo(1M)` command as follows:

```
pkginfo -l SGIcxfs
```

For example, the following output indicates that the CXFS package installed properly:

```

% pkginfo -l SGIcxfs
  PKGINST: SGIcxfs
    NAME: SGI CXFS MultiOS client software
  CATEGORY: system
    ARCH: sparc

```

```
VERSION:  release 2.2
BASEDIR:  /
VENDOR:   Silicon Graphics Inc.
```

Postinstallation Steps for Solaris: Creating the I/O Fencing File

Note: Solaris nodes use the *I/O fencing* feature, which isolates a problem node so that it cannot access I/O devices and therefore cannot corrupt data in the shared CXFS filesystem. Solaris nodes do not have reset lines and therefore require I/O fencing to protect data integrity. I/O fencing can only be used with an SGI-sold and supported Brocade Fibre Channel switch; therefore, the Brocade switch is a required piece of hardware in a multiOS cluster.

To use I/O fencing, you must create the Solaris `/etc/fencing.conf` file, which enumerates the worldwide port name for all of the JNI host bus adapters (HBA) that will be used to mount a CXFS filesystem. There must be a line for the JNI HBA worldwide port name (WWPN) as a 64-bit hexadecimal number.

You must update the `/etc/fencing.conf` file whenever the JNI HBA configuration changes, including the replacement of a JNI HBA.

To determine the HBA WWPN, you must first set up the Brocade Fibre Channel switch and JNI HBA according to the directions in Chapter 3, "Brocade Fibre Channel Switch Verification", page 15, and "JNI Fibre Channel Host Bus Adapter Installation and Configuration", page 26.

At this point, you might be able to determine the HBA WWPN by running the EZ Fibre Configuration GUI: see "Installing and Running the EZ Fibre Configuration GUI", page 28, and Figure 5-6, page 35. If so, and you are **completely certain** that you can determine the correct WWPN of the HBA (and **not** that of any of the SAN targets), you can enter this value in the `/etc/fencing.conf` file.

Note: The WWPN is that of the JNI HBA itself, **not** any of the devices that are visible to that HBA in the fabric.

If you are **not** completely certain which number you should use, do the following:

1. Follow the Fibre Channel cable on the back of the Solaris host to determine the port to which it is connected in the Brocade Fibre Channel switch. Ports are

numbered beginning with 0. (For example, if there are 8 ports, they will be numbered 0 through 7.)

2. Use the `telnet(1)` command to connect to the Brocade Fibre Channel switch and log in as user `admin` (the password is `password` by default).
3. 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 `/etc/fencing.conf` file).

4. Edit or create the `/etc/fencing.conf` file and add the WWPN for the port determined in step 1. (Comment lines begin with #.)

For example, if you determined that port 0 is the port connected to the Brocade Fibre Channel switch, your `/etc/fencing.conf` file should appear as follows:

```
# WWPN of the JNI HBA installed on this system
#
2000000173002c0b
```

5. After the Solaris node is added to the cluster (see Chapter 7, "Cluster Configuration", page 87), enable the fencing feature by using the CXFS GUI or `cmgr` command on an IRIX node; for more information, see the *CXFS Version 2 Software Installation and Administration Guide*.

Manual CXFS Startup/Shutdown for Solaris

The `/etc/init.d/cxfs_cluster` script will be invoked automatically during normal system startup and shutdown procedures. This script starts and stops the processes required to run CXFS.

To start up CXFS processes manually on your Solaris node, enter the following:

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

To stop CXFS processes manually, enter the following:

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

Software Maintenance for Solaris

This section contains the following:

- "Upgrading the CXFS for Solaris Software"
- "Modifying the CXFS for Solaris Software", page 51

Upgrading the CXFS for Solaris Software

Before upgrading CXFS software, ensure that no applications on the node are accessing files on a CXFS filesystem. You can then run the new CXFS software package, which will automatically upgrade all CXFS software.

Modifying the CXFS for Solaris Software

You can modify the CXFS client service (`/usr/cxfs_cluster/bin/cxfs_client`) by placing options in the `/etc/cxfs_cluster/bin/cxfs_cluster.options` file. The available options are documented in the `cxfs_client(1M)` 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.

Windows NT and Windows 2000 Platforms

The information in this chapter applies to both Windows NT and Windows 2000 nodes unless otherwise noted. This chapter contains the following:

- "CXFS on Windows"
- "QLogic Fibre Channel Host Bus Adapter Installation for Windows", page 60
- "Preinstallation Steps for Windows", page 64
- "Client Software Installation Steps for Windows", page 69
- "Postinstallation Steps for Windows", page 73
- "Manual CXFS Startup/Shutdown for Windows", page 81
- "Software Maintenance for Windows", page 82

CXFS on Windows

This section contains the following:

- "Requirements Specific to Windows", page 54
- "CXFS Commands Installed on Windows", page 54
- "Windows Log Files", page 55
- "Functional and Performance Limitations Specific to Windows", page 55

Requirements Specific to Windows

In addition to the items listed in "Requirements", page 5, using a Windows node to support CXFS requires the insertion of a Windows host with at least the following:

- An Intel Pentium or compatible processor.
- 128 MB of RAM (more will improve performance)
- A minimum of 10 MB of free disk space.
- A QLogic 2200 or QLogic 2310 host bus adapter.
- The following QLogic software from the <http://www.qlogic.com> website:
 - QLogic 2200/2310 Driver version 8.1.5.12 or later
 - QLogic 2200 BIOS version 1.76 or later, or QLogic 2310 BIOS version 1.17 or later
 - QLogic SanSurfer (SANblade NT Client and SANblade Manager) 2.0.15 or later
- If two QLogic HBAs are installed, you should also install the QLdirect Filter (8.1.3 or later) in order to facilitate HBA failover and load balancing. If two different model HBAs are installed, you must install drivers for both models.
- Windows NT 4.0 Service Pack 6 or Windows 2000 Service Pack 2.

CXFS Commands Installed on Windows

A single CXFS service and a single CXFS filesystem driver are installed as part of the Windows installation. The service and the CXFS filesystem driver can be configured to run automatically when the first user logs into the node.

The command `C:\Program Files\CXFS\cxfs_license` is installed to assist with license validation; see "Configuring the FLEXlm License for Windows", page 74".

The command `C:\Program Files\CXFS\cxfs_info` is installed to display in a human-readable format the current state of the node in the cluster; see "Verifying the Cluster", page 92.

Windows Log Files

The Windows node will log important events in the system event log. You can view these events by selecting the following:

- For Windows NT:

```
Start
  > Programs
    > Administrative Tools
      > Event Viewer
```

- For Windows 2000:

```
Start
  > Settings
    > Control Panel
      > Administrative Tools
        > Event Viewer
```

For information about the log files created on IRIX nodes, see the *CXFS Version 2 Software Installation and Administration Guide*. The CXFS Client service will also log important information to the following file:

```
C:\Program Files\CXFS\log\cxfs_client.log
```

Functional and Performance Limitations Specific to Windows

There are a number of unique limitations in the CXFS for Windows software when accessing and manipulating files on a CXFS filesystem. These limitations are described from both a UNIX and a Windows perspective in the following sections.

UNIX Perspective of CXFS on a Windows Node

This section describes the differences and limitations of a CXFS filesystem on a Windows node from a UNIX perspective:

- Windows nodes can support multiple CXFS filesystems mounted under a single drive letter. Only one CXFS drive letter may be configured on a Windows node.

The top-level file structure under the CXFS drive letter consists of an in-memory directory structure that mimics the mount points on the IRIX server. The CXFS software creates these directories before mounting the CXFS filesystems. This file

structure supports only creating and deleting directories; there is no support for creating and deleting regular files, renaming directories, and so on. Attempts to perform unsupported actions will generally result in an invalid parameter error. You can perform normal filesystem operations on files and directories beneath the mount points.

- A Windows node can support regular files, directories, and links. However, it does not support other XFS file types.
- Symbolic links cannot be distinguished from normal files or directories on a Windows node. Opening a symbolic link will open the target of the link, or will report `file not found` if it is a dangling link.

You can move, rename, or delete a symbolic link; however, you cannot copy a symbolic link. Copying a valid symbolic link will result in copying the file or directory that the link refers to, rather than the normal UNIX behavior that copies the link itself.

- Setting the read-only flag on a file will remove all (user, group, and other) write mode bits, while unsetting the flag will add the write mode bit for the owner only.
- A user's UNIX user ID (UID) is determined by mapping the user's Windows username to a username in the `passwd` file, as described in "User Administration for CXFS", page 4. The user's UNIX group ID (GID) is the default GID in the `passwd` file for the mapped UID and is not based on a Windows group mapped to a UNIX group name.
- Access controls are enforced on the metadata server using the mapped UID and GID of the user attempting to access the file. Therefore, a user can expect the same access on a Windows node as any other node in the cluster when mounting the same filesystem.
- Access control lists (ACLs) that are set on any files or directories are also enforced as they would be on any other node in the cluster. A new file will inherit the parent directory default ACL, if one is defined.
- The file permissions displayed in Windows Explorer is derived from a Windows ACL generated from the file's IRIX ACL if it has an ACL, or the file's mode bits. The "other" mode bits on a file are represented by the `Everybody` group and the IRIX ACL mask is represented by the `CXFS ACL Mask` group. See the `chacl(1)` man page for an explanation of IRIX ACLs.
- Only the owner of the file may change the file's permissions.

- Setting the read-only flag on a file will remove all (user, group, and other) write mode bits, while unsetting the flag will add the write mode bit for the owner only.
- When changing and applying a Windows ACL to a CXFS file, a new IRIX ACL is applied to the file if the file already had an ACL or if the new IRIX ACL contains user and groups other than the file's owner and group. Otherwise, only the file's mode bits are updated.
- CXFS Windows nodes cannot use the `edquota(1)` or `repquota(1)` commands to view or edit user and group quotas because CXFS administration must be performed on an IRIX node. However, user and group quotas will be enforced correctly, independent of the node in the CXFS cluster where files are created and deleted.

Windows Perspective of CXFS on a Windows Node

This section describes the differences and limitations of a CXFS filesystem on a Windows node in comparison to other Windows filesystems from a Windows perspective:

- The read-only flag is the only file attribute supported by CXFS filesystems on a Windows node; other flags like `hidden` are ignored. Only the owner will be permitted to change the read-only flag on a file or directory. Other users, including the Administrator, will receive an `access denied` error message.
- CXFS filesystems should not be shared to another Windows host from a CXFS Windows node because there is no support for the opportunistic locking that is used by Windows to guarantee data coherency. If this is required, it is recommended that an IRIX server shares the filesystem via Samba.
- Avoid using duplicate filenames in the same directory that vary only in case. CXFS is case sensitive, but some Windows applications may not maintain the case of all filenames, which may result in unexpected behavior.
- Avoid using completely uppercase 8.3 filenames. If you use completely uppercase 8.3 filenames, some applications (including Windows Explorer) may incorrectly assume that only 8.3 filenames are supported by the filesystem and will not preserve case.
- Take care when using Disk Manager and other similar Microsoft and third-party disk tools. These tools assume particular disk formats (FAT or NTFS) and may even write to the super block of a Fibre Channel disk and write a Windows signature on the disk. This will corrupt the XVM volume information on the disk.



Caution: Writing a Windows signature to the disk used for CXFS will corrupt the volume and filesystem structure on that disk, potentially resulting in the loss of data on the disk.

Install the CXFS software components onto a NTFS partition rather than a FAT partition. The security of the following files cannot be guaranteed if these files are installed onto a FAT filesystem:

- C:\Program Files\CXFS\passwd
 - C:\Program Files\CXFS\group
 - C:\Program Files\CXFS\fencing.conf
- There is no recycle bin; deleted files will be permanently deleted.
 - There is no automatic notification of directory changes performed by other nodes in the cluster. Applications (such as Windows Explorer) will not automatically update their display if another node adds or removes files from the directory currently displayed.
 - CXFS software does not export 8.3 alternative filenames. Older Windows applications that only support 8.3 filenames may be unable to open files with longer filenames.
 - A CXFS filesystem cannot be used as the boot partition of a Windows host.
 - User and group names on the Windows node must match the equivalent user and group names on other nodes in the cluster that run operating systems based on UNIX. The Administrator user is automatically mapped to root on the UNIX node, and the Administrators group is mapped to the UNIX sys group.
 - The ownership of a file may not be changed on a Windows node because a Windows user takes ownership of a file, rather than the owner giving ownership to another user (which is supported by the UNIX access controls).
 - If the file permissions include an entry for CXFS ACL Mask, then the effective permissions of all entries except the owner and the Everybody group will be the intersection of the listed permissions for that user or group and the mask permissions. Therefore, changing the CXFS ACL Mask permissions will set the

maximum permissions other listed users and groups may have. Their access may be further constrained in the specific entries for those users and groups.

- If the users and groups listed in a file's permissions (whether mode bits and/or ACLs) cannot be mapped to users and groups on the Windows node, attempts to display the file permissions in a file properties window will fail with an unknown user or group error.
- Only the read, write, and execute bits in a Windows ACL will be set, if appropriate, when displaying the file permissions; the other bits will not be set, and attempts to apply those bits will be ignored. The only exception to this is that the owner of the file will have the "change permissions" and "take ownership" bits set.
- If the users and/or groups in a Windows ACL being applied to a CXFS file cannot be mapped to a UNIX UID or GID, the application of that ACL will fail with an unknown user or group error.
- The volume properties window for the CXFS drive letter will display the total capacity of all mounted filesystems and the largest free space on any one of those filesystems.
- The alignment block size for direct I/O will be the largest block size of the filesystems that are mounted, which may be larger than the block size of the filesystem in use.

Performance Considerations on a CXFS Windows Node

The following are performance considerations on a CXFS Windows node:

- If you open the Windows Explorer Properties window on a directory, it will attempt to traverse the filesystem in order to count the number and size of all subdirectories and files; this action is the equivalent of running the UNIX `du(1)` command. This can be an expensive operation, especially if performed on directories between the drive letter and the mount points, because it will traverse all mounted filesystems.
- Virus scanners, Microsoft Find Fast, and similar tools that traverse a filesystem are very expensive on a CXFS filesystem. Such tools should be configured so that they do not automatically traverse the CXFS drive letter.
- The mapping from Windows user and group names to UNIX identifiers occurs as the CXFS software starts up. In a Windows domain environment, this process can take a number of seconds per user for usernames that do not have accounts within the domain. If the `/etc/passwd` file contains a number of such users, you should

remove users who do not have accounts on the Windows nodes from the `passwd` file that is installed on the Windows nodes.

This issue has less impact on Windows nodes in a workgroup than on those in a domain because the usernames can be quickly resolved on the node itself, rather than across the network to the domain controller.

Forced Unmount on a Windows Node

SGI recommends that you enable the forced unmount feature on CXFS filesystems (see "Recommendations", page 8, and "Forced Unmount of CXFS Filesystems", page 93).

On UNIX machines, a forced unmount using the `-k` option to the `umount(1M)` command causes all processes that have open files on the specified filesystem to be unconditionally killed. However, on Windows nodes, a forced unmount has a slightly different effect.

When issued on a Windows node, a forced unmount does not terminate user processes by default. Instead, all open filehandles that refer to the filesystem enter a special "zombie" state that blocks any further I/O to the filesystem. The filesystem is then unmounted out from underneath the running processes. Processes continue to run normally with the exception that any further I/O operations on the filesystem using a stale file handle will be failed with a `STATUS_INVALID_HANDLE` error. File handles invalidated by a forced unmount will remain invalid indefinitely. If the filesystem is later remounted, a process must open a fresh filehandle before it can resume I/O operations to the filesystem.

There is one exception to this behavior. If a file is memory-mapped by any process at the time of the forced unmount, then it, and every other process that has an open filehandle to that file, will be unconditionally killed. This is necessary for memory mapped files on Windows nodes as it is the only way to free the reference to the filesystem and to guarantee that no stale file data remains in memory after the unmount.

QLogic Fibre Channel Host Bus Adapter Installation for Windows

The QLogic Fibre Channel host bus adapter (HBA) should be installed according to the QLogic hardware and driver installation instructions.

Information regarding large logical unit (LUN) support under Windows can be found in the QLogic documentation and also in Microsoft's support database:

<http://support.microsoft.com/default.aspx?scid=kb;en-us;Q310072>

<http://support.microsoft.com/default.aspx?scid=kb;en-us;Q245637>

This section discusses the following:

- "Confirming the QLogic HBA Installation"
- "Support for More than 8 LUNs under Windows NT ", page 61
- "Configuring Two HBAs for Failover Operation", page 63

Confirming the QLogic HBA Installation

To confirm that the QLogic HBA and driver are correctly installed, select the following to display all of the logical units (LUNs) attached to the Fibre Channel switch:

```
Start
  > Settings
      > Control Panel
          > SCSI Adapters
              > QLogic QLA2200
                  > Bus 0
```

Support for More than 8 LUNs under Windows NT

The Qlogic HBA will require further configuration to support more than 8 LUNs under Windows NT. This is performed by using the SANsurfer software that is distributed with the QLogic hardware or that can be downloaded from the following QLogic website:

http://www.qlogic.com/support/drivers_software.asp

The installation instructions for SANsurfer are also located on the QLogic website.

Upon successful installation of the HBA hardware and software, start the SANsurfer manager application. Connect to the local host and change the LUNs per target to 0 by selecting the following:

Select NVRAM section
> Advanced NVRAM Settings
> LUNS per Target

Save this configuration and then reboot when prompted.

To confirm that the QLogic HBA and driver are correctly installed after the reboot, check the SCSI registry by selecting the following:

Start
> Run
> regedit

Select the following key:

`HKEY_LOCAL_MACHINE\HARDWARE\DEVICEMAP\Scsi\Scsi Port x\Scsi Bus 0\Target Id 0\Logical Unit Id x-xxx`

You should see all of the LUNs available in the RAID device, where `Scsi Port x` represents the `QL2xxx` driver.

If you have fewer than 72 LUNs, you should also see the HBA and LUNs in the SCSI adapters list by selecting the following:

Start
> Settings
> Control Panel
> SCSI Adapters
> QLogic QLA2xxx
> Bus 0

The limitation of the SCSI adapters list with a large number of LUNs is further described in the following:

<http://support.microsoft.com/default.aspx?scid=kb;en-us;Q190834>

If you are unable to see the SCSI devices in the registry, check the following, as described in "Windows QLogic Problems", page 102:

- The HBA is properly seated in the host
- Cables are connected correctly
- The node is not fenced

Configuring Two HBAs for Failover Operation

Note: This procedure assumes that the CXFS driver is already installed and working properly with one host bus adapter (HBA).

To configure two HBAs for failover operation under Windows, do the following:

1. Install the QLdirect driver v8.01.03 (NT4) by following all the default settings for the installation and verify that the CXFS client still operates normally.
2. Disable fencing for this node. You can do this using the CXFS GUI or the `cmgr(1M)` command.
3. Determine the world wide port name (WWPN) of the current adapter:
 - a. Install SANsurfer Qlogic SANblade NT client v2.0.15.
 - b. Install SANsurfer Qlogic SANblade Manager v2.0.15.
 - c. Run SANsurfer to determine the WWPN.
 - d. Record the WWPN on paper.
4. Shut down Windows.
5. Install the second HBA and start Windows.
6. If the second HBA is a different model from the original one, install its mini port driver (for example, `q12300.sys`).
7. Start the Qlogic SANblade Manager and verify that two HBAs are detected. Verify that both of them mirror the same devices and logical units (LUNs). Notice that both HBAs have the same world wide node name (WWNN) but different WWPNs. The original HBA can be recognized by its WWPN recorded in step 3.
8. Make the NVRAM settings (or at least LUNs per target) of the new HBA adapter the same as the original one.
9. Configure the HBA port (click on the **Configure** button).

Note: Ignore the following message, which appears when HBA/LAN configuration is done for the first time (line breaks added here for readability):

An invalid device and LUN configuration has been detected. Auto configure run automatically.

Click on **OK** to continue.

The HBA0 devices are automatically set to be visible for Windows application (notice the open eye) and HBA1 devices are set to be invisible (notice the closed eye).

10. Select the first device in the table, right click, and then select **Configure LUN(s)**. In the new window, select the following:

Tools

- > **Load Balance**
- > **All LUNs**

This will statically distribute the LANs traffic load associated with this device between the two HBAs.

Repeat this step for each of the other HBA devices.

11. Click on **Apply** to save the new configuration.
12. If fencing is used, add the new HBA WWPN to the `fencing.conf` file and update the switch port information using the CXFS GUI or the `cmgr(1M)` command. Enable fencing if required.
13. Reboot Windows.

For more information about using the CXFS GUI or the `cmgr(1M)` command to perform these tasks, see *CXFS Version 2 Software Installation and Administration Guide*.

Preinstallation Steps for Windows

When you install the CXFS software on the client-only node, 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 or a qualified Windows service representative will perform on your Windows nodes prior to installing the CXFS software. It contains the following:

- "Hostname Resolution and Network Configuration Rules for Windows", page 65
- "Adding a Private Network for Windows Nodes", page 65
- "Adding a Private Network for Windows 2000 Nodes", page 67
- "Verifying the Private and Public Networks for Windows", page 68

Hostname Resolution and Network Configuration Rules for Windows



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

The following hostname resolution rules and recommendations apply to Windows nodes:

- Hostnames cannot begin with an underscore (_) or include any whitespace characters.
- The private network IP addresses on a running node in the cluster cannot be changed while cluster 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 for that network.

Adding a Private Network for Windows Nodes

The steps to add a private network are platform-specific.

Adding a Private Network for Windows NT Nodes

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

Note: A private network is **required** for use with CXFS. Only the private network is used by CXFS for heartbeat/control messages.

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

1. Install the second network adapter in the Windows node as per the network adapter vendor instructions. In some cases you must remove all network setups, restart, and then add network services to each network adapter from scratch.
2. Ensure that the Windows NT node recognizes two network adapters in the system. Select the following:

Start
 > **Settings**
 > **Control Panel**
 > **Network**
 > **Adapters**

3. Specify the private network settings (IP address, subnet mask, default gateway) on one of the network adapters. Select the following:

Start
 > **Settings**
 > **Control Panel**
 > **Network**
 > **Protocols**
 > **TCP/IP Protocol**
 > **Properties**
 > **IP Address**
 > **Adapter**

Note: The private network IP address must be a fixed address and cannot be configured by DHCP.

4. Specify DNS settings as required in the **DNS** tab.

5. Select the software components to be installed. All components should be installed.
6. Skip the **WINS Address tab** (a WINS server is not required).
7. Ensure that **IP Forwarding** is *not* selected in the **Routing** tab.
8. In the **IP Address** tab, ensure that the other network adapter is configured with a different IP address and a different subnet; this second network is the public network for all other network traffic.

Note: The public network interface can be configured with DHCP.

Adding a Private Network for Windows 2000 Nodes

The following procedure provides an overview of the steps required to add a private network to the Windows 2000 node.

Note: A private network is **required** for use with CXFS. Only the private network is used by CXFS for heartbeat/control messages.

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

1. Install the second network adapter in the Windows node as per the network adapter vendor instructions. In some cases you must remove all network setups, restart, and then add network services to each network adapter from scratch.
2. Ensure that the Windows 2000 node recognizes two network adapters in the system. Select the following:

Start

> **Settings**
> **Network and Dial-up Connections**

3. Specify the private network settings (IP address, subnet mask, default gateway) on one of the network adapters. Select the following:

Start

> **Settings**
> **Network and Dial-up Connections**

Then right-mouse click on **Properties** and select the following:

**Internet Protocol (TCP/IP)
> Properties**

Note: The private network IP address must be a fixed address and cannot be configured by DHCP.

4. Specify the static IP address and DNS server.

Verifying the Private and Public Networks for Windows

You can confirm that the previous procedures to add private networks were performed correctly by using the `ipconfig` command in a DOS command shell. In the following example, the 10 network is the private network and the 192.168.63 network is the public network on a Windows system:

```
> ipconfig /all
Windows NT IP Configuration
    Host Name . . . . . : cxfs1
    DNS Servers . . . . . :
    Node Type . . . . . : Hybrid
    NetBIOS Scope ID. . . . . :
    IP Routing Enabled. . . . . : No
    WINS Proxy Enabled. . . . . : No
    NetBIOS Resolution Uses DNS : No

Ethernet adapter E190x1:
    Description . . . . . : 3Com EtherLink PCI
    Physical Address. . . . . : 00-01-03-46-2E-09
    DHCP Enabled. . . . . : No
    IP Address. . . . . : 10.0.0.201
    Subnet Mask . . . . . : 255.0.0.0
    Default Gateway . . . . . : 10.0.0.255

Ethernet adapter E190x2:
    Description . . . . . : 3Com EtherLink PCI
    Physical Address. . . . . : 00-B0-D0-31-22-7C
    DHCP Enabled. . . . . : No
    IP Address. . . . . : 192.168.63.201
    Subnet Mask . . . . . : 255.255.255.0
    Default Gateway . . . . . : 192.168.63.254
```

Primary WINS Server : 192.168.63.254

Client Software Installation Steps for Windows

The CXFS software will be initially installed and configured by SGI personnel. This section provides an overview of those procedures. You can use the information in this section to verify the installation.

Note: This procedure assumes that the CXFS software is installed under the default path `C:\Program Files\CXFS`. If a different path is selected, then that path should be used in its place in the following instructions.

To install the CXFS client software on a Windows node, do the following:

1. Log onto the Windows node as Administrator.
2. Verify that the node has been updated to the correct service pack:

- Windows NT 4.0 Service Pack 6:

```
Start
  > Programs
      > Administrative Tools
          > Windows NT Diagnostics
              > Version
```

- Windows 2000 Service Pack 2:

```
Start
  > Programs
      > Accessories
          > System Tools
              > System Information
```

3. Insert the CXFS for Windows CD-ROM into the Windows host. Normally the setup program will automatically run, otherwise run `winnt/setup.exe` from the CD-ROM.
4. Acknowledge the software license agreement when prompted. The release notes will be displayed automatically.

5. Install the CXFS software, as shown in Figure 6-1. If the software is to be installed in a nondefault directory, click on the **Browse** button to select another directory. Click on **Next** when finished.

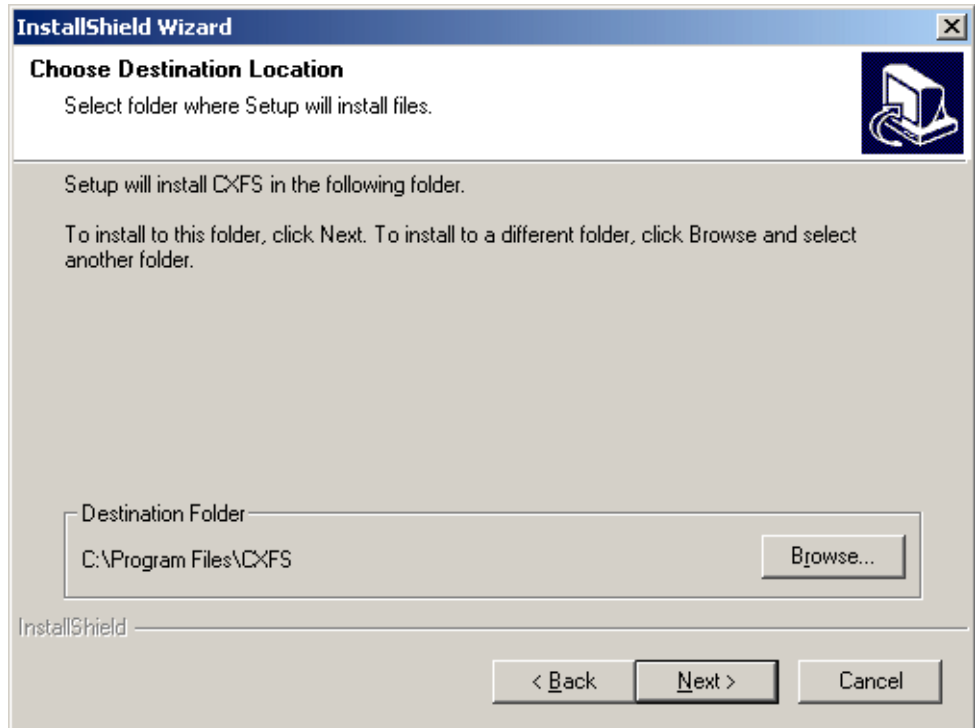


Figure 6-1 Select All Software Components

6. Enter details for the following fields as shown in Figure 6-2 and click the **Next** button when finished:
 - **Select drive letter for CXFS volumes to be mounted on:** specify the **drive letter** under which all CXFS filesystems will be mounted. You cannot select a drive letter that is currently in use.
 - **Specify location of fencing, UNIX /etc/passwd and /etc/group files:** specify the path where the configuration files will be installed. The default is the same location as the software under C:\Program Files\CXFS.

- **Select or enter the IP address of the heartbeat network adapter:** specify the IP address of the private network adapter on the Windows node.
- **Enter any additional command line arguments:** enter arguments that may be passed to the CXFS Client service. For most configurations, this can be left empty. See "Modifying the CXFS for Solaris Software", page 51.

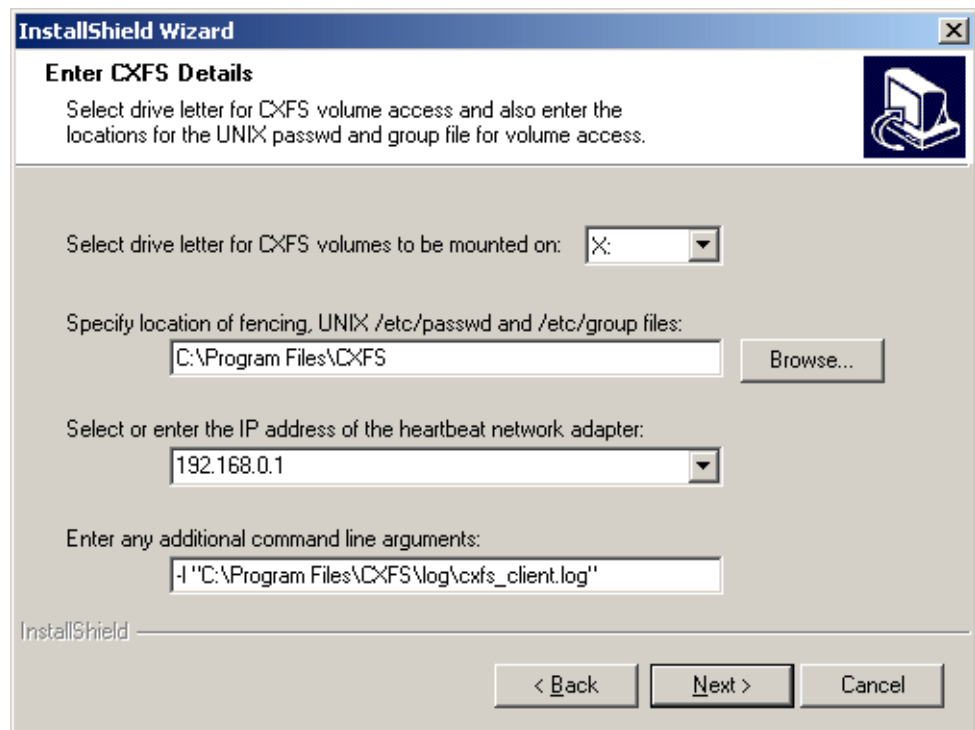


Figure 6-2 Enter CXFS Details

7. Review the settings, as shown in Figure 6-3. If they appear as you intended, click the **Next** button. If you need to make corrections, click the **Back** button.

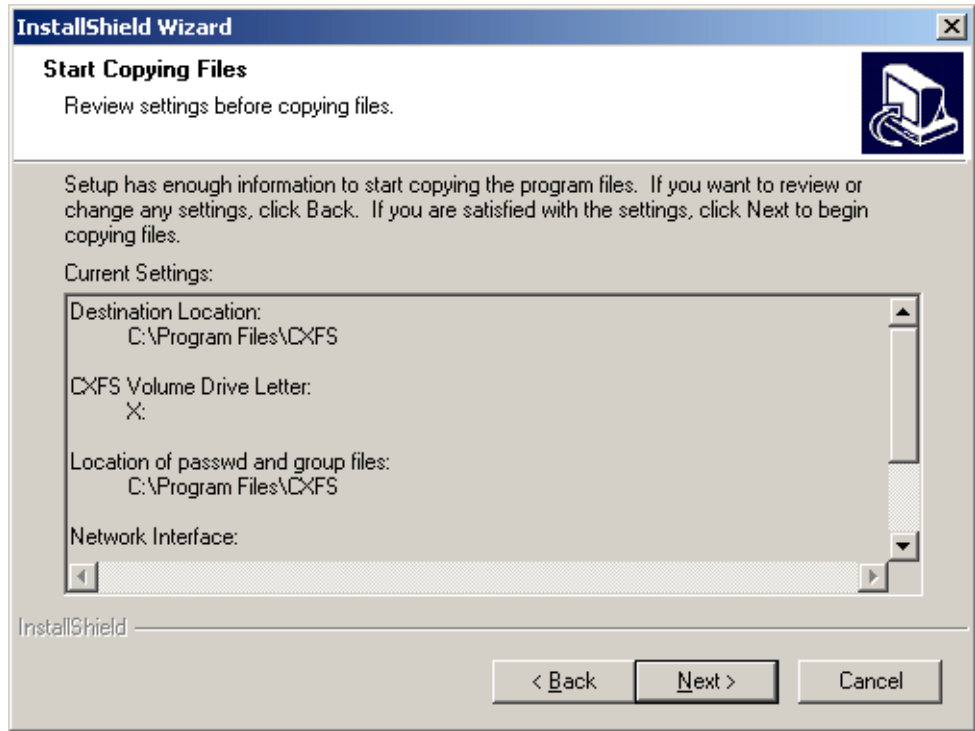


Figure 6-3 Review the Settings

After you click the **Next** button, the CXFS software will be installed.

8. You will be given the option to start the driver at system start-up or now, as show in Figure 6-4.

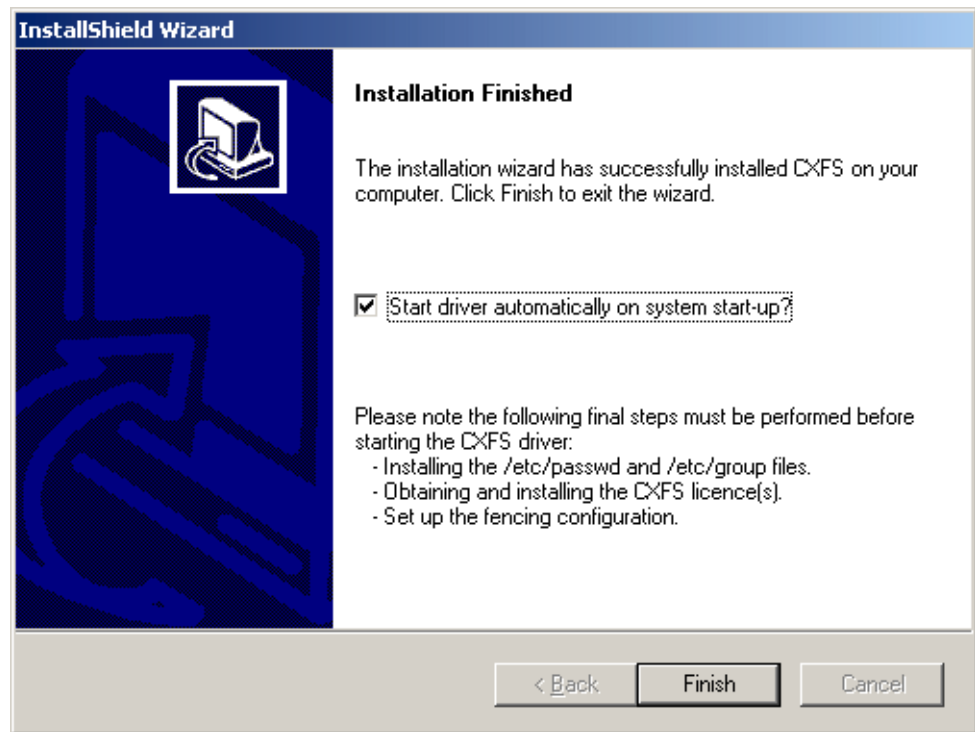


Figure 6-4 Start CXFS Driver

Because there are some important postinstallation steps, do not start the CXFS driver now. Choose **Start driver automatically on system start-up** and click the **Finish** button.

Postinstallation Steps for Windows

This section discusses the configuration steps that you should perform after installing CXFS software but before restarting a Windows node.

The following postinstallation steps are required to ensure the correct operation of the CXFS software:

- "Configuring the FLEXlm License for Windows"

- "Creating the Windows I/O Fencing File", page 75
- "Performing User Configuration", page 77
- "Checking Permissions on the Password, Group, and Fencing Files", page 78
- "Creating a New Hardware Profile", page 78

Configuring the FLEXlm License for Windows

Note: Windows NT licenses cannot be used under Windows 2000 and vice versa. If you are upgrading a Windows NT node to Windows 2000, you must obtain a new license.

You must configure a FLEXlm license before you restart the Windows node by following these steps:

1. Add the mandatory CXFS license and the optional XVM license to the following file:

```
C:\Program Files\CXFS\lib\license.dat
```

For more information, see Chapter 4, "Obtaining CXFS and XVM FLEXlm Licenses", page 19.

2. Validate these licenses by running the `cxfslicense` command in a DOS command shell.

Create a DOS command shell with the following sequence:

- Windows NT:
 - Start**
 - > **Programs**
 - > **Command Prompt**

- Windows 2000:

Start
 > Programs
 > Accessories
 > Command Prompt

To run `cxfslicense`, enter the following command:

```
C:\Program Files\CXFS\cxfslicense
```

If a valid license has been correctly specified, the following will be displayed:

```
Found valid license for feature CXFS_NT version 2.000  
The CPU count specified in the license is OK.
```

If the Windows node has the optional XVM mirroring license, you will also see the following:

```
Found valid license for feature XVM_NT version 3.000  
The CPU count specified in the license is OK.
```

Note: Licenses for Windows 2000 have the feature names `CXFS_W2K` and `XVM_W2K`.

Creating the Windows I/O Fencing File

Note: Windows nodes use the *I/O fencing* feature, which isolates a problem node so that it cannot access I/O devices and therefore cannot corrupt data in the shared CXFS filesystem. Windows nodes do not have reset lines and therefore require I/O fencing to protect data integrity. I/O fencing can only be used with an SGI-sold and supported Brocade Fibre Channel switch; therefore, the Brocade switch is a required piece of hardware in a multiOS cluster.

To use I/O fencing, you must create the Windows fencing file, which enumerates the worldwide port name for all of the QLogic host bus adapters (HBA) that will be used to mount a CXFS filesystem. The path to the I/O fencing file is as follows:

```
C:\Program Files\CXFS\fencing.conf
```

The I/O fencing file must contain a line for the QLogic HBA worldwide port name (WWPN) as a 64-bit hexadecimal number. You must update the fencing file whenever the QLogic HBA configuration changes, including the replacement of a QLogic HBA.

To determine the HBA WWPN, do the following:

1. Set up the Brocade Fibre Channel switch and QLogic HBA according to the directions in Chapter 3, "Brocade Fibre Channel Switch Verification", page 15, and "QLogic Fibre Channel Host Bus Adapter Installation for Windows", page 60.
2. Follow the Fibre Channel cable on the back of the Windows host to determine the port to which it is connected in the Brocade Fibre Channel 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(1)` command to connect to the Brocade Fibre Channel 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.

```
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 `C:\Program Files\CXFS\fencing.conf` file).

5. Edit the following file to add the WWPN for the port determined in step 2 (comment lines begin with #):

```
C:\Program Files\CXFS\fencing.conf
```

For example, if you determined that port 0 is the port connected to the Brocade Fibre Channel switch, the fencing file should appear as follows:

```
C\> type C:\Program Files\CXFS\fencing.conf
#
# WWPN of the JNI HBA installed on this system
#
2000000173002c0b
```

6. After the Windows node is added to the cluster (see Chapter 7, "Cluster Configuration", page 87), enable the fencing feature by using the CXFS GUI or `cmgr` command on an IRIX node; for more information, see the *CXFS Version 2 Software Installation and Administration Guide*.

Performing User Configuration

Note: If you do not install the `passwd` and `group` files properly, the CXFS software will treat all filesystem operations as user `nobody`.

After installing the CXFS software onto the Windows node, but before the CXFS node is restarted, you must install the `/etc/passwd` and `/etc/group` files to the location specified during installation, which defaults to the following:

- `/etc/passwd` as `C:\Program Files\CXFS\passwd`
- `/etc/group` as `C:\Program Files\CXFS\group`

Note: The default `passwd` and `group` files that are installed are invalid files containing comments; these invalid files will cause CXFS Client to generate warnings in its log file. You must remove the comments in these files when you install the `/etc/passwd` and `/etc/group` files.

Checking Permissions on the Password, Group, and Fencing Files

The permissions on the `fencing.conf`, `passwd`, and `group` files must restrict access so that only the system administrator can modify the files. This can be done by right-clicking on the file names in Windows Explorer and selecting the following:

- Windows NT:

Properties
 > **Security**
 > **Permissions**

- Windows 2000:

Properties
 > **Security**

Verify that the permissions are **Read for Everyone** and **Full Control for Administrators**.



Caution: Failure to set permissions on the `passwd` and `group` files would allow users to change their UID/GUI at will and even gain superuser access to the files on the CXFS filesystem. Failure to set permissions on the `fencing.conf` file would allow users to misconfigure and even disable I/O fencing, which could result in data corruption on the CXFS filesystems.

Creating a New Hardware Profile

It is strongly recommended that you create a new hardware profile and that you disable the CXFS software in the current hardware profile, in order to have a backup profile available. If the CXFS software causes the host to crash on startup, you can easily switch back to the original hardware profile and successfully return to the configuration before the CXFS software was installed.

To create a new hardware profile, right-click the **My Computer** icon and select the following:

- Windows NT:

Properties
 > **Hardware Profiles**
 > **Copy**

- Windows 2000:

Properties
 > **Hardware**
 > **Hardware Profiles**
 > **Copy**

This action copies the current hardware profile, most likely called the **Original Configuration (Windows NT)** or **Profile 1 (Windows 2000)**. You should call this new profile **CXFS Configuration** to distinguish it from other profiles. You can make the **CXFS Configuration** the default profile chosen on startup by selecting the up arrow button and moving the **CXFS Configuration** profile to the top of the list.

To remove the CXFS driver from the current hardware profile, which should be the original profile, select the following:

- Windows NT:

Start
 > **Settings**
 > **Control Panel**
 > **Devices**
 > **CXFS**
 > **HW Profile**
 > **Disable**

- Windows 2000:

Start
 > **Settings**
 > **Control Panel**
 > **Administrative Tools**
 > **Computer Management**
 > **System Tools**
 > **Device Manager**

To show non-plug-and-play devices, select the following:

View
 > **Show hidden devices**

To show the CXFS driver, select the following:

Non-Plug and Play Devices
 > **CXFS**
 > **Properties**
 > **Device Usage**
 > **Do not use this device in the current hardware profile**

You should also disable the CXFS Client service for the current profile by selecting the following:

- Windows NT:

Start
 > **Settings**
 > **Control Panel**
 > **Services**
 > **CXFS Client**
 > **HW Profile**
 > **Disable**

- Windows 2000:

```
Start
  > Settings
    > Control Panel
      > Administrative Tools
        > Services
          > CXFS Client
            > Properties
              > Log On
                > Disable
```

When the Windows host boots, you may choose **CXFS Configuration** to automatically start CXFS or choose the previous profile (most likely **Original Configuration**) to start without CXFS.

Manual CXFS Startup/Shutdown for Windows

The CXFS processes are automatically started when a Windows node is restarted. This behavior may be altered by changing the configuration of the CXFS filesystem driver and the CXFS Client service.

By default, the driver is configured to start manually and the Client service is configured to start automatically. Because the CXFS Client service depends on the CXFS filesystem driver, the driver will be started by the service.

It is recommended that the CXFS driver configuration remains manual.

You can change the CXFS Client service configuration to start manually, so that CXFS does not automatically start, by selecting the following:

- Windows NT:

```
Start
  > Settings
    > Control Panel
      > Services
```

- Windows 2000:

Start
 > **Settings**
 > **Control Panel**
 > **Administrative Tools**
 > **Services**

Change **CXFS Client** to manual rather than automatic. CXFS can then be started and stopped manually by the Administrator using the same selection sequence.

Software Maintenance for Windows

This section contains the following:

- "Modifying the CXFS for Windows Software", page 82
- "Upgrading the CXFS for Windows Software", page 84
- "Removing the CXFS for Windows Software", page 84
- "Downgrading the CXFS for Windows Software", page 85

Modifying the CXFS for Windows Software

To change the location of the software and other configuration settings that were requested in "Client Software Installation Steps for Windows", page 69, perform the following steps:

1. Select the following:

Start
 > **Settings**
 > **Control Panel**
 > **Add/Remove Programs**
 > **CXFS**
 > **Add/Remove**
 > **Modify**

Figure 6-5, page 83, shows the screen that lets you modify the software.

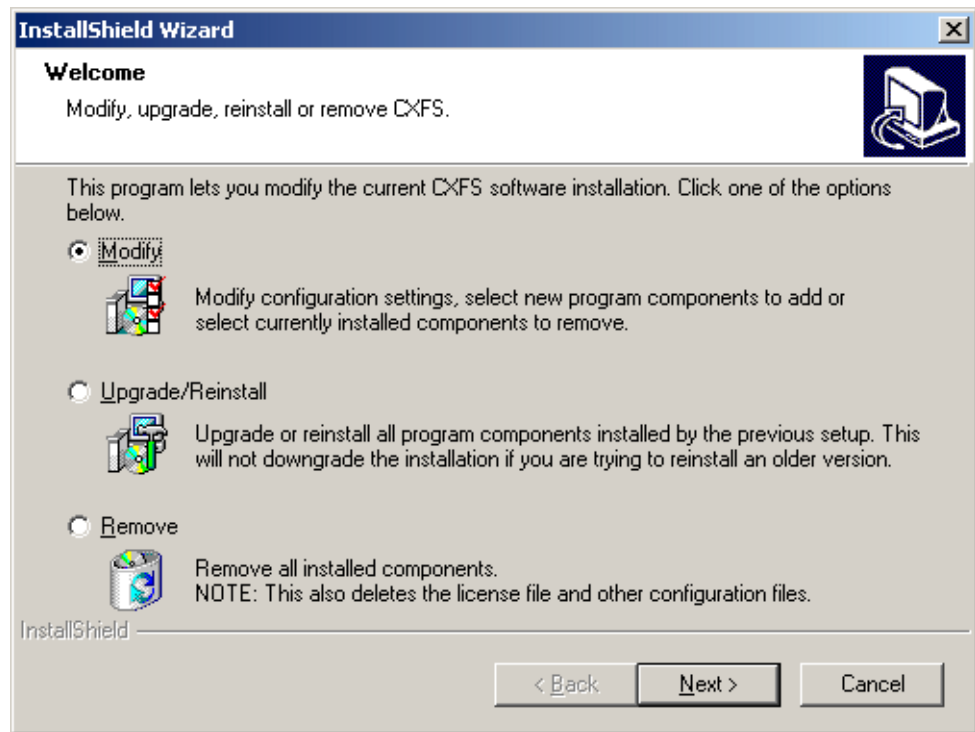


Figure 6-5 Modify the CXFS for Windows

2. Make the necessary configuration changes.

You can display the list of possible command line arguments supported by the CXFS Client service by running the service from a DOS command shell as follows:

```
> C:\Winnt\system32\cxfs_client.exe -h
```

3. Restart the Windows node, which causes the changes to take effect.

Upgrading the CXFS for Windows Software

To upgrade the CXFS for Windows software, perform the following steps:

1. Insert the CD-ROM containing the upgraded software to run the setup program. If the setup program does not automatically start, run `winnt/setup.exe` from the CD-ROM.
2. Select **Upgrade/Reinstall** and follow the prompts as discussed in "Client Software Installation Steps for Windows", page 69.

Figure 6-5, page 83, shows the screen that lets you modify the software.

3. Restart the Windows node. The upgraded software will not activate until the Windows node is restarted.

Removing the CXFS for Windows Software

To remove the CXFS for Windows software, first ensure that no applications on this host are accessing files on a CXFS filesystem. Then, select the following sequence to remove all installed files and registry entries:

```
Start
  > Settings
    > Control Panel
      > Add/Remove Programs
        > CXFS
          > Add/Remove
            > Remove
```

Figure 6-5, page 83, shows the screen that lets you remove the software.

Note: The `passwd`, `group`, and `fencing.conf` files will be removed.

You should then restart the Windows node. This will cause the changes to take effect.

Downgrading the CXFS for Windows Software

To downgrade the CXFS software, follow the instructions to remove the software in "Removing the CXFS for Windows Software", page 84, and then install the older version of the software as directed in "Client Software Installation Steps for Windows", page 69.

Note: The removal process may remove the configuration and license files. You should back up these files before removing the CXFS software so that you can easily restore them after installing the downgrade.

Cluster Configuration

This chapter provides an overview of the procedures to add the client-only nodes to an established IRIX cluster. It assumes that you already have an IRIX cluster installed and running with mounted filesystems. These procedures will be performed by you or by SGI service personnel.

All CXFS administrative tasks other than restarting the Windows node must be performed on an IRIX node, using either the CXFS GUI (invoked by the `cxtask` command) or the `cmgr(1M)` command. The GUI also provides a guided configuration for defining a cluster.

This section discusses the following tasks in cluster configuration:

- "Defining the Client-Only Nodes"
- "Adding the Client-Only Nodes to the Cluster", page 89
- "Defining the Switch for I/O Fencing", page 90
- "Starting CXFS Services on the Client-Only Nodes", page 90
- "Mounting Filesystems on the Client-Only Nodes", page 91
- "Restarting the Windows Node", page 91
- "Verifying the Cluster", page 92
- "Forced Unmount of CXFS Filesystems", page 93

For detailed configuration instructions, see the *CXFS Version 2 Software Installation and Administration Guide*.

Defining the Client-Only Nodes

To add a client-only node to a CXFS cluster, you must define it as a node in the pool. You can do this on an IRIX node using the CXFS GUI or `cmgr(1M)` command.

Do the following to determine the value for the hostname field in the GUI:

- Solaris: use the value displayed by `/etc/nodename`, which must match the node's primary hostname in the `/etc/inet/hosts` (or `/etc/hosts`) file; that is,

the first field after the node's IP address in `/etc/inet/hosts` (or `/etc/hosts`). This field can be either the hostname or the fully qualified domain name.

- Windows NT:

```
Start
  > Settings
      > Control Panel
          > Network
              > Identification
                  > Computer Name
```

- Windows 2000:

```
Start
  > Settings
      > Network and Dial-up Connections
          > Advanced
              > Network Identification
```

When you specify that a node is running an operating system other than IRIX, the node will automatically be given a weight of 0, and you cannot change it. (These nodes cannot be potential metadata servers, and always have a weight of 0 so that they are not counted when calculating the CXFS membership quorum.) For client-only nodes, you must specify a unique node ID.

For example, the following shows the entries used to define a Solaris node named `solaris1` using the `cmgr` command in prompting mode:

Note: The following example shows the `cmgr(1M)` prompts used with IRIX 6.5.17, in which `NT` should be selected for any Windows system. As of IRIX 6.5.18, the prompt will change as follows:

```
Operating System <IRIX|Solaris|Windows> ?
```

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

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|Solaris|NT> ? solaris
Node ID ? 7
Do you wish to define failure hierarchy[y/n]:y
Hierarchy option 0 <System|Fence|Shutdown>[optional] ? fence
Hierarchy option 1 <System|Fence|Shutdown>[optional] ? shutdown
Hierarchy option 2 <System|Fence|Shutdown>[optional] ?
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)
NIC 1 - (use network for control messages) true|false> ? true
NIC 1 - Priority <1,2,...> ? 1
```

For details about these commands, see the “Define a Node” sections of the GUI or `cmgr` reference chapters in the *CXFS Version 2 Software Installation and Administration Guide*.

Adding the Client-Only Nodes to the Cluster

After you define all of the client-only nodes, you must add them to the cluster using either the CXFS GUI or the `cmgr(1M)` command on an IRIX node.

For example, if you have already defined a cluster named `cxfscluster` and want to add the Solaris nodes `solaris1` and `solaris2`, you could use the following `cmgr` command:

```
cmgr> modify cluster cxfscluster

cxfscluster ? add node solaris1
cxfscluster ? add node solaris2
cxfscluster ? done
```

For details, see the “Modify a Cluster” sections of the GUI or `cmgr` reference chapters in the *CXFS Version 2 Software Installation and Administration Guide*.

Depending upon your filesystem configuration, you may also need to add the node to the list of clients that have access to the volume. See “Mounting Filesystems on the Client-Only Nodes”, page 91.

Defining the Switch for I/O Fencing

You are required to use I/O fencing on client-only nodes in order to protect data integrity. I/O fencing requires a Brocade switch. To define the switch for the cluster database, use either the CXFS GUI or the `cmgr(1M)` command on an IRIX node.

For example:

```
cmgr> define switch ptg-brocade username admin password password
```

Note: The masking feature applies only to IRIX nodes. (IRIX nodes automatically discover the available HBAs, and the masking feature is used to restrict the HBAs that are eligible for fencing.) For Solaris nodes and Windows nodes, the fencing file determines the ports that may be fenced. For more information, see "Postinstallation Steps for Solaris: Creating the I/O Fencing File", page 48, and "Creating the Windows I/O Fencing File", page 75.

For details, see the "Define a Switch" sections of the GUI or `cmgr` reference chapters in the *CXFS Version 2 Software Installation and Administration Guide*.

Starting CXFS Services on the Client-Only Nodes

After adding the client-only nodes to the cluster, you must start CXFS services on them. You can do this using either the CXFS GUI or the `cmgr(1M)` command on an IRIX node.

For example:

```
cmgr> start cx_services on node solaris1 for cluster cxfscluster
cmgr> start cx_services on node solaris2 for cluster cxfscluster
```

For details, see the "Start CXFS Services" sections of the GUI or `cmgr` reference chapters in the *CXFS Version 2 Software Installation and Administration Guide*.

Mounting Filesystems on the Client-Only Nodes

If you have specified that the filesystems are to be automatically mounted on any newly added nodes, then you do not need to specifically mount the filesystems on the new client-only nodes that you added to the cluster.

Otherwise, you can mount the filesystems on the new client-only nodes by unmounting the currently active filesystems, enabling the mount on the required nodes, and then performing the actual mount. You can do this using the GUI or the `cmgr(1M)` command on an IRIX node.

For example, to mount the `fs1` filesystem on all nodes in the cluster except `solaris2`, you could use the following commands:

```
cmgr> admin cxfs_unmount cxfs_filesystem fs1 in cluster cxfscluster
cmgr> modify cxfs_filesystem fs1 in cluster cxfscluster

cxfs_filesystem fs1 ? set dflt_local_status to enabled
cxfs_filesystem fs1 ? add disabled_node solaris2
cxfs_filesystem fs1 ? done
```

Note: SGI recommends that you enable the *forced unmount* feature for CXFS filesystems, which is turned off by default; see "Recommendations", page 8, and "Forced Unmount of CXFS Filesystems", page 93.

For details, see the "Define a Filesystem" and "Mount a Filesystem" sections of the GUI or the `cmgr` reference chapters in the *CXFS Version 2 Software Installation and Administration Guide*.

Restarting the Windows Node

After completing the steps in "Postinstallation Steps for Windows", page 73, and this chapter, you should restart the Windows node. This will automatically start the driver and the Client service.

When you log into the node after restarting it, Windows Explorer will list the CXFS drive letter, which will contain the CXFS filesystems configured for this node.

Verifying the Cluster

To verify that the client-only nodes have been properly added to the cluster and that filesystems have been mounted, use the view area of the CXFS GUI, the `clconf_info` command, and the `cluster_status` command on an IRIX node.

For example:

```
irix# /var/cluster/cmgr-scripts/cluster_status

+ Cluster=cxfscluster  FailSafe=Not Configured  CXFS=ACTIVE  15:15:33
  Nodes =   cxfs6   cxfs7   cxfs8   solaris1   solaris2
FailSafe =
  CXFS =      UP      UP      UP      UP      UP

CXFS      DevName      MountPoint      MetaServer      Status
fs1       /dev/cxvm/fs1      /fs1             cxfs7           UP
fs2       /dev/cxvm/fs2      /fs2             cxfs6           UP
```

On client-only nodes, the `cxfs_info` command serves a similar purpose. The command path is as follows:

- Solaris: `/usr/cxfs_cluster/bin/cxfs_info`
- Windows: `\program files\CXFS\cxfs_info.exe`

On Solaris nodes, you can use the `-e` option to wait for events, which keeps the command running until you kill the process and the `-c` option to clear the screen between updates.

On Windows nodes, these options are enabled by default, and the window will stay up until you close it. To disable these options on Windows and get the standard UNIX behavior, use the `-D` option.

For example, on a Solaris node:

```
solaris# /usr/cxfs_cluster/bin/cxfs_info

cxfs_client status [timestamp Jul 19 13:30:22 / generation 21604]

Cluster:
  zagato (1) - enabled
Local:
  thump (2) - enabled, state: stable, cms: up, xvm: up, fs: up
```

Nodes:

```
leesa      enabled up    0
thump      enabled up    2
thunderbox enabled up    1
```

Filesystems:

```
bigstripe0 enabled mounted          bigstripe0          /mnt/bigstripe0
concat0    enabled mounted          concat0             /mnt/concat0
mirror0    enabled mounted          mirror0             /mnt/mirror0
r0lun0s0   enabled mounted          r0lun0s0            /mnt/cxfs0
r0lun0s1   enabled mounted          r0lun0s1            /mnt/cxfs1
r0lun0s2   enabled mounted          r0lun0s2            /mnt/cxfs2
stripe0    enabled mounted          stripe0              /mnt/stripe0
```

Forced Unmount of CXFS Filesystems

Normally, an unmount operation will fail if any process has an open file on the filesystem. However, a *forced unmount* allows the unmount to proceed regardless of whether the filesystem is still in use. To enable forced unmount, define or modify the filesystem to unmount with force and then unmount the filesystem, using the following `cmgr(1M)` commands:

```
define cxfs_filesystem logical_filesystem_name [in cluster clustername]
    set force to true

modify cxfs_filesystem logical_filesystem_name [in cluster clustername]
    set force to true

admin cxfs_unmount cxfs_filesystem filesystemname [on node nodename] [in cluster clustername]
```

For example, the following set of commands modifies the `fs1` filesystem to allow forced unmount, then unmounts the filesystem on all nodes in the `cxfscluster` cluster:

```
cmgr> modify cxfs_filesystem fs1 in cluster cxfscluster
Enter commands, when finished enter either "done" or "cancel"cmgr>

cxfs_filesystem fs1 ? set force to true
cxfs_filesystem fs1 ? done
Successfully defined cxfs_filesystem fs1
```

```
cmgr> admin cxfs_unmount cxfs_filesystem fs1 in cluster cxfscluster
```

For details, see the “CXFS Filesystems Tasks with the GUI” sections of the GUI or the `cmgr` reference chapters in the *CXFS Version 2 Software Installation and Administration Guide*.

Troubleshooting

This chapter contains the following:

- "Identifying Problems on Solaris Nodes"
- "Identifying Problems on Windows Nodes", page 96
- "Common Problems and Solutions", page 101
- "Reporting Problems", page 104

Identifying Problems on Solaris Nodes

The following sections will help you identify problems with Solaris client-only nodes:

- "Is the Solaris Node in the Cluster?"
- "Are there Error Messages for the Solaris Node?"

Is the Solaris Node in the Cluster?

To determine if the node is in the cluster, use the `cluster_status` command or the CXFS GUI on an IRIX node. See "Verifying the Cluster", page 92.

Are there Error Messages for the Solaris Node?

Look at the `/var/log/cxfs_client` log to see if there are any error or warning messages. These include any messages containing the words `ERROR` or `Warning`. Specific cases in which error messages will occur include the following:

- The fencing file was not found, therefore the fencing configuration will not be updated on the server. For example:

```
cxfs_client: cis_get_hba_wwns warning: fencing configuration file "fencing.conf" not found
```

- A filesystem mount has failed and will be retried. For example:

```
cxfs_client:op_failed ERROR: Mount failed for concat0
```

For more information about these files, see "Solaris Log Files", page 23. Also see the log files on the IRIX node; for more information, see the *CXFS Version 2 Software Installation and Administration Guide*.

Identifying Problems on Windows Nodes

The following sections will help you identify problems with Windows client-only nodes:

- "Is the CXFS Software Running Correctly on the Windows Node?"
- "Is the Windows Node in the Cluster?", page 98
- "Are There Error Messages for the Windows Node?", page 98

Is the CXFS Software Running Correctly on the Windows Node?

The methods used to verify that the CXFS software is running correctly varies by Windows platform.

Windows NT CXFS Software Verification

To verify that the CXFS software is running correctly on a Windows NT node, do the following:

- Verify that the CXFS driver has started by selecting the following:

```
Start
  > Settings
    > Control Panel
      > Devices
```

- Verify that the CXFS Client service has started by selecting the following:

Start
 > **Settings**
 > **Control Panel**
 > **Services**

Windows 2000 CXFS Software Verification

To verify that the CXFS software is running correctly on a Windows 2000 node, do the following:

- Verify that the CXFS driver has started by selecting the following:

Start
 > **Settings**
 > **Control Panel**
 > **Administrative Tools**
 > **Computer Management**
 > **System Tools**
 > **Device Manager**

To show non-plug-and-play devices, select the following:

View
 > **Show hidden devices**

To show the CXFS driver, select the following:

Non-Plug and Play Devices
 > **CXFS**
 > **Properties**

- Verify that the CXFS Client service has started by selecting the following:

Start
 > **Settings**
 > **Control Panel**
 > **Administrative Tools**
 > **Services**

Is the Windows Node in the Cluster?

To determine if the Windows node is in the cluster, use the `cluster_status` command or the CXFS GUI on an IRIX node, and the `cxfs_info` command on the Windows node. See "Verifying the Cluster", page 92.

Are There Error Messages for the Windows Node?

Look in the following file to see if there are any error or warning messages:

`C:\Program Files\CXFS\log\cxfs_client.log`

You can also view the **System Event** log by selecting the following:

- Windows NT:

Start
 > **Programs**
 > **Administrative Tools**
 > **Event Viewer**

- Windows 2000:

Start
 > **Settings**
 > **Control Panel**
 > **Administrative Tools**
 > **Event Viewer**

Windows Error Message Explanations

Following are typical Windows error messages and their meanings:

`cis_get_hba_wwns` warning: fencing configuration file "fencing.conf" not found

The fencing file `fencing.conf` was not found, therefore the fencing configuration will not be updated on the server.

`op_failed` ERROR: Mount failed for `concat0`

A filesystem mount has failed and will be retried.

`cis_generate_userid_map` warning: could not open `passwd` file

The `passwd` file could not be found.

`cis_generate_userid_map` warning: could not open `group` file

The `group` file could not be found.

Even with `passwd` and `group` warnings above, filesystem mounts should proceed; however, all users will be given `nobody` credentials and will be unable to view or modify files on the CXFS filesystems. For more information about these files, see "Solaris Log Files", page 23, and "Windows Log Files", page 55. Also see the log files on the IRIX node; for more information, see the *CXFS Version 2 Software Installation and Administration Guide*.

could not get location of `passwd/group` files
could not retrieve fencing configuration file name from registry
error retrieving `passwd` filename
error retrieving `group` filename
error retrieving fencing filename

The registry entries for the location of the `passwd`, `group`, or `fencing.conf` files may be missing, or the path provided on the command line to the CXFS Client service is badly formed. Reset these values by modifying the current installation as described in "Modifying the CXFS for Windows Software", page 82.

could not open `passwd` file
could not open `group` file
fencing configuration file not found

Check that the `passwd`, `group` and `fencing.conf` files are in the configured location and are accessible as described in "Checking Permissions on the Password, Group, and Fencing Files", page 78.

Unix user is something other than a user on the NT domain/workgroup
Unix group is something other than a group on the NT domain/workgroup

This warning indicates that a username or groupname is not a valid user or group on the Windows node, which may be confusing when examining file permissions.

no valid users configured in passwd file

No users in the passwd file could be matched to users on the Windows node. All users will be treated as user nobody for the purpose of all access control checks.

no valid groups configured in group file

No groups in the group file could be matched to groups on the Windows node. Attempts to display file permissions will most likely fail with the message Unknown Group Errors.

cis_driver_init() failed: could not open handle to driver
cis_driver_init() failed: could not close handle to CXFS driver

The CXFS driver may not have successfully started. Check the system event log for errors.

unable to create mount point
Configured drive letter may already be in use

Check that the configured drive letter is not already in use by a physical or mapped drive.

unable to join multicast group on interface
unable to create multicast socket
unable to allocate interface list
unable query interfaces
failed to configure any interfaces
unable to create multicast socket
unable to bind socket

Check the network configuration of the Windows node, ensuring that the private network is working and the Windows node can at least reach the metadata server by using the ping command from a command shell.

Common Problems and Solutions

This section contains the following common problems and their solutions:

- "Incorrect Configuration"
- "Determining If A Client-Only Node Is Fenced", page 101
- "Common Solaris Problems: JNI Problems", page 101
- "Common Windows Problems", page 102

Incorrect Configuration

To avoid having trouble with the CXFS client-only node, ensure you have the correct configuration. See "Requirements", page 5.

Determining If A Client-Only Node Is Fenced

To determine if a client-only node is fenced, log in to IRIX and use the `hafence(1M)` command. For more details, see the *CXFS Version 2 Software Installation and Administration Guide*.

Common Solaris Problems: JNI Problems

If you have difficulty with the JNI verification steps, consult the following checklist to help you identify the problem:

- Is the HBA firmly seated in its PCI slot?
- Are all cables undamaged and connected?
- Is power applied to all devices?
- Do the link lights illuminate on all units?
- Is the problem confined to just one TP9400 unit? If so, check the cabling between the switch and the unit; if no units are being shown, suspect cabling from the HBA.
- Is the Brocade switch properly licensed?
- Did you enable fabric mode? See step 4, in "Installing the JNI HBA", page 26.

For more information, see the *Installation Guide, FCE-6460 and FCE2-6460 PCI-to-Fibre Channel Host Bus Adapters (Solaris, Windows NT/2000, Novell, AIX, HP-UX, Mac OS, Linux) JNI FibreStar* or contact your SGI service representative.

Common Windows Problems

This section contains the following common Windows problems:

- "Windows QLogic Problems"
- "Filesystems Are Not Displayed on a Windows Node", page 103
- "Large Log Files on Windows", page 103
- "Windows Failure on Restart", page 104
- "Memory Configuration of the Windows Node", page 104

Windows QLogic Problems

If you have difficulty with the QLogic verification steps, consult the following checklist to help you identify the problem:

- Is the HBA firmly seated in its PCI slot?
- Are all cables undamaged and connected?
- Is power applied to all devices?
- Do the link lights illuminate on all units?
- Is the problem confined to just one TP9400 unit? If so, check the cabling between the switch and the unit; if no units are being shown, suspect cabling from the HBA.
- Is the Brocade switch properly licensed?
- Check the QLogic management tool event and alarm logs. Select the following:

Start

> **Programs**

> **QLogic Management Suite**

> **SANsurfer**

For more information, see the following QLogic documentation or contact your SGI service representative:

- *Hardware Installation Guide for the QLA2xxx Board Family*
- *Software Installation Guide for the QLA2xxx Board Family*

Also see the QLogic website at:

<http://www.qlogic.com>

Filesystems Are Not Displayed on a Windows Node

If the CXFS drive letter is visible in Windows Explorer but no filesystems are mounted, do the following:

- Run `C:\Program Files\CXFS\cxfs_info` to ensure that the filesystems have been configured for this node.
- Verify the filesystems that should be mounted by using the `cmgr(1M)` command on an IRIX node. For more information, see "Mounting Filesystems on the Client-Only Nodes", page 91.
- Ensure that the CXFS metadata server is up and that the Windows node is in the cluster membership; see "Verifying the Cluster", page 92.
- Check that the CXFS Client service has started. See "Is the CXFS Software Running Correctly on the Windows Node?", page 96, and "Manual CXFS Startup/Shutdown for Windows", page 81.
- Check the following file for warnings and errors regarding licenses or mounting filesystems:

`C:\Program Files\CXFS\log\cxfs_client.log`

- Check the cluster configuration to ensure that this node is configured to mount one or more filesystems.

Large Log Files on Windows

The CXFS Client service creates the following log file:

`C:\Program Files\CXFS\log\cxfs_client.log`

This log file may become quite large over a period of time if the verbosity level is increased. The service does not perform any automatic log rotation, so the service must be stopped in order to move or truncate this file, then restarted. See "Manual CXFS Startup/Shutdown for Windows", page 81, on how to stop and start the CXFS Client Service.

Windows Failure on Restart

If the CXFS Windows node fails to start and terminates in a blue screen, restart your computer, and select the backup hardware profile (with CXFS disabled). Alternatively, pressing **L** at the **Hardware Profile** menu will select the last configuration that was successfully started and shut down. If the node has only one hardware profile, press the spacebar after selecting the boot partition to get to the **Hardware Profile** menu.

Memory Configuration of the Windows Node

A Windows problem may affect Windows CXFS nodes performing large asynchronous I/O operations. If the Windows node crashes with a `NO_MORE_SYSTEM_PTES` message, the work-around described in the following link should be considered (line break added here for readability):

http://www.microsoft.com/technet/treeview/default.asp?url=/technet/prodtechnol/winxp/pro/reskit/prmd_stp_fztl.asp

Reporting Problems

When reporting a problem with a client-only node, it is important to retain the appropriate information; having access to this information will greatly assist SGI in the process of diagnosing and fixing problems. The methods used to collect required information for problem reports are platform-specific.

Reporting Solaris Problems

When reporting a problem about a CXFS Solaris node to SGI, you should retain the following information:

- System core files in `/var/crash/hostname` on Solaris nodes.
- `mdb(1M)` modular debugger output:

- For panics or generated dumps, use the following commands and save the output:
 - `$c` (or `$C`)
 - `$r`
 - `$<msgbuf`
- For dumps from hangs:
 - `$<threadlist`
 - `$c` (or `$C`)
 - `$r`
 - `$<msgbuf`
- Output from the `crash(1M)` utility.
- A list of the installed CXFS packages. Use the `pkginfo(1)` command as follows:

```
# pkginfo -l SGIcxfs
```
- A list of the Solaris patches that have been installed. Use the `showrev(1M)` command. The `showrev` command without options prints a summary and the `-p` option lists the revision information about patches.
- A list of the loaded Solaris kernel modules and versions. Use the `modinfo(1M)` command.
- Output about the cluster obtained from the IRIX `cdump` utility run on an IRIX node. This utility is available from your SGI service representative.

If any of the above Solaris tools are not currently installed on your Solaris system, you should install them.

Reporting Windows Problems

To report problems about a Windows node, you should retain platform-specific information and save crash dumps.

Retain Windows NT Information

When reporting a problem about a CXFS Windows NT node to SGI, you should retain the following information:

- The version and memory of the current diagnostics. Select the following:

Start

> **Programs**

> **Administrative Tools**

> **Windows NT Diagnostics**

Record the version displayed in the **Version** tab and the total physical memory in the **Memory** tab.

- The build date and firmware versions. Using Windows Explorer, open the following directory:

C:\Winnt\system32\drivers

Then do the following:

- Right click on `cxfs.sys` and select the following:

Properties

> **Version**

Record the values of **BuildDate** and **Product Version**.

- Right click on `ql2200.sys` and select the following:

Properties

> **Version**

Record the values of **Firmware** and **Product Version**.

- The contents of the following file:

C:\Program Files\CXFS\log\cxfs_client.log

Compress this file with `winzip` if it is large.

- The contents of the crash dump if one was generated, Compress this file with `winzip`. For more information, see "Save Crash Dumps for Windows", page 108.

Retain Windows 2000 Information

When reporting a problem about a CXFS Windows 2000 node to SGI, you should retain the following information:

- The configuration of the machine. Select the following:

Start

- > **Programs**
 - > **Accessories**
 - > **System Tools**
 - > **System Information**
 - > **Action**
 - > **Save As System Information File**

This will create a file that describes all of the installed hardware and configured drivers on the machine.

Alternatively, you could dump information about each item in the hardware tree to a text file by using the following selection:

Action

- > **Save As Text File**

However, you must repeat this action for each item.

- The build date and firmware versions. Using Windows Explorer, open the following directory:

`C:\Winnt\system32\drivers`

Then do the following:

- Right click on `cxfs.sys` and select the following:

Properties

- > **Version**

Record the values of **BuildDate** and **Product Version**.

- Right click on `ql2200.sys` and select the following:

Properties
 > **Version**

Record the values of **Firmware** and **Product Version**.

- The contents of the following file:
`C:\Program Files\CXFS\log\cxfs_client.log`
Compress this file with `winzip` if it is large.
- The contents of the crash dump if one was generated, Compress this file with `winzip`. For more information, see "Save Crash Dumps for Windows", page 108.

Save Crash Dumps for Windows

If you are having problems, you should configure the Windows node to save crash dumps to a filesystem that is not a CXFS filesystem. To do this, click the right mouse button on the **My Computer** icon and select the following:

- Windows NT:

Properties
 > **Startup**
 > **Shutdown**
 > **Write debugging information to**

- Windows 2000:

Properties
 > **Advanced**
 > **Startup and Recovery**
 > **Write debugging information to**

Enter a path on a filesystem other than a CXFS filesystem. This change will take affect only after the node is restarted.

Glossary

active metadata server

A weighted IRIX node chosen from the list of potential metadata servers. There can be multiple active metadata servers, one for each file system.

client

See *CXFS client*.

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. Thus, if two clusters will be sharing the same network for communications, then they must have unique cluster IDs. In the case of multiOS clusters, both the names and IDs must be unique if the clusters share a network.

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 database

Contains configuration information about nodes, filesystems, and the cluster. The database is managed by the `fs2d` daemon and is stored on IRIX nodes in the pool of nodes that are running the `fs2d` daemon.

cluster ID

A unique number within your network in the range 1 through 128. The cluster ID is used by the IRIX kernel to make sure that it does not accept cluster information from

any other cluster that may be on the network. The kernel does not use the database for communication, so it requires the cluster ID in order to verify cluster communications. This information in the kernel cannot be changed after it has been initialized; therefore, you must not change a cluster ID after the cluster has been defined. Clusters that share a network must have unique names and IDs.

cluster node

A node that is defined as part of the cluster.

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.

A node's control networks should not be set to accept control messages if the node is not a dedicated CXFS node. Otherwise, end users who run other jobs on the machine can have their jobs killed unexpectedly when CXFS resets the node.

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

A node that is part of the cluster and is a potential metadata server, but is currently not acting as the active metadata server. See also *IRIX node* and *CXFS client-only node*.

CXFS client-only node

A node that is part of the cluster but is not a potential metadata server. Solaris nodes and Windows nodes are client-only nodes. See also *IRIX node*.

CXFS database

See *cluster database*.

CXFS 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 membership. Once accepted, the node can share the filesystems of the cluster. (Also known as *kernel-space membership*.)

database

See *cluster database*.

GUI

Graphical user interface.

heartbeat messages

Messages that cluster software sends between the nodes that indicate a node is operational. Heartbeat messages and control messages are sent through the node's network interfaces that have been attached to a control network.

I/O fencing

The failure action 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.

IRIX node

A CXFS node that is running the IRIX operating system. An IRIX node can be a potential metadata server if it is configured as such and has weight, or it can be a CXFS client. In a multiOS cluster, all IRIX nodes must be running a minimum of 6.5.16f.

membership

See *CXFS membership*.

membership weight

A number (usually 0 or 1) that is assigned to a node for purposes of calculating the CXFS membership quorum. 1 indicates that the node is eligible to be a potential

metadata server. IRIX nodes may have a weight of 0 or 1. CXFS client-only nodes always have a weight of 0.

metadata

Information that describes a file, such as the file's name, size, location, and permissions.

metadata server

The IRIX 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. See also *active metadata server* and *potential metadata server*.

multiOS cluster

A cluster that is running multiple operating systems, such as IRIX and Solaris.

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 pool (and therefore) only one cluster.

A node can run the IRIX operating system or another operating system, such as Solaris, as defined in the CXFS Version 2 for CXFS Client-Only Nodes: Installation and Configuration Guide.

node membership

The list of nodes that are active (have CXFS membership) in a cluster.

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 or `cmgr(1M)` command. 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).

potential metadata server

A weighted IRIX node that is listed in the metadata server list when defining a filesystem; there can be multiple potential metadata servers, but only one node in the list will be chosen as the active metadata server for one filesystem.

recovery

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

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.

SAN

Storage area network, a high-speed, scalable network of servers and storage devices that provides storage resource consolidation, enhanced data access/availability, and centralized storage management.

tree view

The portion of the CXFS GUI window that displays components graphically.

quorum

The number of nodes required to form a cluster.

weight

See *membership weight*.

Index

8-port switch, 15
16-port switch, 15
100baseT TCP/IP network, 6
4774 and 4884 units, 13

A

adapter parameters, 33, 36
add
 client-only node to the cluster, 89
 package, 46
 private network
 Solaris, 40
 Windows 2000, 67
 Windows NT, 66
adjacent IRIX releases, 8
admin cxfs_unmount, 91
administrative tasks, 4

B

bandwidth, 2, 5
BIOS version, 54
block size, 24
boot command, 27
boot Solaris in single-user mode, 40
Brocade
 license, 3
 switch, 6, 15

buffered I/O and CXFS, 5
build date for Windows NT, 106, 107

C

\$c or \$C, 105
C:\Program Files\CXFS directory, 69
C:\Program Files\CXFS\log\cxfs_client.log
 file, 98, 103
C:\ProgramFiles\CXFS\group file, 78
C:\ProgramFiles\CXFS\passwd file, 77
C:\Winnt\system32\drivers directory, 106, 107
cables and JNI HBA, 101, 102
client software installation
 Solaris, 45
 Windows, 69
client-only node configuration
 add to the cluster, 89
 define the node, 87
 define the switch, 90
 modify the cluster, 89
 mount filesystems, 91
 permit fencing, 87
 start CXFS services, 90
 verify the cluster, 92
cluster
 configuration, 87
 size, 8
 verification, 92
cluster_status command, 92
cmgr command, 87
command shell for DOS, 74
commands installed
 Solaris, 22
 Windows, 54

- common problems
 - Solaris, 101
 - Windows, 102
- concepts, 1
- configuration overview, 11
- configuration problems, 101
- configure
 - cluster, 87
 - network interfaces, 39, 65
 - Solaris system files, 39
 - Windows system files, 65
- contacting SGI with problems
 - Solaris, 104
 - Windows 2000, 107
 - Windows NT, 106
- controller firmware, 13
- core files, 104
- crash dumps
 - Solaris, 104
 - Windows, 108
- crash utility and gathering output, 105
- crontab, 8
- CXFS
 - GUI and cmgr, 87
 - license, 3
 - metadata server, 3
 - software removal on Windows, 84
 - startup/shutdown
 - Solaris, 50
 - Windows, 81
- CXFS Client service
 - command line arguments, 71
- CXFS startup/shutdown
 - Windows, 81
- cxfs_client command, 22
- cxfs_cluster command, 50
- cxfslicense command, 19, 22, 25, 47

D

- define a client-only node, 87
- defragmentation software and CXFS, 8
- dflt_local_status, 91
- DHCP, 67
- direct-access I/O, 2
- disk device verification for Solaris, 37
- display LUNs for QLogic HBA, 61
- distributed applications and CXFS, 5
- dmesg command, 42
- DNS, 41
- DOS command shell, 74
- driver software for JNI HBA, 27
- dumps and output to gather, 105

E

- entitlement ID, 19
- Entitlement Sheet, 6
- error messages
 - Solaris, 95
 - Windows, 98
- /etc/fencing.conf file, 35, 48
- /etc/hostname.<interface> file, 43
- /etc/hosts file, 40, 88
- /etc/inet/hosts file, 40, 88
- /etc/init.d/cxfs_cluster command, 50
- /etc/netmasks file, 43
- /etc/nodename file, 43
- /etc/nsswitch.conf file, 40, 41
- /etc/sys_id file, 43
- Event Viewer, 55
- examples
 - add a client-only node to the cluster, 89
 - CXFS software installation
 - Solaris, 46
 - Windows, 70
 - define a node, 88
 - define a switch, 90

- /etc/hosts file, 41
- /etc/inet/hosts file, 41
- fabric enable, 27
- ifconfig for Solaris, 45
- JNI GUI screens, 30
- license properly installed, 19
- modify the cluster, 89
- modifying the CXFS
 - Solaris, 51
 - Windows, 82
- mount filesystems, 91
- name services, 41
- ping output for Solaris, 45
- private network interface test for Solaris, 45
- .rhosts, 44
- start CXFS services, 90
- upgrade CXFS for Windows software, 84
- verify the cluster configuration, 92
- Windows Client service command line
 - options, 83
- EZ Fibre GUI, 28
- ezf, 29

F

- fabric mode, 27
- fail action hierarchy, 88
- FailSafe coexecution, 7
- failure on restart, 104
- FcFabricEnabled, 27
- FcLoopEnabled, 27
- FCO 1056, 13
- fence specification in node definition, 88
- fencing feature, 15
- fencing file, 75
 - See "I/O fencing file", 75
- fencing.conf file, 35, 48
- Fibre Channel HBA, 26
- Fibre Channel requirements for Solaris, 22
- FibreStar HBA, 26

- file size, 5
- filesystem network access, 2
- filesystems not displayed on Windows, 103
- find and crontab, 8
- FLEXlm
 - license key, 6
 - license requirements for CXFS, 19
 - license verification
 - Solaris, 25
 - Windows, 74
 - licenses, 3
- forced unmount, 9
- format command, 37
- free disk space required, 54
- fuser, 10

G

- generated dumps and output to gather, 105
- good practices, 9
- group file on Windows, 78
- guided configuration, 87

H

- hangs and output to gather, 105
- hardware profile, 78
- hardware requirements
 - all platforms, 5
 - Solaris, 21
 - Windows, 54
- HBA, 26, 27
- heartbeat/control, 40
- hierarchy of fail actions, 88
- host bus adapter, 26
- hostid command, 19
- hostname command, 19
- hostname resolution, 39, 65
- hostname.<interface> file, 43

hosts file, 40, 88
hub, 8

I

I/O fencing, 6, 15
I/O fencing file
 Solaris, 35, 48
 Windows, 75
I/O operations, 2
icons on JNI GUI, 31
identifying problems
 Solaris, 95
 Windows, 96
ifconfig command, 41, 42, 45
incorrect configuration, 101
initial setup services, 1
initialize the private network's interface for
 Solaris, 42
install.sh script, 28
installation overview, 11
installed packages, 105
installed patches, 105
integrated Ethernet, 42
Intel Pentium processor, 54
interface for the private network, 42
internode communication, 39, 65
introduction, 1
IP address, changing, 39, 65
IP forwarding, 67
ipconfig, 68
IRIX
 labels in warning messages, 38
 node required, 4, 6
 operating system level, 6

J

JBOD, 6

JNI

HBA installation, 26
HBA requirement, 22
troubleshooting, 101, 102
JNIC146x, 27
jnic146x.conf file, 27
JNISnia, 27

K

kernel modules and versions, 105
/kernel/drv/jnic146x.conf file, 27, 28

L

large files, 2
license
 Brocade, 3, 15
 CXFS, 3, 19
 FLEXlm for CXFS and XVM, 19
 installation, 20
 mirroring, 19
 SGI webpage, 20
 verification
 Solaris, 25
 verification on Solaris, 25
 XVM, 3, 19
licenseshow command, 16
licensing, 6
link lights and JNI HBA, 101, 102
log files
 Solaris, 23
 Windows, 55, 103
LUN
 logical unit, 36
 LUN 31 and mapping, 36
 mapping, 36
 maximum with TP9400, 13
 zoning, 35

M

- maintenance and cluster services, 9
- manual CXFS startup/shutdown
 - Solaris, 50
 - Windows, 81
- mapping
 - LUNs, 36
- masking feature and Solaris, 90
- mdb, 105
- memory mapped shared files and CXFS, 5
- messages
 - See "error messages", 95, 98
- metadata
 - concept, 2
 - server, 3
 - server weight, 9
 - transaction examples, 5
 - transfer and performance considerations, 5
- minimum block size, 24
- mirroring feature and license, 3, 19
- missing filesystems on Windows, 103
- modify cluster command, 89
- modify the CXFS software
 - Solaris, 51
 - Windows, 82
- modinfo, 105
- modules and versions, 105
- Mojave code, 13
- mount filesystems, 91
- mount-point nesting on Solaris, 22
- msgbuf, 105
- \$(msgbuf, 105
- multiOS cluster, 1

N

- name restrictions, 39, 65
- name service daemon, 41
- nested mount points on Solaris, 22

- netmasks file, 43
- network
 - information service, 41
 - interface configuration, 39, 65
 - issues, 8
 - requirements, 6
 - switch, 8
- NFS and CXFS, 5, 22
- NIS, 41
- nobody (user name on Windows), 77
- node weight, 8
- nsd, 41
- nsswitch.conf file, 40, 41
- number of nodes supported, 7
- NVSRAM files, 13

O

- O2, 6
- operating system level, 6
- order desk, 20
- OS level, 6
- output to gather
 - Solaris, 104
 - Windows 2000, 107
 - Windows NT, 106
- overview, 11

P

- packages installed, 105
- panics and output to gather, 105
- partitioned system licensing, 6
- partitioning key for TP9400, 13
- passwd file on Windows, 77
- patches installed, 105
- PCI slot and JNI HBA, 101, 102
- performance considerations
 - Solaris, 5

- Windows, 5
- permissions for Windows passwd and group files, 78
- ping, 45
- pkgadd command, 22, 46
- pkginfo command, 47, 105
- plexing license, 3
- plumb, 42
- postinstallation steps
 - Solaris, 48
 - Windows, 73
- power management software and CXFS, 9
- power to JNI HBA, 101, 102
- preinstallation steps
 - Solaris, 38
 - Windows, 64
- premount and postmount scripts, 4
- primary hostname
 - Solaris, 40
 - Windows, 68
- private network
 - heartbeat and control, 39, 65
 - interface test for Solaris, 45
 - required, 6
 - Solaris, 40
- problem reporting
 - Solaris, 104
 - Windows 2000, 107
 - Windows NT, 106
- public network
 - Solaris, 42

Q

- Qlogic HBA installation, 60
- QLogic HBA model numbers and driver versions, 54

R

- \$r, 105
- RAID firmware, 13
- READ_CAPACITY SCSI command, 24
- recommendations for CXFS
 - general, 8
- recovery, 9
- relocation, 9
- remove CXFS software
 - Windows, 84
- reporting problems
 - Solaris, 104
 - Windows 2000, 107
 - Windows NT, 106
- requirements for CXFS
 - all platforms, 5
 - Solaris, 21
 - Windows, 54
- reset lines, 6
- /.rhosts file, 43, 44

S

- serial reset lines, 6
- service pack, 54
- set dflt_local_status, 91
- setup program for Windows, 69
- setup services, 1
- SGIcxf package, 46
- shell for DOS, 74
- showrev command, 105
- Silicon Graphics O2, 6
- Silkworm switch, 15
- single-user mode in Solaris, 40
- size of the cluster, 7
- small files, 5
- SNIA API package, 27
- software maintenance
 - Solaris, 50

- Windows, 82
- software requirements
 - all platforms, 5
 - Solaris, 21
 - Windows, 54
- software upgrades
 - Solaris, 51
 - Windows, 84
- Solaris
 - client software installation, 45
 - commands installed by CXFS, 22
 - error messages, 95
 - FLEXlm license verification, 25
 - I/O fencing file, 48
 - identifying problems, 95
 - JNI HBA installation, 26
 - JNI problems, 101, 102
 - kernel
 - modules and versions, 105
 - limitations, 23
 - log files, 23
 - manual CXFS startup/shutdown, 50
 - minimum block size, 24
 - modify the CXFS software, 51
 - non-disk devices and JNI controllers, 24
 - operating system version, 22
 - performance considerations, 5
 - postinstallation steps, 48
 - preinstallation steps, 38
 - problems, 101
 - recommendations, 8
 - requirements, 5, 21
 - software
 - maintenance, 50
 - upgrades, 51
 - space requirements, 46
 - user
 - administration for CXFS, 4
 - user and group quotas, 23
 - solutions to common problems
 - Solaris, 101

- Windows, 102
- space requirements for Solaris, 46
- start
 - CXFS processes
 - Solaris, 50
 - Windows, 81
 - CXFS services, 81, 90
- stop CXFS processes on Solaris, 50
- Storage Networking Industry Association
 - application programming interface package, 27
- subnet, 6
- Sun hardware, 22
- switch, 8, 15
- switch definition, 90
- switchshow
 - Solaris, 49
 - Windows, 76
- sys_id file, 43
- system core files, 104
- System Event log, 98

T

- TCP/IP network requirements, 6
- telnet, 76
- telnet port and I/O fencing, 7
- \$(threadlist), 105
- TP9100, 13
- TP9100 RAID firmware, 13
- TP9400 RAID firmware, 13
- TP9400 RAID troubleshooting, 101, 102
- TRIX and Solaris nodes, 7
- troubleshooting, 95
- Trusted IRIX and Solaris nodes, 7

U

- UFS and CXFS, 22
- Ultra Enterprise platforms, 22

- umount, 10
- uname, 46
- unmount the filesystems, 9
- upgrade CXFS software
 - Solaris, 51
 - Windows, 84
- user and group quotas, 23
- user configuration for Windows, 77
- /usr/bin/hostid command, 19
- /usr/bin/hostname command, 19
- /usr/bin/showrev command, 105
- /usr/cxfs_cluster/bin/cxfs_client command, 22
- /usr/cxfs_cluster/bin/cxfslicense command, 19, 22, 25
- /usr/cxfs_cluster/bin/xvmprobe command, 22

V

- /var/cluster/cmgr-scripts/cluster_status, 92
- /var/crash/<hostname>, 104
- /var/log/cxfs_client log file, 23
- verify
 - Brocade license and firmware, 16
 - cluster, 92
 - FLEXlm license
 - Solaris, 25
 - Windows, 74
- verify networks
 - Windows, 68
- version command, 16
- version of Windows NT, 106, 107
- versions of modules installed, 105
- volume plexing license, 3

W

- warning message and IRIX labels, 38
- weighted nodes, 6, 8
- when to use CXFS, 2

Windows

- client software installation steps, 69
- crash dumps, 108
- CXFS commands installed, 54
- CXFS software removal, 84
- debugging information, 108
- DHCP, 67
- error messages, 98
- failure on restart, 104
- filesystems not displayed, 103
- FLEXlm license verification, 74
- hardware profile, 78
- heartbeat and control, 65
- hostname resolution, 65
- I/O fencing file, 75
- identifying problems, 96
- IP address, changing, 65
- IP forwarding, 67
- ipconfig, 68
- large log files, 103
- log files, 55
- LUNs, 61
- manual CXFS startup/shutdown, 81
- modify the CXFS software, 82
- name restrictions, 65
- network interface configuration, 65
- performance
 - considerations, 5
- postinstallation steps, 73
- preinstallation steps, 64
- private network, 65
- problems, 102
- Qlogic HBA installation, 60
- recommendations, 8
- requirements, 5, 54
- software maintenance, 82
- software upgrades, 84
- user, 4
 - configuration, 77
- verify networks, 68
- version, 106, 107

WINS server, 67
Windows NT
 build date, 106, 107
winnt/setup.exe command, 69, 84
WINS server, 67
worldwide node name, 28
worldwide port name, 28, 48, 76
WWNN, 28
WWPN, 28, 48, 76

X

 xfs_repair, 8
XVM mirroring license, 3, 19
xvmprobe command, 22