

Origin™ FibreVault and Fibre Channel RAID Owner's Guide

Document Number 007-3428-004

CONTRIBUTORS

Written by Mark Schwenden and Carolyn Curtis

Illustrated by Dan Young, Cheri Brown, and Carolyn Curtis

Production by Heather Hermstad

Engineering contributions by James Oliver, Sammy Wilborn, Radek Aster,
Paul Stolle, Henry Ortiz, Ron Naminski, James Bergman, Frank Demcak,
Jeremy Higdon, Bob Housholder, and Paul Tsien

St. Peter's Basilica image courtesy of ENEL SpA and InfoByte SpA. Disk Thrower
image courtesy of Xavier Berenguer, Animatica.

Some information in this document was derived from the Fibre Channel Loop
Community (FCLC) Web site, and the Fibre Channel Association (FCA) technical
information Web site.

The authors gratefully acknowledge the assistance of the publications group at
CLARiiON Storage Systems Inc. in developing this manual.

© 1998, Silicon Graphics, Inc.— All Rights Reserved

The contents of this document may not be copied or duplicated in any form, in whole
or in part, without the prior written permission of Silicon Graphics, Inc.

RESTRICTED RIGHTS LEGEND

Use, duplication, or disclosure of the technical data contained in this document by
the Government is subject to restrictions as set forth in subdivision (c) (1) (ii) of the
Rights in Technical Data and Computer Software clause at DFARS 52.227-7013
and/or in similar or successor clauses in the FAR, or in the DOD or NASA FAR
Supplement. Unpublished rights reserved under the Copyright Laws of the United
States. Contractor / manufacturer is Silicon Graphics, Inc., 2011 N. Shoreline Blvd.,
Mountain View, CA 94043-1389.

Attention

This product requires the use of external shielded cables in order to maintain
compliance pursuant to Part 15 of the FCC Rules.

FCC Warning

This equipment has been tested and found compliant with the limits for a Class A digital device, pursuant to Part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at personal expense.

VDE 0871/6.78

This equipment has been tested to and is in compliance with the Level A limits per VDE 0871.

European Union Statement

This device complies with the European Directives listed on the “Declaration of Conformity” which is included with each product. The CE mark insignia displayed on the device is an indication of conformity to the aforementioned European requirements.



International Special Committee on Radio Interference (CISPR)

This equipment has been tested to and is in compliance with the Class A limits per CISPR publication 22.

Canadian Department of Communications Statement

This digital apparatus does not exceed the Class A limits for radio noise emissions from digital apparatus as set out in the Radio Interference Regulations of the Canadian Department of Communications.

Attention

Cet appareil numérique n'émet pas de perturbations radioélectriques dépassant les normes applicables aux appareils numériques de Classe A prescrites dans le Règlement sur les interférences radioélectriques établi par le Ministère des Communications du Canada.

Japanese Compliance Statement

この装置は、情報処理装置等電波障害自主規制協議会 (VCCI) の基準に基づくクラス A 情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

Chinese Compliance Statement

警告使用者：

這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻干擾，在這種情況下，使用者會被要求採取某些適當的對策。

Silicon Graphics, InfiniteReality, IRIS, IRIX, OCTANE, and Onyx are registered trademarks, and the Silicon Graphics logo, IRIS FailSafe, Onyx2, Origin, Origin200, and Origin2000 are trademarks of Silicon Graphics, Inc. CLARiiON is a registered trademark of Data General Corporation. CRAY is a registered trademark and CrayLink is a trademark of Cray Research, Inc.

Contents

List of Figures ix

List of Tables xi

About This Guide xiii

Audience xiii

Structure of This Document xiv

Additional Useful Manuals xiv

Typographical Conventions xv

- 1. Fibre Channel Features and Capabilities** 1
 - Fibre Channel Technology Introduction 1
 - Networks, Channels, and Fibre Channel 2
 - Applications 5
 - High Availability and RAID Applications 6
 - Online Administrative Disk Swapping 9
 - Fibre Channel Host Interface Boards 9
 - Fibre Channel XIO Board 10
 - Fibre Channel PCI Board 12
 - Fibre Channel Board Connectors 13
 - Fibre Channel Board Logic 13
 - Copper or Optical Media Interface Adapter 16
 - Chassis Grounding Issues 17
 - Fibre Channel Storage 18
 - Fibre Channel Hardware Components 18
 - FC-AL Disk Modules 18
 - FibreVault and Fibre Channel RAID Enclosures 21
 - The Fibre Channel Rack 24
 - Fibre Channel Deskside Systems 26

- FibreVault Options in Origin and Onyx2 Racks 28
- Fibre Channel Administrative Interface 29
 - High-Level Drivers 30
- 2. Fibre Channel Architecture 31**
 - Networks, Channels, and Fibre Channel 31
 - Fabric Topology 32
 - Point-to-Point Topology 32
 - Arbitrated Loop Topology 32
 - Silicon Graphics Fibre Channel Ports and Topologies 33
 - Data Transfer in Fibre Channel 36
 - Fibre Channel Layers 37
 - FC-0 37
 - FC-1 38
 - FC-2 38
 - FC-3 42
 - FC-4 42
- 3. Fibre Channel Storage 43**
 - The FibreVault Enclosures 44
 - FibreVault Components 47
 - Link Controller Card (LCC) 47
 - Disk Module Fan Assembly 49
 - The Enclosure Power Supplies 50
 - The FibreVault Midplane 52
 - Fibre Channel RAID Enclosure Components 52
 - Fibre Channel RAID Deskside Tower 53
 - Storage Processor Board Assembly 54
 - Storage Processor Fan Module 56
 - Fibre Channel Rack 58
 - Optional Standby Power Supply (SPS) 59
 - Monitoring Storage System Status 61
 - Shutting Down the Fibre Channel Rack and Disk Enclosures 62

4.	Fibre Channel Component Replacement	65
	Adding or Replacing a Disk Module	66
	General Guidelines for Adding or Removing Disks	67
	Confirming a Failed Disk Module	68
	Physically Replacing a Disk	68
	Setting an Enclosure Address	71
	Replacing a Fan Assembly	75
	Replacing or Adding a Power Supply in an Enclosure	77
	Preparing the New Power Supply	78
	Removing a FibreVault Power Supply or Filler	80
	Inserting the New FibreVault Power Supply	83
	Removing and Replacing a Fibre Channel RAID Power Supply	84
	Replacing or Adding a Link Control Card (LCC)	89
5.	Configuring a Fibre Channel System	93
	XIO Interface Configurations	93
	Getting System Status	94
A.	Technical Specifications	95
	Enclosure Height Specifications	95
	Site Requirements for the Fibre Channel Rack and FibreVault Deskside	96
	Enclosure Power Requirements	96
	Enclosure Thermal Requirements	97
	Technical Specifications for the FibreVault Enclosure	97
	Technical Specifications for the FibreVault Deskside (Tower)	98
	Technical Specifications for the Fibre Channel RAID Enclosure	99
	Technical Specifications for the Fibre Channel RAID Deskside Tower	100
	Technical Specifications for the Fibre Channel Rack	102
	Cabling Interfaces	103
	Glossary	107
	Index	113

List of Figures

Figure 1-1	Fibre Channel in the Overall Origin2000 and Onyx2 I/O Structure	10
Figure 1-2	Fibre Channel XIO Board and Connectors	11
Figure 1-3	Fibre Channel PCI Interface Board	12
Figure 1-4	XIO Board Connector Pin Assignments	13
Figure 1-5	Fibre Channel XIO Card Block Diagram	15
Figure 1-6	Optical Module (MIA)	16
Figure 1-7	FC-AL Disk Module	19
Figure 1-8	FibreVault Enclosure (Rear View)	21
Figure 1-9	Fibre Channel RAID Disk Enclosure (Front View)	22
Figure 1-10	Fibre Channel Rack (Front View)	25
Figure 1-11	Fibre Channel Deskside Storage Systems	27
Figure 1-12	FibreVault Enclosures in an Origin Rack (Rear View)	28
Figure 1-13	Fibre Channel Software Interface	29
Figure 2-1	Fibre Channel Arbitrated Loop (FC-AL)	33
Figure 2-2	Port Arbitration Activity Example	35
Figure 2-3	Fibre Channel Layers and OSI Layers	37
Figure 3-1	FibreVault Back View with Fan Module	44
Figure 3-2	FibreVault Deskside (Tower) Unit	45
Figure 3-3	Rackmount FibreVault Enclosure (Fan Module Removed)	46
Figure 3-4	Origin FibreVault Enclosure Front Panel	46
Figure 3-5	FibreVault Link Controller Card (LCC)	48
Figure 3-6	Fan Assembly Features	50
Figure 3-7	Power Supply Features	51
Figure 3-8	Fibre Channel RAID Disk Enclosure (Rear View)	53
Figure 3-9	Deskside Fibre Channel RAID Tower (Rear View)	54
Figure 3-10	Storage Processor Connectors	55
Figure 3-11	Storage Processor Board Components	56

Figure 3-12	Storage Processor Fan Module	57
Figure 3-13	FC-AL Storage Subsystem Example Diagram	58
Figure 3-14	Fibre Channel RAID Standby Power Supply (SPS) Front Panel	59
Figure 3-15	Two Rackmounted SPS Examples With Cabling (Rear View)	60
Figure 3-16	Disconnecting Rack Power Cables	63
Figure 3-17	Removing an Enclosure Disk Fan Module Example	64
Figure 4-1	The Disk Module	70
Figure 4-2	Setting the FibreVault ID	72
Figure 4-3	Fibre Channel RAID Enclosure Channel Address Switches	73
Figure 4-4	Fibre Channel RAID Enclosure Rear View	74
Figure 4-5	Removing the Disk Module Fan Assembly	76
Figure 4-6	Fan Assembly Status LED Orientation	77
Figure 4-7	Unlatching a New Power Supply	79
Figure 4-8	Power Supply Power Cord	81
Figure 4-9	Unlatching a Power Supply	82
Figure 4-10	Removing a FibreVault Power Supply	83
Figure 4-11	Turn Off Power and Unplug Cord from the Supply	85
Figure 4-12	Removing an Existing Fibre Channel RAID Power Supply or Filler Panel	87
Figure 4-13	Installing the Bottom Fibre Channel RAID Power Supply	88
Figure 4-14	LCC Removal Example	90
Figure 4-15	LCC Installation Example	91

List of Tables

Table 1-1	Networks and Channels	3
Table 1-2	Silicon Graphics Storage Interface Solutions	5
Table 1-3	Fibre Channel Disk Drive Descriptions	20
Table 2-1	Silicon Graphics Fibre Channel Cable Characteristics	38
Table 3-1	SPS LEDs	59
Table 3-2	LEDs Visible From the Front of the Fibre Enclosure	61
Table 3-3	LEDs Visible From the Rear of the Fibre Enclosure	62
Table 4-1	Fibre Channel Disk Drive Descriptions	66
Table 4-2	FibreVault or Fibre Channel RAID Disk Slot Numbering	73
Table 5-1	Configuration Maximums	93
Table A-1	Enclosure Heights	95
Table A-2	FibreVault Enclosure Technical Specifications	97
Table A-3	Deskside FibreVault Enclosure Technical Specifications	98
Table A-4	Fibre Channel RAID Enclosure Technical Specifications	99
Table A-5	Fibre Channel RAID Deskside Tower Technical Specifications	100
Table A-6	Fibre Channel Rack Specifications	102
Table A-7	Cable Characteristics	104
Table A-8	Silicon Graphics Fibre Interface Cables	105

About This Guide

This manual describes the general function, configuration, and use of the Silicon Graphics fibre channel (also known as FC) options for Origin and Onyx2 systems. Fibre channel options for the OCTANE workstation are limited to FibreVaults only at the time this document was published.

This product consists of an XIO or PCI fibre channel controller board, copper cables or optional optical cables with media interface adapter (MIA) conversion modules, and FibreVault or fibre channel RAID storage enclosures. The fibre channel product also includes driver support software, a command-line interface (CLI), and a graphical user interface (GUI), for configuring and administering the fibre channel options. These options are documented in the *Origin FibreVault and Fibre Channel RAID Administrator's Guide*.

This guide introduces fibre channel architecture concepts as well as describing the available hardware. The spelling of fibre (as opposed to fiber) designates a communications and mass storage technology. Fibre channel components do not always connect with fiber (optical) cabling. Silicon Graphics fibre channel products come standard with copper cable connections. Optional optical cables are also available. The connection technology of fiber optic communication cables is a separate topic.

A fibre channel RAID expansion enclosure is a FibreVault that connects to the fibre channel RAID loop by way of an LCC cable(s). It uses fibre channel RAID disks only.

Audience

This guide is written primarily for users of a fibre channel system. Some fibre channel option components can be installed, removed, or replaced by the end user; this guide provides instructions for these procedures. This guide presumes the user has some familiarity with Origin or Onyx2 systems, but no special training is presumed. This guide also provides general information that may be of interest to system administrators.

Structure of This Document

This guide contains the following chapters and appendixes:

- Chapter 1, “Fibre Channel Features and Capabilities,” introduces the main hardware and software features of the fibre channel options; it compares FC with other solutions for interfacing to storage.
- Chapter 2, “Fibre Channel Architecture,” describes fibre channel architecture and topology in general, and Silicon Graphics fibre implementation more specifically.
- Chapter 3, “Fibre Channel Storage,” describes the fibre channel storage system, giving details on the Origin FibreVault and the Origin fibre channel rack. It explains the Origin FibreVault LEDs and describes how to shut down and power on the Origin FibreVault and the Origin fibre channel rack.
- Chapter 4, “Fibre Channel Component Replacement,” explains how to replace or add Origin fibre channel RAID or FibreVault owner-replaceable components.
- Chapter 5, “Configuring a Fibre Channel System,” explains how to configure an existing fibre channel system. It presents various system setups and shows ways to customize them.
- Appendix A, “Technical Specifications,” gives specifications for the fibre channel enclosures, desktside systems, and rack.

A glossary and an index complete this guide.

Additional Useful Manuals

Other manuals that may provide additional useful information include:

- *Origin FibreVault and Fibre Channel RAID Administrator's Guide* (p/n 007-3715-*nnn*)
- *Fibre Channel XIO and PCI Option Board Owner's Guide* (p/n 007-3633-*nnn*)
- *OCTANE Workstation Owner's Guide* (p/n 007-3435-*nnn*)

Typographical Conventions

These type conventions and symbols are used in this guide:

Helvetica Bold Hardware labels

Italics Executable names, filenames, IRIX commands, manual or book titles, new terms, program variables, tools, utilities, variable command-line arguments, variable coordinates, and variables to be supplied by the user in examples, code, and syntax statements

Bold Function names

Fixed-width type

Error messages, prompts, and onscreen text

Bold fixed-width type

User input, including keyboard keys (printing and nonprinting); literals supplied by the user in examples, code, and syntax statements

"" (Double quotation marks) Onscreen menu items and references in text to document section titles or chapter titles

[] (Brackets) Optional syntax statement arguments

Fibre Channel Features and Capabilities

This chapter introduces you to the Silicon Graphics fibre channel mass storage options available with your Origin2000 or Onyx2 host system.

This chapter includes the following main sections:

- “Fibre Channel Technology Introduction” on page 1
- “Networks, Channels, and Fibre Channel” on page 2
- “Applications” on page 5
- “Fibre Channel Host Interface Boards” on page 9
- “Fibre Channel Storage” on page 18
- “Fibre Channel Hardware Components” on page 18
- “FibreVault Options in Origin and Onyx2 Racks” on page 28
- “Fibre Channel Administrative Interface” on page 29

Fibre Channel Technology Introduction

Fibre channel is the general name of an integrated set of standards being developed by the American National Standards Institute (ANSI).¹ The fibre channel standard defines a high-speed data transfer interface that can be used to connect workstations, mainframes, supercomputers, storage devices, and displays. The fibre channel standard addresses the need for fast transfers (up to 1 Gbit per second) of large amounts of information. Currently, fibre channel’s main use is as an interface to storage.

¹ For information on the ANSI Fibre Channel standards, contact Global Engineering Documents, 15 Inverness Way East, Englewood, CO 80112 USA (303)-397-0271 or (800)-854-7179 (U.S. & Canada).

Conceived as a generic, efficient physical transport system that can support multiple protocols, the standard also relieves system manufacturers of the burden of supporting the variety of channels and networks currently in place, because it provides one standard for networking, storage, and data transfer. Note that the Silicon Graphics initial implementation is for communication with *mass storage* systems *only*.

The fibre channel provides a general transport vehicle for upper level protocols (ULPs), including the intelligent peripheral interface (IPI) and small computer system interface (SCSI) command sets; high-performance parallel interface (HIPPI) data framing; internet protocol (IP); and IEEE 802.2. Proprietary and other command sets can also use and share the fibre channel, although such use has not been defined as part of the fibre channel standard and is not supported by Silicon Graphics host systems. The Silicon Graphics implementation currently supports only the SCSI fibre channel protocol.

Networks, Channels, and Fibre Channel

In recent years, technical developments have precipitated a greater need for extremely fast data links. Performance improvements have spawned increasingly data-intensive and high-speed networking applications, such as multimedia and scientific visualization. However, the existing network interconnects between computers and I/O devices have been unable to run at the speeds needed. Fibre channel development has been aimed at practical, inexpensive, yet expandable means of quickly transferring data between workstations, mainframes, supercomputers, desktop computers, storage devices, displays, and other peripherals.

The two basic types of data communication between processors and peripherals have been *channels* and *networks*:

- A channel provides a direct or switched point-to-point connection between the communicating devices. A channel is typically hardware intensive and transports data at high speed with low overhead.
- A network is an aggregation of distributed nodes (workstations, file servers, or peripherals) with its own protocol that supports interaction among these nodes. A network has relatively high overhead because it is software-intensive; consequently, it is slower than a channel. Networks can handle a more extensive range of tasks than channels because they operate in an environment of unanticipated connections, whereas channels link only a few devices with predefined addresses.

Table 1-1 summarizes the differences between networks and channels.

Table 1-1 Networks and Channels

Network	Channel
Unanticipated connections; must be able to connect any pair of devices; addresses and routes are usually not predefined	Defined domain, usually a closed-system, small set of devices in close proximity; every device address is predefined
Software intensive; high protocol overhead	Hardware intensive; low overhead
Interconnects processors at moderate to high speeds using a form of switching over relatively large distances	Interconnects processors and high-performance peripherals at high speeds, relatively short point-to-point links
Uses various protocols for error detection, error correction, and data retransmission	Simple error correction or data retransmission on receipt of a busy signal; message is passed back to central processor if more complex error recovery is required

The fibre channel (FC) standard attempts to combine the simplicity and speed of channel communications with the flexibility and interconnectivity of protocol-based network communications. Despite the name fibre channel, the FC architecture represents neither a channel nor, for that matter, a real network topology.

Fibre channel ports can be connected in three topologies: as *point-to-point* links, in an *arbitrated loop*, or to a *fabric*, or switch. The Silicon Graphics fibre channel option topology is the arbitrated loop (FC-AL). Chapter 2, "Fibre Channel Architecture," explains this topology in detail. Some of the benefits of FC-AL are

- Low connectivity cost.
 - FC-AL supports up to 125 devices per single host adapter, reducing costs as well as cabinet or subsystem space. (Silicon Graphics supports up to 110 disks per fibre channel loop.)
 - Copper cabling is used for device connection within each subsystem (cabling one enclosure to another in a loop).
 - Host-to-storage enclosure connections can be copper or (optionally) fiber optic.
 - Fibre channel cables use fewer connection pins and conductors.

- High performance. FC-AL features make it well suited to both high-speed and transaction-processing applications:
 - The dual-ported fibre channel XIO interface delivers up to 200-MB-per-second peak data transfer rates.
 - FC-AL supports host-to-primary enclosure connection distances up to 25 meters using copper cable and up to 300 meters with optical cable options.
- Silicon Graphics FC implementation relies on dual-ported fibre disk drives that provide redundant FC data paths from the enclosure backplane to the drive.
- Reliability and ease of use. FC-AL devices are quick to install or reconfigure:
 - The high reliability is advantageous in densely populated cabinets or subsystems.
 - The FC-AL design enables higher data availability for today's systems and applications.
 - Widely available standard components are used (no jumpers are required to configure the drive).
 - Sophisticated error detection code provides better data integrity than traditional parity bit schemes.
- Support for SCSI protocols.
- Low latency for decreased overhead and more throughput.
- Storage. The Silicon Graphics implementation supports a rack with 11 enclosures and up to 110 disks total on a single FC-AL loop
- Data transfer. DC-balanced 8b/10b signals using odd or even disparity; variable-length data frames (maximum 2 KB) with complete error checking (32-bit CRC).
- Shortwave optical interface adapters (media interface adapters) make it possible using optional fibre optic cable, to place the storage up to 300 meters away from the host.

Applications

Fibre channel is useful for a variety of applications:

- Scientific graphics and video markets, using high-bandwidth rates with large I/O requests.
- File server and database markets, which must support large amounts of I/Os per second (IOPS) with relatively small random I/Os.

Some database applications also require access to vast amounts of storage, which can be in the terabytes range. High availability is also a requirement for much of this market. The Silicon Graphics fibre channel implementation addresses these needs by supporting up to 100 RAID disks in an IOPS environment on a single FC-AL.

For applications requiring large capacity with lower IOPS capability, the Silicon Graphics fibre channel implementation supports up to 110 non-RAID disks on a single FC-AL.

- The emerging video-on-demand market, which must provide many random data streams at a guaranteed data rate.

High-availability RAID is also a requirement for this market. Fibre channel RAID addresses these needs by supporting up to 100 RAID disks in a well-balanced access pattern on a single arbitrated loop. The arbitrated loop also provides an arbitration fairness scheme that prevents high-priority requests from starving low-priority requests.

Table 1-2 summarizes how fibre channel compares with other storage interfaces.

Table 1-2 Silicon Graphics Storage Interface Solutions

Feature	Fast-10 Wide SCSI-2	Fast-20 Wide (Ultra) SCSI	Silicon Graphics Gigabit Fibre Channel Implementation
Bandwidth	20 MBps	40 MBps	100 MBps
Duplex	Half	Half	Full
Distance (copper)	25 m	25 m	25 m (27.3 yards)
Distance (optical)	N/A	N/A	300 m (328 yards)
Connectivity	16	16	110 +1 FC-AL

Table 1-2 (continued) Silicon Graphics Storage Interface Solutions

Feature	Fast-10 Wide SCSI-2	Fast-20 Wide (Ultra) SCSI	Silicon Graphics Gigabit Fibre Channel Implementation
Administrative method for hot-plugging a disk	No	Yes	Yes
Fair arbitration	No	No	Yes
Optical support	No	No	Yes (optional)
Dual port	No	No	Yes (optional)
Conductors in cable	68	68	4 in copper or 2 in optical fiber

High Availability and RAID Applications

Fibre channel can address high availability in several ways. At the simplest level, the dual ports of one or more FC disks or disk enclosures can each be connected to a different host so that, if a host fails, the other host can still access the disks. Optional secondary fibre channel link control cards (LCCs) installed in the fibre enclosures provide an additional physical connection to the second port of each disk in the vault. High-availability loop configurations require the use of an optional FailSafe configuration with a fibre channel hub (or switch).

Note: Special optional IRIS FailSafe fibre channel software and hardware support is required to take advantage of maximum high-availability configurations. Customer-specific application software may also cover some high-availability requirements. High availability is a software and hardware option package; it is not a standard IRIX feature.

Optional redundant power supplies for FC disk enclosures that can be connected to separate power sources are available. They provide protection against power supply failure and some insurance against power circuit failure.

Each enclosure uses a fan pack with a redundant fan. The pack can be removed and replaced immediately without shutting off the disk enclosure.

RAID Data Redundancy

RAID technology provides redundant disk resources in disk-array configurations that make the storage system more highly available. A RAID array maintains parity or mirrored data that lets the disk group survive a single disk module failure without losing data.

RAID disk arrays continue operating even if one of the disk modules fail. When a disk module in the array does fail, the data is still available and the storage processor (SP) can *reconstruct* it from the surviving disk(s) in the array. The RAID subsystem begins reconstructing data as soon as

- A hot spare (dedicated replacement disk module) is available
- The failed disk module is replaced with a new disk module

If a disk module has been configured (bound) as a hot spare (see “RAID Hot Spare” on page 8), it is available as a replacement for a failed disk module immediately. Reconstruction begins as soon as the failure is detected. The SP automatically writes to the hot spare and rebuilds the group using the information stored on the surviving disks. Performance is degraded while the SP rebuilds the data and parity on the new module. However, the storage system continues to function, giving users access to all data, including data stored on the failed module.

Similarly, when a new disk module is inserted to replace a failed one, the SP automatically writes to it and rebuilds the group using the information stored on the surviving disks. The length of the rebuild period, during which the SP re-creates the second image after a failure, can be specified when RAID levels are set and disks are bound into RAID units. These processes are explained in the *Origin FibreVault and Fibre Channel RAID Administrator’s Guide*.

Supported RAID Arrays

The Silicon Graphics fibre channel RAID product line supports most types of RAID disk groups. The *Origin FibreVault and Fibre Channel RAID Administrator's Guide* defines and describes these supported configurations.

Note: RAID functionality is dependent on the release level of your hardware, firmware, and software. Upgrades may be required to support certain RAID functions. Contact your service provider for information regarding upgrades to your fibre channel system. See your system release notes for any changes in the RAID support levels.

RAID Hot Spare

A hot spare is a dedicated replacement disk unit on which users cannot store information. The hot spare is an additional disk added to the enclosure that already has a RAID disk array installed. The capacity of a disk module that you bind as a hot spare must always be at least as great as the capacity of the largest disk module it might replace.

If any single disk in a fibre channel RAID group fails, the SP automatically begins rebuilding the failed disk module's structure on the hot spare. When the SP finishes rebuilding, the disk group functions as usual, using the hot spare instead of the failed disk. When you replace the failed disk, the SP starts copying the data from the former hot spare onto the replacement disk. When the copy is done, the disk group consists of disk modules in the original slots, and the SP automatically frees the hot spare to serve as a hot spare again.

A hot spare is most useful when you need the highest data availability. Note that the SP finishes rebuilding the disk module onto the hot spare before it begins copying data to a replacement disk, even if you replace the failed disk during the rebuild process.

Online Administrative Disk Swapping

The Silicon Graphics fibre channel implementation allows you to shut down and remove individual disk modules in the non-RAID FibreVault using a command-line interface. This administrative “warm swapping” of disk modules must be executed in a specific sequence to guarantee data integrity.

A single RAID disk module can be removed from an array and replaced without special administrative procedures. This procedure should be done only with full knowledge of the status of the LUN that the drive is part of.

See the *Origin FibreVault and Fibre Channel RAID Administrator's Guide* for more information. See “Adding or Replacing a Disk Module” in Chapter 4 for mechanical information on removing or replacing a fibre disk.

Caution: Pulling and replacing (hot swapping) a drive without implementing the proper administrative sequence could result in data loss if the drives are improperly removed or replaced.

Fibre Channel Host Interface Boards

Two types of fibre channel interface boards are available to plug into host systems. The XIO and PCI interface boards each use the same standard or optional cable connections. The following sections describe their features and options:

- “Fibre Channel XIO Board” on page 10
- “Fibre Channel PCI Board” on page 12
- “Fibre Channel Board Connectors” on page 13
- “Fibre Channel Board Logic” on page 13
- “Copper or Optical Media Interface Adapter” on page 16
- “Chassis Grounding Issues” on page 17

Fibre Channel XIO Board

The Silicon Graphics host system can use an XIO-based fiber controller board for fibre channel communication. Figure 1-1 shows the XIO fibre channel controller in the overall Origin2000 and Onyx2 I/O structure.

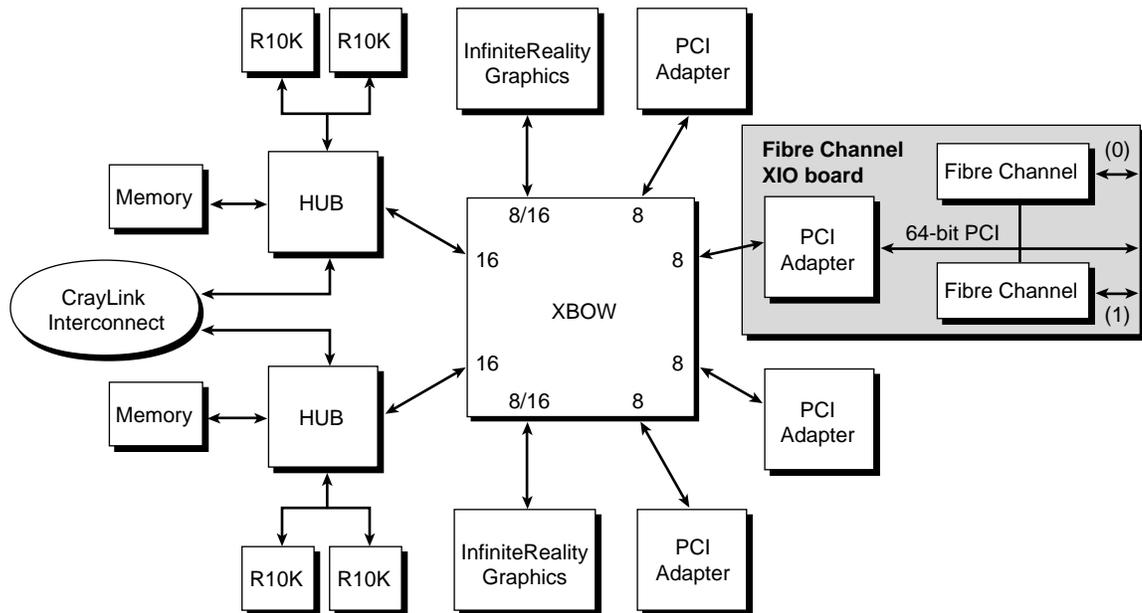


Figure 1-1 Fibre Channel in the Overall Origin2000 and Onyx2 I/O Structure

Each XIO fibre board has two connectors, and each connector can be used to control up to 110 fibre disks in a fibre rack. Figure 1-2 shows the fibre channel XIO board and connectors. The 9-pin connector farthest from the locking handle is channel 0; the connector closest to the handle is channel 1.

Note: The XIO boards used in the OCTANE workstation use a different hook locking mechanism than the one shown in Figure 1-2. Do not install a fibre channel XIO board from an Origin2000 or Onyx2 into an OCTANE workstation. The board may not come out and partial disassembly of the OCTANE workstation may be required to remove it. The boards are functionally identical in all other respects.

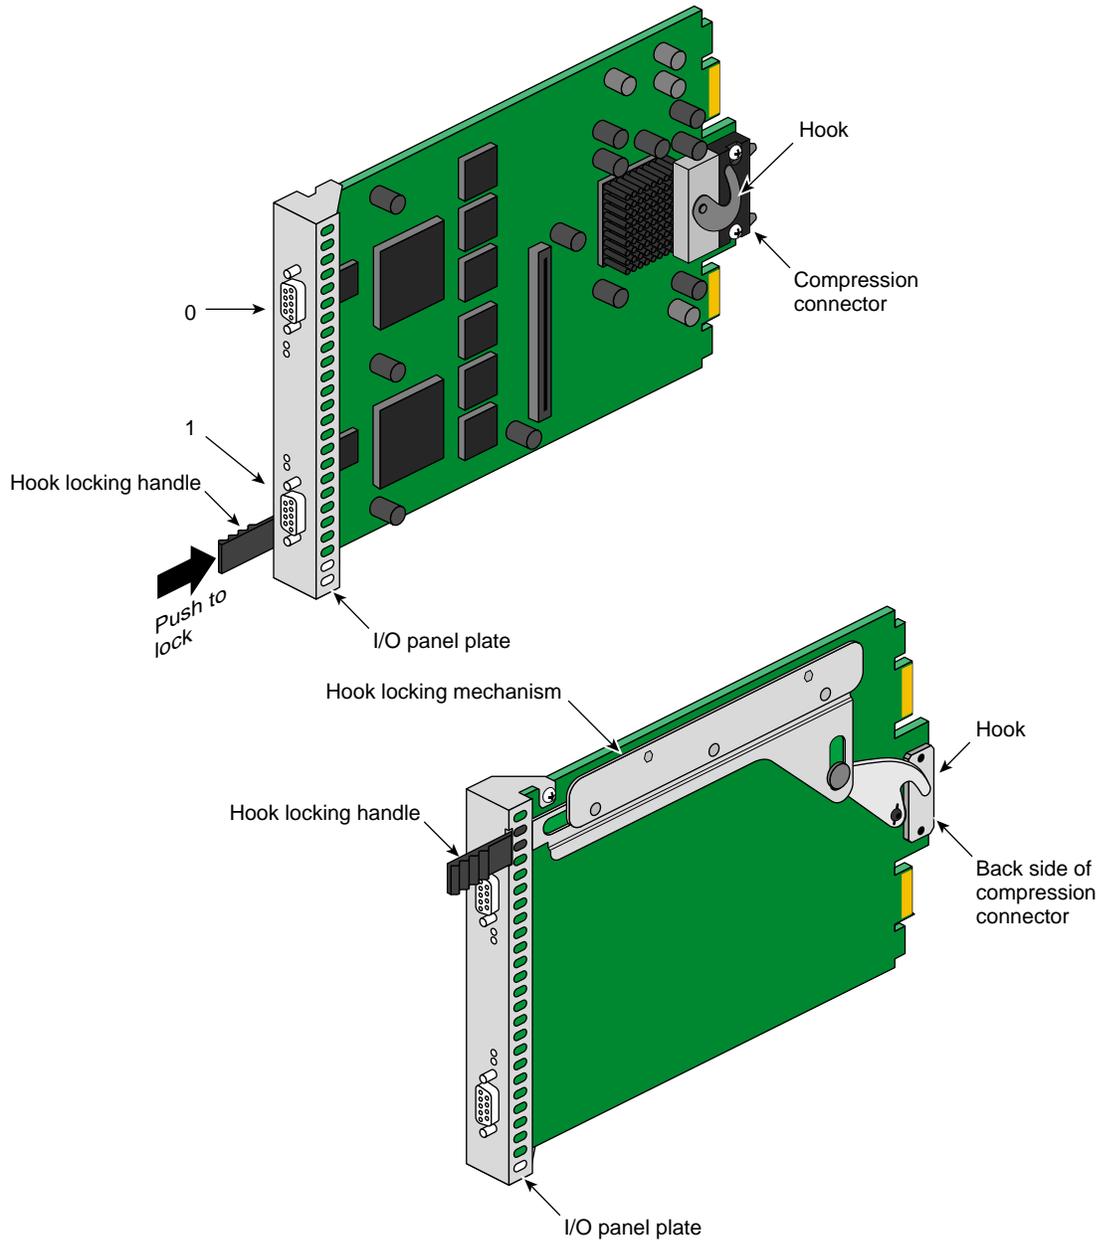


Figure 1-2 Fibre Channel XIO Board and Connectors

Fibre Channel PCI Board

The fibre channel PCI board is a half-size PCI option board that provides a single-connection high-performance FC-AL interface between the host system and the fibre enclosure(s).

This board can be inserted in an Origin200 chassis or in the optional PCI module of an Origin2000, Onyx2 system, or OCTANE workstation.

- For an Origin2000 or Onyx2 system, only a qualified Silicon Graphics service engineer should install the board. The board installs in the optional PCI module in XIO slot 2, or by using an optional XIO to PCI adapter board (p/n 030-1275-xxx).
- Customers who have an Origin200 server should follow the installation instructions in the latest version of the *Origin200 and Origin200 GIGACHannel Owner's Guide* (p/n 007-3708-*nnn*).
- Customers with an OCTANE workstation must follow the installation instructions in the *OCTANE PCI Module Installation Guide* (p/n 007-3547-*nnn*) or *OCTANE Workstation Owner's Guide*.

Figure 1-3 shows the fibre channel PCI board.

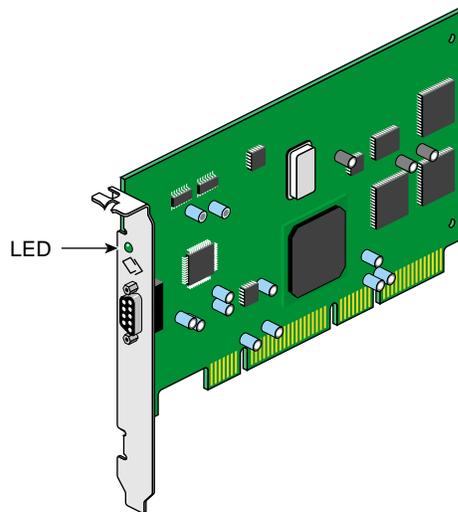


Figure 1-3 Fibre Channel PCI Interface Board

Fibre Channel Board Connectors

Each fibre channel XIO board has two 9-pin connectors. The fibre channel PCI interface board has a single 9-pin connector. Note that each fibre channel loop uses one fibre channel connector.

Figure 1-4 shows the fibre channel connector pin assignments.

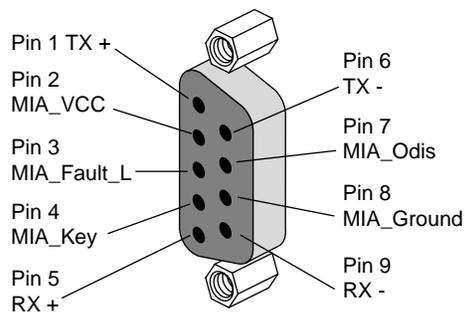


Figure 1-4 XIO Board Connector Pin Assignments

Fibre Channel Board Logic

The fibre channel option board provides bus-protocol conversion between an 8-bit (400-MB-per-second) XIO interconnect link and a 64-bit PCI interface. This conversion is provided by the fibre channel board's "bridge" interface ASIC, see Figure 1-5. The second (PCI-to-fibre-channel) conversion is provided by the PCI-64 fibre channel ASIC. The board also provides an interface between a little endian 64-bit PCI interface and a 1.062-Gb fibre channel interface.

The FC option board supports fibre channel Class 3 operations as a loop port (L_Port). The firmware supports Class 3 and FC-AL transfers only. For details on ports, Class 3, and FC-AL, see Chapter 2, "Fibre Channel Architecture."

If the fibre channel XIO board loses power or the physical fibre channel connection is broken, the link that the board is attached to becomes inoperable. High-availability loop configurations require the use of an optional FailSafe configuration with a fibre channel hub (or switch).

Transceivers (SerDes) embedded on the board convert 10-bit parallel data using 8b/10b encoding to differential serial signals. They also provide frame synchronization, word alignment, and clock recovery for incoming serial data.

For incoming serial data, two recovered clocks at 53.125 MHz for odd or even bytes are provided as outputs on two pins, 180 degrees out of phase. Any required equalization to compensate for high-frequency losses for copper cables (by attenuating the lower frequencies to match) is supplied externally to the fibre channel option board. Serial data in both directions between the transformer and the external connector is AC coupled using a capacitor.

Figure 1-5 is a fibre channel XIO card block diagram.

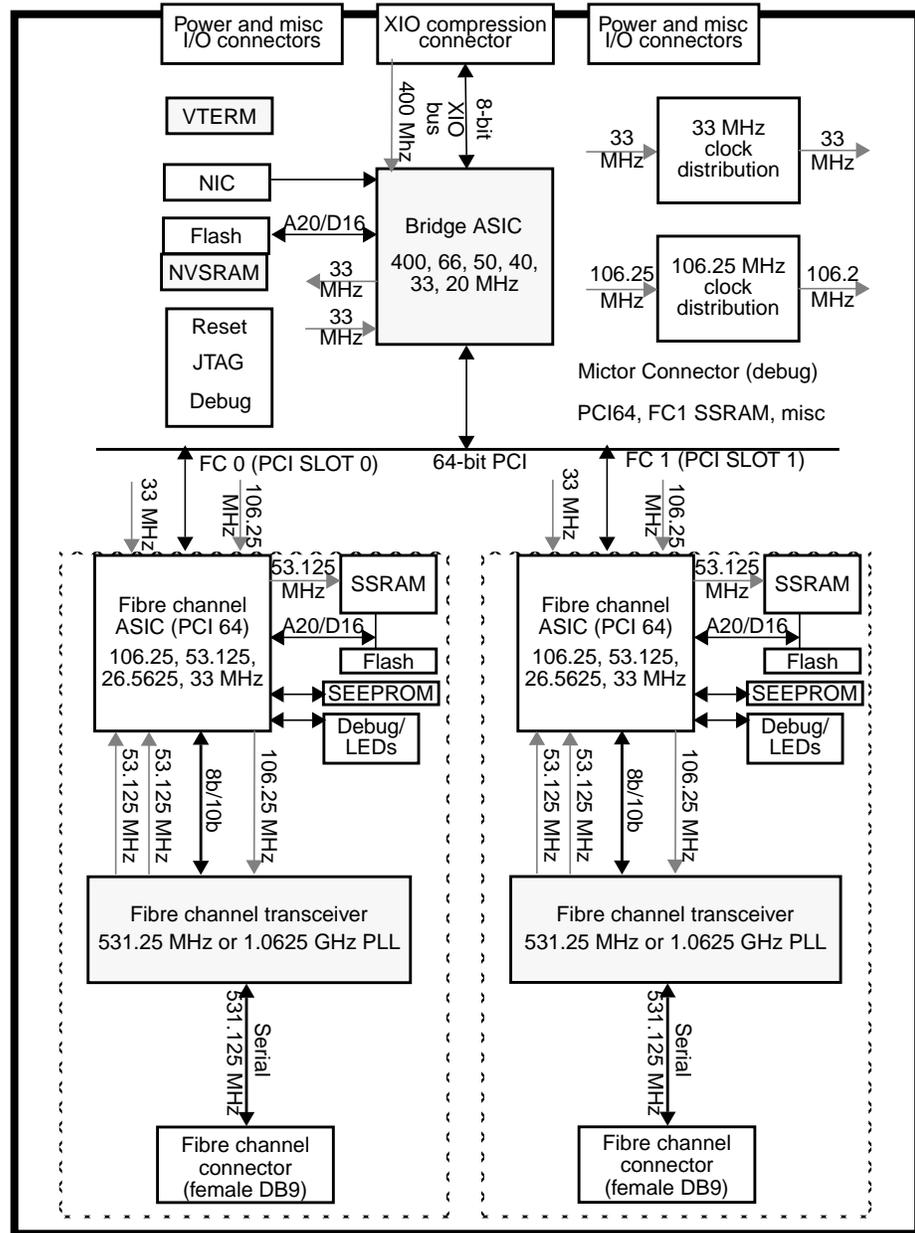


Figure 1-5 Fibre Channel XIO Card Block Diagram

Copper or Optical Media Interface Adapter

Both a copper and an optical interface are supported. The copper interface has a female DB9 connector. Fibre channel arbitrated loop (FC-AL) disk drives have a copper native interface.

For distances up to 300 meters, an optional fiber-optic interface option using a media interface adapter (MIA) module is available. A hot-pluggable MIA external module with shortwave CD-ROM lasers provides the physical interface. The MIA is a full-duplex module that converts photons to electrons in one direction and converts electrons to photons in the other direction. An industry-standard duplex SC connector, shown in Figure 1-6, supplies the external fiber-optic connection on the optional MIA.

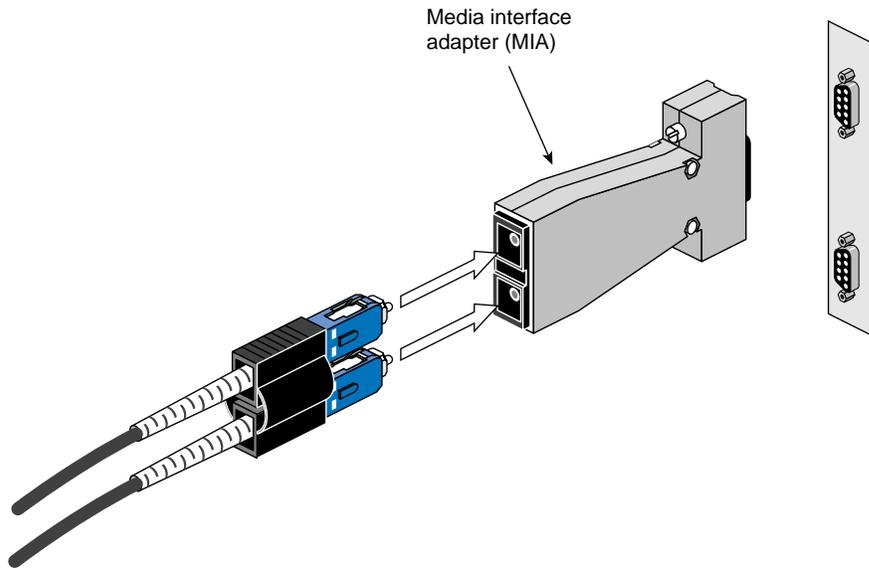


Figure 1-6 Optical Module (MIA)

Chassis Grounding Issues

Grounding issues are very important in Origin and Onyx2 systems. Each chassis must be well grounded through its power connector. All chassis connected by copper cables must share the same transformer, be grounded through the same earthing rod, and be on the same branch circuit. If you have any doubts about the quality of the ground connection, consult with a qualified electrician. Using an optical cable between the fibre enclosure(s) and the host connection eliminates any problems related to common grounding.

Caution: Any difference in ground potential greater than 500 millivolts (0.5 volts) between two chassis connected by copper cables can cause severe equipment damage and create hazardous conditions.

The branch circuit wiring should be provided with an insulated grounding conductor that is identical in size, insulation material, and thickness to the earthed and unearthed branch-circuit supply conductors. The grounding conductor should be green, with or without one or more yellow stripes. This grounding or earthing conductor should be connected to earth at the service equipment or, if supplied by a separately derived system, at the supply transformer or motor-generator set. The power receptacles in the vicinity of the systems should all be of an earthing type, and the grounding or earthing conductors serving these receptacles should be connected to earth at the service equipment.

Fibre Channel Storage

Fibre channel makes it possible to configure an array along a single interface instead of across many separate interfaces, as is done with SCSI storage. Since the drives constituting the FC array unit are organized along an interface, this arrangement is sometimes described as a longitudinal array.

Chapter 4, “Fibre Channel Component Replacement,” is an end-user description of the installation and replacement of FC-AL storage disks. Up to five fibre channel RAID modules can be installed in a rack (each requires a separate loop connection). Fibre channel RAID enclosures use the same XIO fibre channel controller boards as non-RAID enclosure and disk combinations.

Fibre Channel Hardware Components

This section introduces

- “FC-AL Disk Modules” on page 18
- “FibreVault and Fibre Channel RAID Enclosures” on page 21
- “The Fibre Channel Rack” on page 24
- “Fibre Channel Deskside Systems” on page 26

FC-AL Disk Modules

Disk modules used in non-RAID fibre channel enclosures can be removed or added using an administrative procedure for taking an individual disk module offline. Sectors on non-RAID drives are 512 bytes, and RAID disks use 520 byte sectors. The two types of drives are *not* interchangeable. Single RAID disks can be removed without special administrative procedures. Always follow the guidelines listed in the *Origin FibreVault and Fibre Channel RAID Administrator’s Guide*.

Figure 1-7 shows features of the FC-AL disk module.

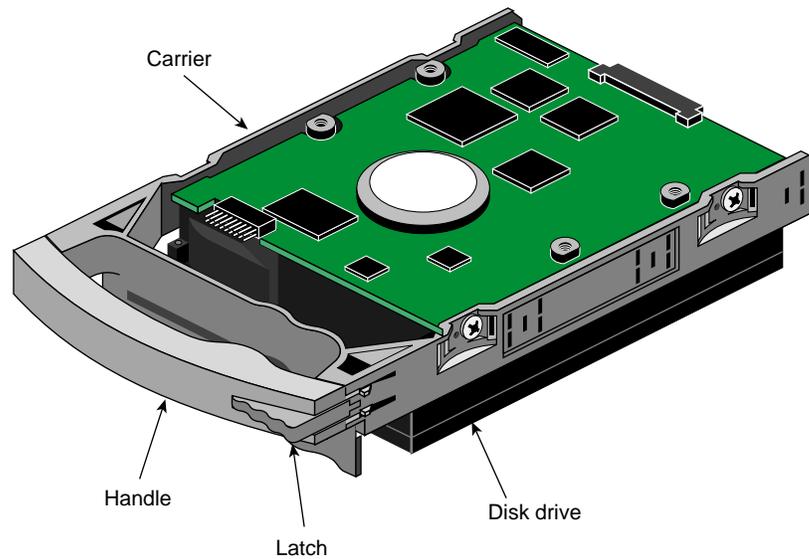


Figure 1-7 FC-AL Disk Module

Each disk module consists of one 3.5-inch FC-AL disk drive in a carrier. The disk drive carrier provides smooth, reliable mating with the chassis slot guides and midplane connectors. The latch on the carrier handle holds the disk module in place to ensure proper connection with the midplane. You must push the latch to remove a disk module from the slot. You can insert the drive module in only one direction. All Silicon Graphics fibre disk drives are dual-port capable.

Always confirm the type, capacity (in gigabytes) and speed of a fibre drive (RAID or non-RAID), before you install it. Each drive module has an identifying sticker with its part number. Table 1-3 provides descriptions and part numbers for the fibre drives.

Table 1-3 Fibre Channel Disk Drive Descriptions

Fibre Drive Type	Drive Capacity	Drive Speed	Sector Sizing	Part Number
Non-RAID	9.1 GB	7200 RPM	512-byte sectors	9470140
RAID	8.8 GB	7200 RPM	520-byte sectors	9470192
Non-RAID	9.1 GB	10,000 RPM	512-byte sectors	9470282
RAID	8.8 GB	10,000 RPM	520-byte sectors	9470281
Non-RAID	18 GB	7200 RPM	512-byte sectors	9470255
RAID	17.8 GB	7200 RPM	520-byte sectors	9470257

Note: Always consult with your service provider before replacing a disk with another that has a different capacity or speed rating (different part number).

Be aware of the following information when adding or replacing disks:

- RAID and non-RAID fibre drives are not compatible and may not be mixed in the same enclosure.
- A lower capacity RAID drive must not be used as a replacement disk or hot spare in an array of higher capacity disks. For example, you must never install an 8.8 GB RAID drive as a hot spare or replacement in an enclosure of 17.8 GB disks.
- Installing a lower speed (RPM) drive in an array of faster RPM disks causes the entire array bandwidth (read and write speed) to slow to the lower speed capacity.

Caution: The disk modules have no electrostatic discharge (ESD) shielding; do not stack them on top of each other. Use a grounding strap whenever handling the drives, and keep them in antistatic containers when not in use. Use care when moving the drives; do not drop or jar them.

FibreVault and Fibre Channel RAID Enclosures

The FibreVault and fibre channel RAID enclosures are high-performance, high-capacity, disk-array subsystems using an FC-AL as an interconnect interface. These modular and scalable designs with optional high-availability features make expansion easy when storage needs increase. The fibre channel RAID expansion is a FibreVault using RAID disk modules.

Each enclosure contains from one to ten disk modules, plus other components. The FibreVault is available in a rackmount or desktower enclosure. Figure 1-8 shows an example of the rackmountable FibreVault enclosure.

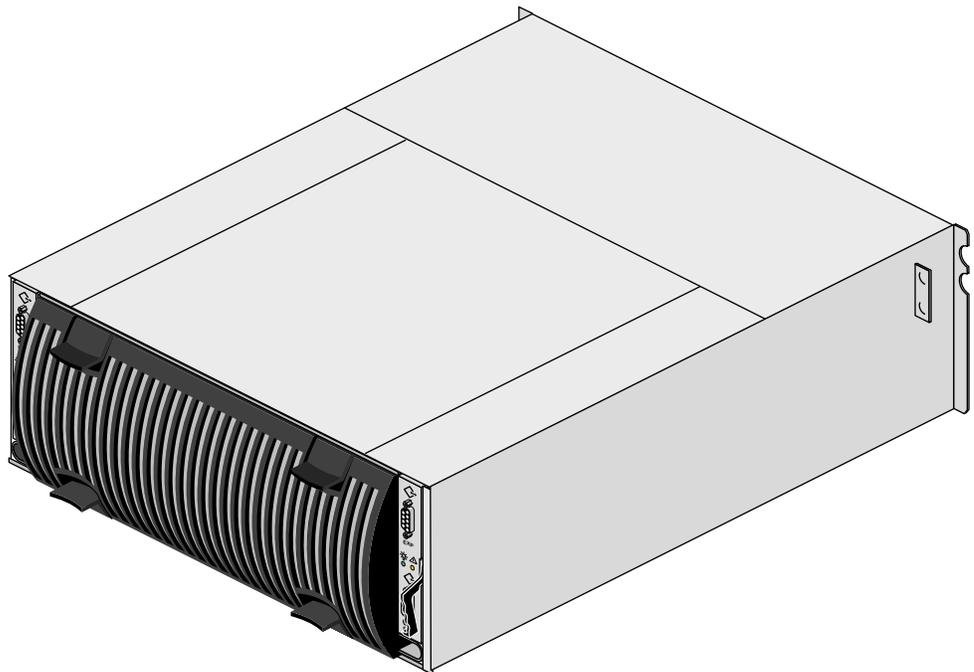


Figure 1-8 FibreVault Enclosure (Rear View)

The front door, which has a built-in electromagnetic interference (EMI) shield, must be closed for the FibreVault or RAID enclosures to be EMI compliant. When the door is open, you can remove or install drive modules and change the enclosure address. Chapter 3, "Fibre Channel Storage," has complete details of enclosure operation; Chapter 4, "Fibre Channel Component Replacement," contains complete procedures for replacing and adding disk modules and other owner-replaceable components.

The key for the locking latch fits any FibreVault or fibre channel RAID door (see Figure 1-9). Figure 1-9 shows an example of the fibre channel RAID enclosure.

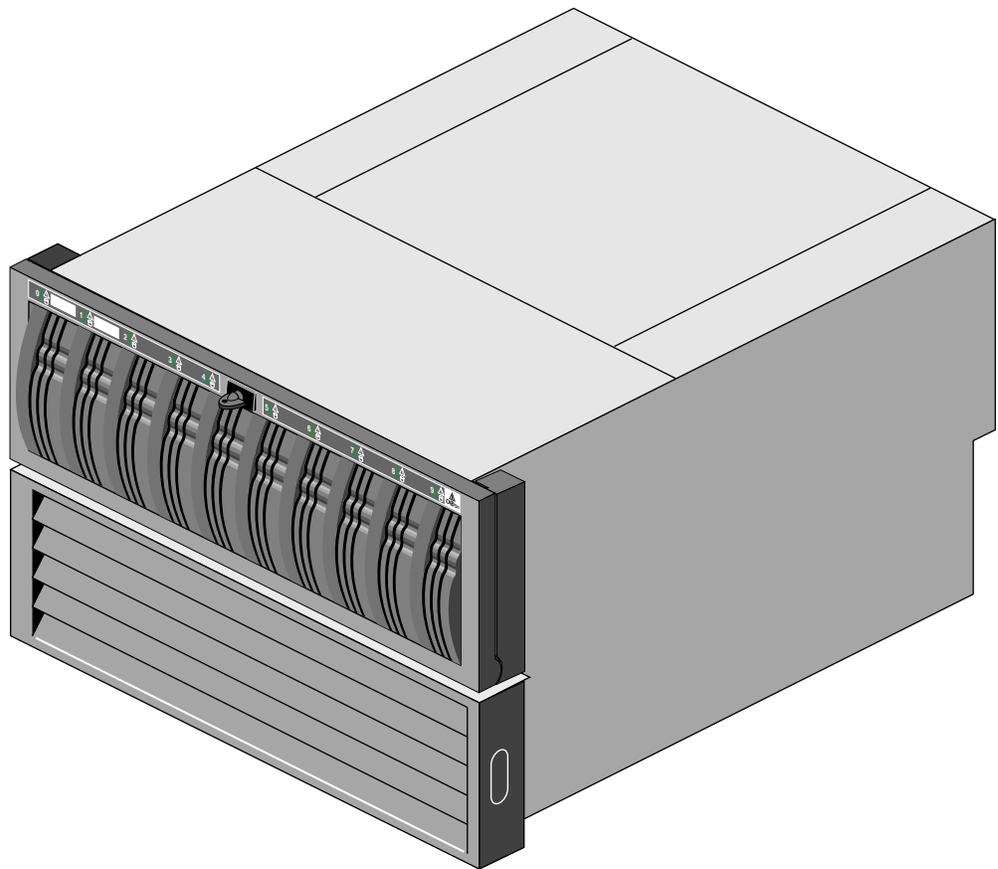


Figure 1-9 Fibre Channel RAID Disk Enclosure (Front View)

For proper cooling and normal operation, each disk slot in a fibre enclosure must contain either an FC-AL disk module or a disk module filler. Likewise, each LCC and power supply position must have either the actual component or a filler panel installed.

Fibre channel enclosure features include the following:

- Two LED status indicators for each disk module slot:
 - Check (amber): If it blinks, a cooling fault occurred; if it stays on, another type of component fault has occurred.
 - Activity (green): The enclosure is powered on.
- One link controller card (LCC); an optional second LCC requires special application software for high-availability system configurations.
- One autoranging power supply, optional second power supply; separate power cords.
- Fan assembly containing three fans.
- Address range (16-bit switch); LEDs to indicate address.
- Disk slot bypass feature for use when a disk module has failed.
- DB9 connectors and 0.3-meter (11.8 inches) “twinax” cable for connection to another enclosure using the same FC-AL. FibreVault options in Origin racks use 0.5 meter (19.7 inch) enclosure connection cables.
- Automatic loopback of signals on the last enclosure in a chain, eliminating the need for a loopback connector or terminator.

Disk modules, redundant LCCs and power supplies, and the fan assembly can be replaced with no tools while the enclosures are powered on. An optional additional power supply and LCC and the multiple fans in the fan assembly provide high-availability configurations. If two LCCs and their cabling are installed, the disk drives support dual-port FC-AL interconnects.

Note: Dual-port (dual-LCC) configurations require IRIS FailSafe or a customer specific software application to use these features fully. The standard IRIX operating system does not currently support high-availability dual-port FC-AL configurations.

Up to 11 non-RAID FibreVault enclosures (or three RAID enclosures) can be mounted in a fibre channel rack. The deskside non-RAID chassis supports one FibreVault enclosure. The fibre channel RAID deskside holds one RAID (DPE) enclosure and one FibreVault RAID expansion (DAE). Enclosures are connected to a host fibre controller board with 10-meter or optional 25-meter copper cable or an optional 25-, 100-, or 300-meter optical cable used with a media interface adapter (MIA) at each end. FibreVaults used in Origin and Onyx2 system racks use a 3-meter (10 foot) copper cable to connect the host's controller board to the fibre channel enclosure.

The Fibre Channel Rack

The 19-inch fibre channel rack houses up to 11 Origin FibreVault enclosures, mounted horizontally. Figure 1-10 shows the fibre channel rack.

For further details on fibre enclosures, see Chapter 3, "Fibre Channel Storage," for more discussion and additional illustrations.

Cables are included for connecting enclosures together in a rack to form one large disk storage loop.

For further details on FibreVault, see Chapter 3, "Fibre Channel Storage."

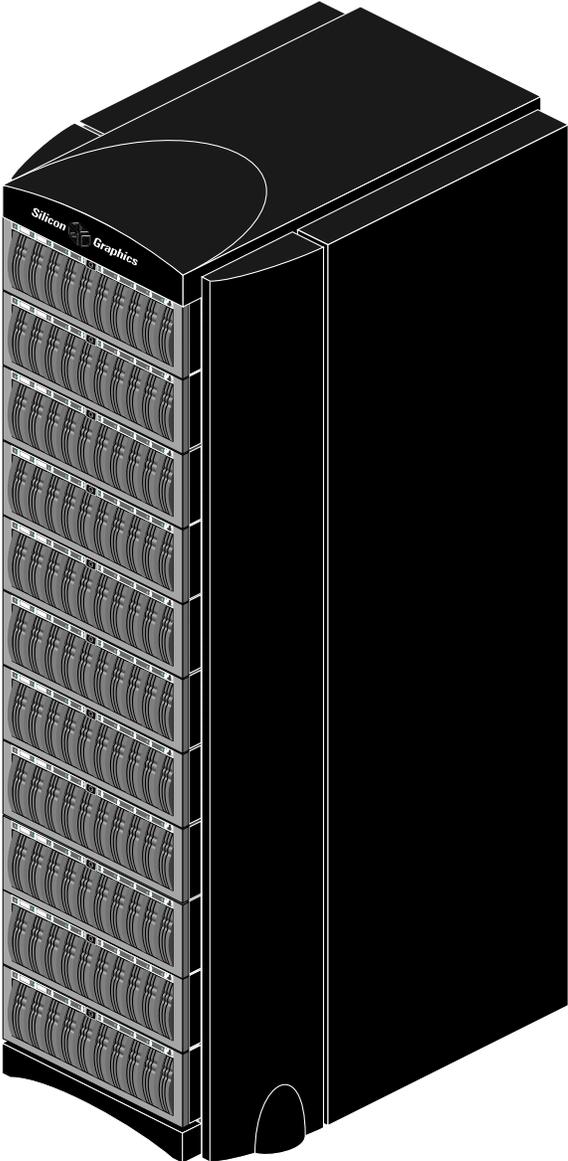


Figure 1-10 Fibre Channel Rack (Front View)

Fibre Channel Deskside Systems

Two deskside fibre channel “tower” systems are available:

- The deskside FibreVault tower system holds one non-RAID fibre enclosure in a vertical orientation.
- The deskside fibre channel RAID tower system holds one fibre channel RAID enclosure and one FibreVault RAID expansion enclosure mounted vertically.

Figure 1-11 shows the two types of fibre channel deskside systems.

Note that the fibre channel enclosures that mount in the fibre channel rack and deskside towers are exactly the same. The mounting orientation (vertical vs. horizontal) is the only difference. All internal components remove, replace, and operate in the same manner.

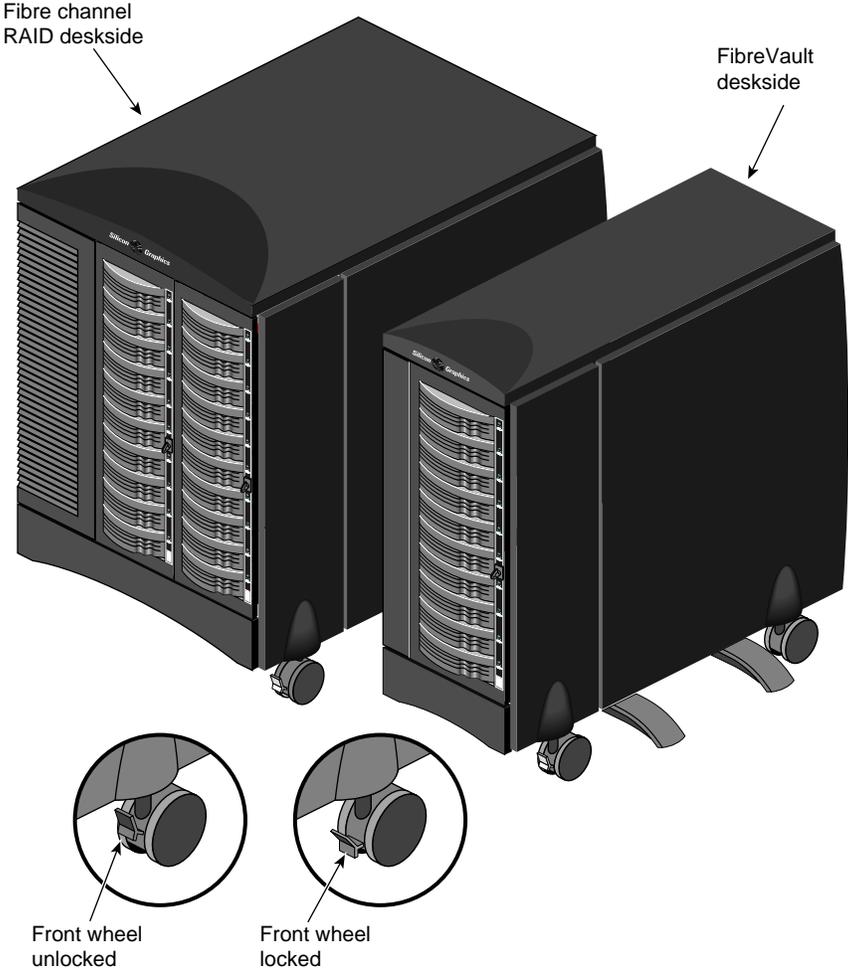


Figure 1-11 Fibre Channel Deskside Storage Systems

FibreVault Options in Origin and Onyx2 Racks

The non-RAID FibreVault enclosure can be mounted in the Origin or Onyx2 rack. The FibreVaults in Origin racks do not have redundant power options. The FibreVault enclosures install in the Origin or Onyx2 rack in the same location(s) and take up approximately the same space as the Origin Vault SCSI disk and tape unit.

A maximum of three FibreVaults may be installed in the Origin rack, one in the Onyx2 rack. All enclosure installations are done by Silicon Graphics trained and approved installers.

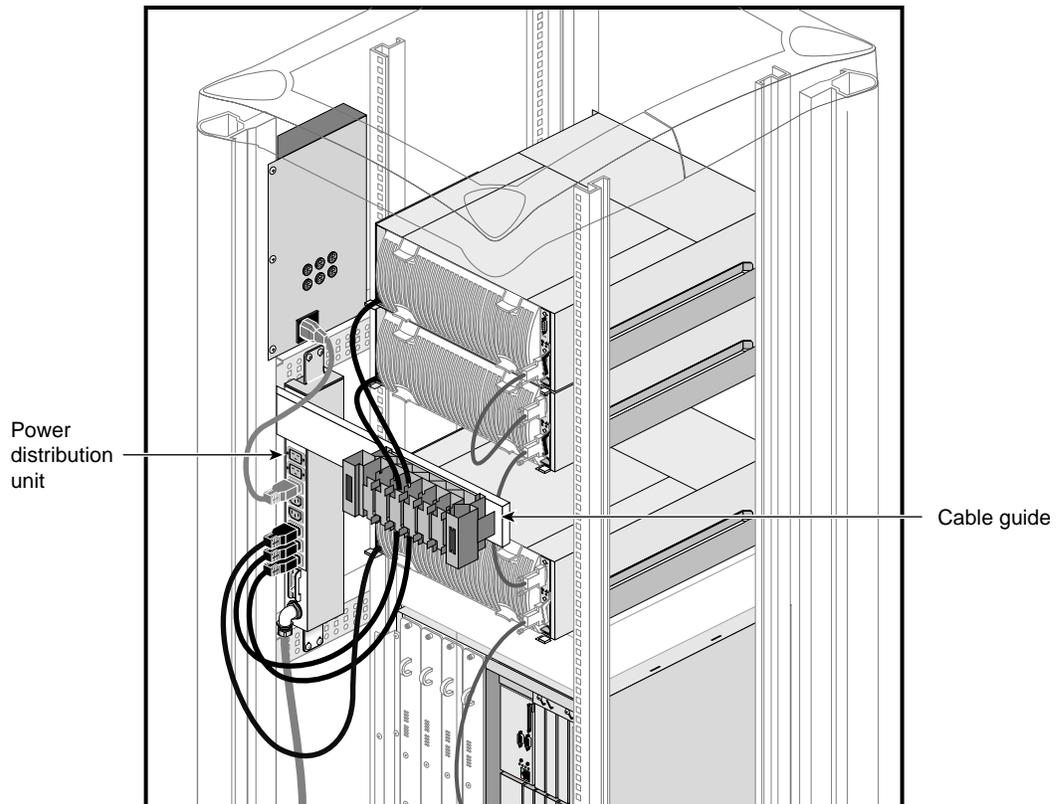


Figure 1-12 FibreVault Enclosures in an Origin Rack (Rear View)

Fibre Channel Administrative Interface

Normal reads and writes go through the *dksc* driver; configuration and environmental status use the *devscsi* driver.

For fibre channel RAID applications, an agent (*ssmagent*) provides the interface to the *devscsi* driver for configuration and environmental status. The agent can be configured to talk to a dedicated logical unit number (LUN) on the RAID. On top of the agent is a graphical user interface (GUI) and a command-line interface (CLI), which can be used for configuration of both RAID and non-RAID fibre arrays.

Figure 1-13 diagrams the software interface. The software driver implementation follows the private loop direct attach (PLDA) profile (see Figure 1-13).

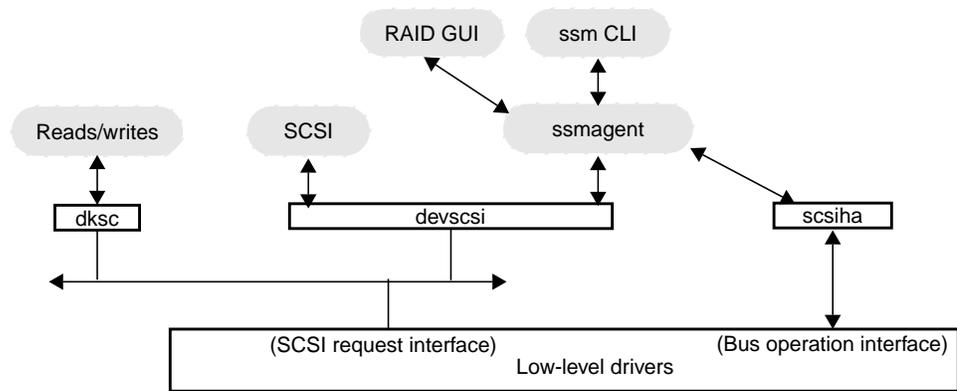


Figure 1-13 Fibre Channel Software Interface

For more information on using the administrative CLI and GUI, see *Origin FibreVault and Fibre Channel RAID Administrator's Guide*.

Chapter 5, "Configuring a Fibre Channel System," explains how to use some of the commands for specific purposes.

High-Level Drivers

The *dksc* and *devscsi* drivers are the same as those for previous implementations of SCSI. The existing SCSI request (high-to low-level) interface is unchanged, but extensions have been added:

- *dksc* driver: The function of this driver is essentially unchanged from its SCSI-only ancestor. It interfaces to both the SCSI and fibre channel low-level drivers and implements access to basic disk I/O functions.
- *devscsi* driver: Like its SCSI-only ancestor, this driver provides an interface to both the SCSI and fibre channel low-level drivers and implements SCSI pass-through command functionality. The *devscsi* driver can send arbitrary SCSI commands to a fibre channel disk just as arbitrary SCSI commands are sent to a SCSI disk; for example, a sequence of SCSI commands for downloading new firmware to a SCSI device.
- *scsiha*: The driver *scsiha* allows access to all fibre channel-specific functions not implemented by *dksc* and *devscsi* for both device-targeted and controller-targeted commands. This interface is used for commands that address a fibre channel loop as a whole, and are not addressed to any specific device or do not require a response from a device. Examples are the Loop Initialization command (LIP) or an enable PBE command to a specific device.

The low-level driver provides resources for up to 1,000 concurrent outstanding commands.

Controller Addressing

Each fibre channel interface (XIO controller board) address is chosen during loop initialization. This ensures that it does not conflict with other devices (including other controllers).

Each controller tries to get ID 125. If ID 125 is already taken, the controller takes the next-highest ID number available. This is automatic, so no system hostid is required.

Fibre Channel Architecture

This chapter includes the following main sections:

- “Networks, Channels, and Fibre Channel” on page 31
- “Silicon Graphics Fibre Channel Ports and Topologies” on page 33
- “Data Transfer in Fibre Channel” on page 36
- “Fibre Channel Layers” on page 37

Networks, Channels, and Fibre Channel

The two most common peripheral protocols for device communication in the computer industry are networks and channels.

Networks

- Involve I/O interfaces that usually support many small transactions with relatively high overhead due to software involvement in the flow of information
- Allow a host or device in the network to communicate with any other device
- Operate in an open, unstructured, and unpredictable environment

Channels

- Supply peripheral I/O interface to a host and transport large amounts of data between the host and peripherals
- Keep data processing overhead to a minimum by handling data transfer in hardware, with little or no software involvement once an I/O operation begins
- Operate in a closed, structured, and predictable environment where all devices that can communicate with a host are known in advance and any change requires host software or configuration table changes

Fibre channel attempts to combine the best of these two methods into an I/O interface that meets the needs of both channel users and network users. Fibre channel communications can be conducted over copper coax, twisted pair, or optical fiber. Note that Silicon Graphics currently supports only copper coax, with optical cable and a media interface adapter (MIA) as an option.

Fibre channel technology provides three different interconnect topologies to serve the combined needs of channel and network usages. These three topologies are:

- Fabric
- Point-to-point
- Arbitrated loop

Currently, Silicon Graphics supports only arbitrated loop technology, (see “Silicon Graphics Fibre Channel Ports and Topologies” on page 33).

Fabric Topology

A fabric topology permits dynamic interconnections between nodes through ports connected to the fabric. This is a standard network configuration. In the future, it may be possible to carry both network traffic and disk I/O over the same wire.

Point-to-Point Topology

In point-to-point host applications, two ports are used, connected to a link. The transmitter of each port is connected directly to the receiver of the opposite port. This topology limits the number of connections that can be made across the wire.

Arbitrated Loop Topology

The arbitrated loop topology used by Silicon Graphics is called fibre channel arbitrated loop (FC-AL). In this topology, each port arbitrates for access to the loop. Ports that “lose” the arbitration act as repeaters of all traffic on the loop. The loop contains a dedicated transmit channel and a dedicated receive channel that are clad together into one cable to form a loop out and back. This protocol allows up to 127 ports connected in a serial loop (one FL_Port and 126 NL_Ports). Silicon Graphics supports a maximum of 110 ports in a single rack.

Ports are called N_Ports (Node_Ports), NL_Ports (Node_Loop Ports), F_Ports (Fabric_Ports), or FL_Ports (Fabric_Loop Ports).

FC-AL ports represent a superset of functions and must work correctly with all three topologies (fabric, point-to-point, and arbitrated loop).

An NL_Port represents each disk in a disk array. Each NL_Port sees all messages and passes messages not addressed to that port. Ports passing messages are said to be in "repeat mode".

An FL_Port provides the connection from the fibre channel loop to the I/O board on the host, (in this case XIO or PCI). In the future, FL_Ports may provide a link from an arbitrated loop to a fabric.

Silicon Graphics Fibre Channel Ports and Topologies

Silicon Graphics fibre channel topology uses an *arbitrated loop* (FC-AL), as diagrammed in Figure 2-1.¹

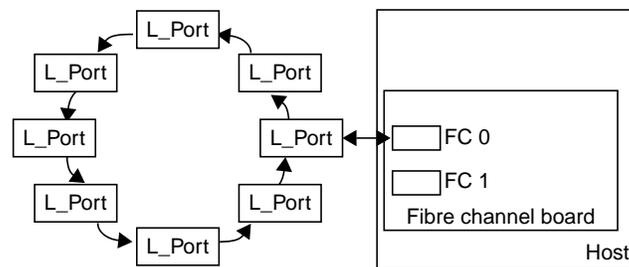


Figure 2-1 Fibre Channel Arbitrated Loop (FC-AL)

FC-AL is a shared-bandwidth distributed topology with arbitration fairness. The Silicon Graphics FC-AL option potentially supports 126 active L_Ports. Current support is limited to a maximum of 110 disks within a single rack. Communication on a loop is between two devices at a time. All nonactive FC-AL ports act as repeaters; all devices on the loop run at the same interface speed.

¹ Silicon Graphics host systems support a specific subset of available fibre channel technology.

The arbitrated loop topology is a distributed topology in which each port includes the minimum necessary functions to establish a circuit. The ports are all L_Ports.

Disk L_Ports use the point-to-point protocol to create a point-to-point circuit between two L_Ports.

Port arbitration begins with one port replacing fill words between frames with an arbitration primitive called "ARBx." If it goes completely around the loop and returns to the sending port, then that port has won arbitration.

A port that wins arbitration stops retransmitting received data and opens the loop between its receiver and transmitter by sending out the open primitive, OPNyx, where "y" represents the physical address of the destination port (AL_PD=Arbitrated Loop Physical Destination) and "x" represents the sending port. See Figure 2-2 for an example of loop functions.

Closing a loop is done by transmitting the close primitive (CLS). The partner port finishes its work and retransmits the CLS, and then the loop is available for arbitration again.

The FC-AL loop interface is an implementation of SCSI-3 architecture that provides high-performance, greater connectivity, and high-availability features. FC-AL was designed as a low-cost method of connecting multiple hosts and storage devices without requiring the use of switches and fabrics. Note that multi-host configurations require special optional software and hardware; ask your Silicon Graphics sales or service representative for information.

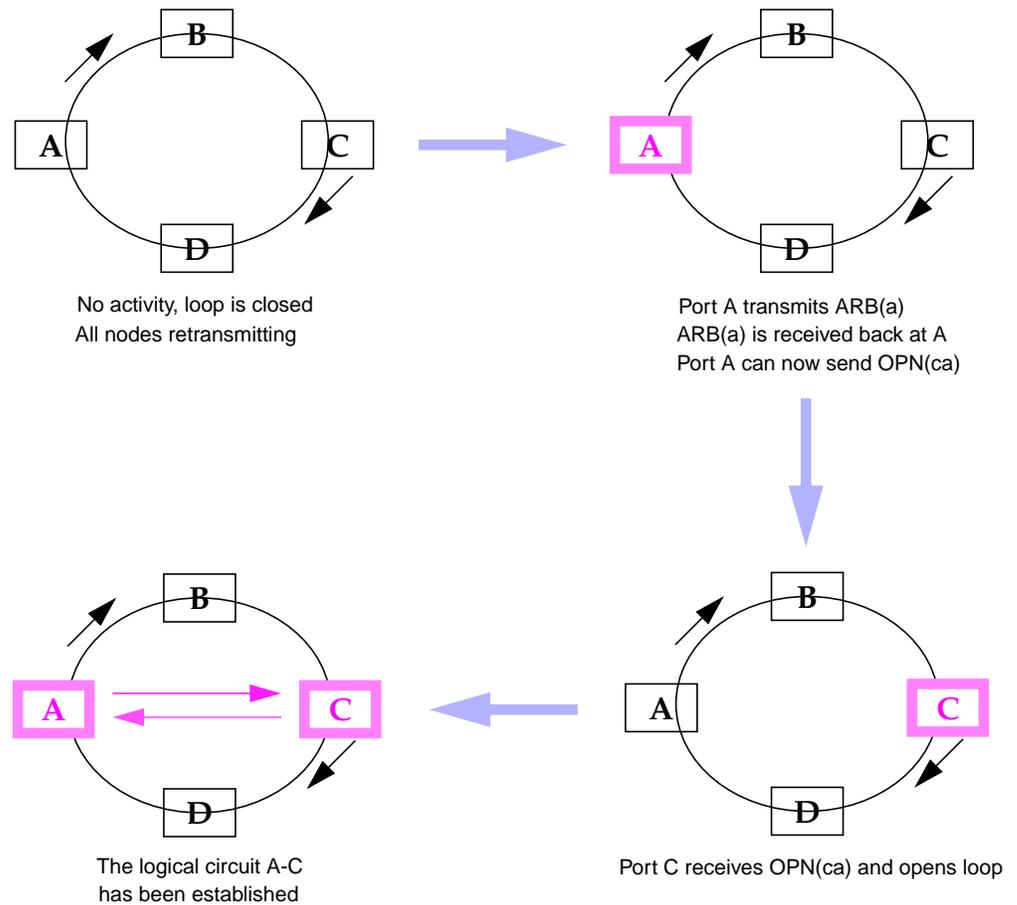


Figure 2-2 Port Arbitration Activity Example

Data Transfer in Fibre Channel

Data transfers in fibre channel use the following parameters

- DC-balanced 8b/10b signals with odd or even disparity
- variable-length data frames (maximum 2 KB)
- 32-bit CRC on frames

Combining channel and LAN characteristics, fibre channel uses buffers, one at the source node (port) and one at the destination. Each buffer can be any size.

FC layers 0, 1, and 2 move data from one buffer to another without regard to the format or meaning of data, thus avoiding the overhead required for handling different network communications protocols. FC layers 0, 1, and 2 provide only control of the transfer and simple error detection, relying on the content of the message header to trigger actions such as routing data to the proper buffer. Except for how FC formats serial data for transmission, protocols are irrelevant; communication protocols between devices can be whatever those devices require. The Silicon Graphics fibre implementation uses SCSI Framing Protocol (FCP) for FC-AL communications between devices and the host system.

Information can flow between two ports—actually between their buffers—in both directions simultaneously. The mechanism coordinating the flow of information between two N_Ports is an *exchange*.

Data transfer takes place in units called *frames*, each a maximum of 2048 bytes of data. The frame contains the information to be transmitted, the address of the source and destination ports, and link control information. Frames are of two types, Data frames and Link_Control frames. Data frames can be used as Link_Data frames and Device_Data frames. A set of related frames for one operation is called a *sequence*.

There are no limits on the size of a transfer. Frame sizes are transparent to upper-level software because the unit of transfer is a sequence. A specific layer of fibre channel (FC-2, described in “FC-2” later in this chapter) is responsible for breaking a sequence into the frame size that was negotiated between ports. This layer of FC also detects transmission errors using 32-bit CRC.

Fibre Channel Layers

Fibre channel is a layered architecture with five layers: FC-0, FC-1, FC-2, FC-3, and FC-4. Figure 2-3 diagrams the relationship between FC layers and OSI layers.

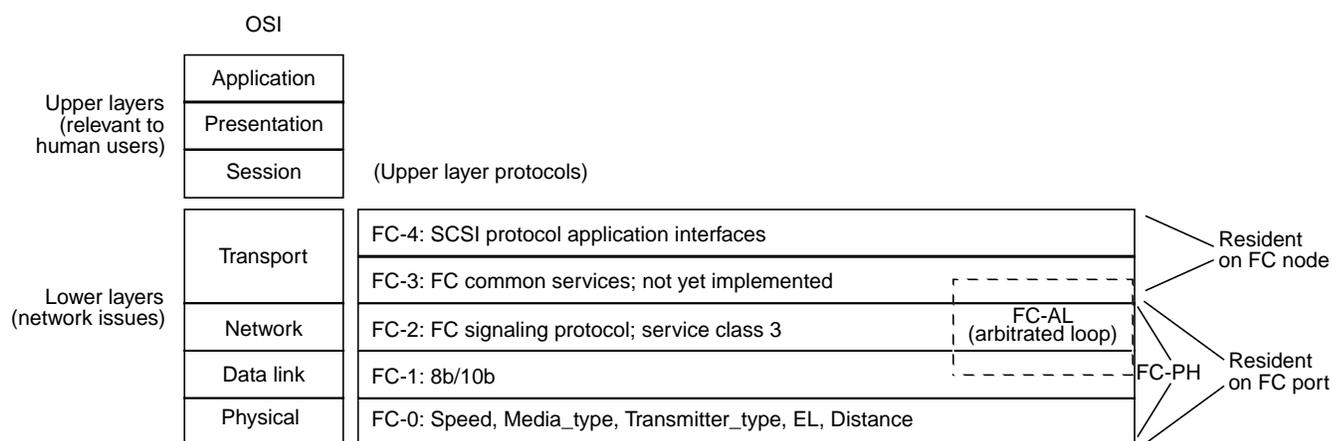


Figure 2-3 Fibre Channel Layers and OSI Layers

FC-PH is the FC physical and signaling interface, revision 4.1, defined in the FC-PH standard (ANSI X4.3). FC-PH is made up of FC-0, FC-1, and FC-2. Each FC layer is described and explained in the following separate sections:

- “FC-0” on page 37
- “FC-1” on page 38
- “FC-2” on page 38
- “FC-3” on page 42
- “FC-4” on page 42

FC-0

FC-0 defines the physical interface (media), or link, encompassing a wide variety of media, speed, and distance combinations. It defines the physical link in the system, including the fibre, connectors, and optical and electrical parameters for a variety of data rates. Table 2-1 summarizes the Silicon Graphics implementation for FC-0.

Note: In FC, there is no change in protocol between copper and fiber optic.

Table 2-1 Silicon Graphics Fibre Channel Cable Characteristics

Parameter	Copper Cable	Fiber Optic Cable Option
Value	100-TW-EL-S full-duplex (quad conductor)	100-M5-SL/N-I
Connector	Male DB9	SC duplex connector
Media type	150 ohm copper	62.5-micron multimode fiber
Speed	(1.0626 Gbaud) physical interfaces	1.0626 Gbaud
Transmitter type	PECL	780 nm optical laser
Distance	Maximum 25 meters from controller to enclosure	Maximum 300 m from controller to enclosure with optical cable

Caution: To maintain acceptable quality and signal integrity, use only Silicon Graphics supplied copper and optical cables with your fibre enclosures.

FC-1

FC-1 defines the transmission protocol, including serial encoding and decoding rules, special characters, and error control. The information transmitted over a fiber is encoded 8 bits at a time into a 10-bit transmission character. The transmission code, which must be DC balanced to support the electrical requirements of the receiving units, improves the transmission characteristic of information. The transmission characters ensure that short-run lengths and enough transitions are present in the serial bitstream to make clock recovery possible.

FC-2

FC-2, the signaling protocol, is the transport mechanism of fibre channel. It defines the framing rules of the data to be transferred between ports, the different mechanisms for controlling service classes, and the means of managing the sequence of a data transfer.

The standard defines these building blocks to aid in the transport of data across the link:

- “Ordered Set” on page 39
- “Frame” on page 40
- “Sequence” on page 40
- “Exchange” on page 40
- “Protocol” on page 41

Ordered Set

An ordered set is a four-byte transmission word containing data and special characters that have a special meaning. Ordered sets provide the availability to obtain bit and word synchronization, which also establishes word boundary alignment. An ordered set always begins with the special character K28.5. Three major types of ordered sets are defined by the signaling protocol:

- Frame delimiters: start-of-frame (SOF) and end-of-frame (EOF), which immediately precede or follow the contents of a frame.
- Two primitive signals (primitives), idle and receiver ready (R_RDY):
 - Idle indicates an operational port facility ready for frame transmission and reception.
 - R_RDY indicates that the interface buffer is available for receiving further frames.

The primitives are transmitted and repeated continuously to indicate specific conditions within a port or conditions encountered by the receiver logic of a port.

- Four primitive sequences:
 - Offline (OLS)
 - Not Operational (NOS)
 - Link Reset (LR)
 - Link Reset Response (LRR)

When a primitive sequence is received and recognized, a corresponding primitive sequence or idle is transmitted in response. Recognition of a primitive sequence requires consecutive detection of three instances of the same ordered set.

Frame

The basic building block of an FC connection is the frame, which contains the information to be transmitted (payload), the address of the source and destination ports, and link control information. Frames are broadly categorized as data frames and Link_control frames. Data frames can be used as Link_Data frames and Device_Data frames. Link_control frames are classified as acknowledge (ACK) and Link_Response (busy and reject) frames.

Each frame begins and ends with a frame delimiter. The frame header immediately follows the SOF delimiter. The frame header controls link applications and device protocol transfers and detects missing or out-of-order frames. An optional header can contain additional link control information. A field up to 2112 bytes (payload) contains the information to be transferred from a source N_Port to a destination N_Port. The four-byte cyclic redundancy check (CRC), which precedes the EOF delimiter, detects transmission errors.

Sequence

A sequence is formed by a set of one or more related frames transmitted unidirectionally from one port to another. Each frame in a sequence is uniquely numbered with a sequence count. Error recovery, controlled by an upper protocol layer is usually performed at sequence boundaries. For more detail, see “Data Transfer in Fibre Channel” on page 36 in this chapter.

Exchange

An exchange is composed of one or more nonconcurrent sequences for a single operation. The exchanges can be unidirectional or bidirectional between two ports. Within a single exchange, only one sequence can be active at any one time, but sequences of different exchanges can be active concurrently.

Protocol

The protocols are related to the services offered by

- Primitive sequence protocols for link failure
- N_Port login protocol: before performing data transfer, an N_Port interchanges its service parameters with another N_Port
- N_Port logout protocol: performed when an N_Port requests removal of its service parameters from the other N_Port (used to free resources at the connected N_Port)
- Data transfer protocol: methods of transferring upper-layer protocol (ULP) data using FC flow control management

Flow control paces the flow of frames between N_Ports and between an N_Port and the fabric to prevent overrun at the receiver. Class 3 uses only buffer-to-buffer flow control. Flow control is managed by the sequence initiator (source) and sequence recipient (destination) ports using Credit and Credit_CNT:

- Credit is the number of buffers allocated to a transmitting port.
- Credit_CNT represents the number of data frames that have not been acknowledged by the sequence recipient.

Buffer-to-buffer flow control is managed between an N_Port and an F_Port or between N_Ports in point-to-point topology. Each port is responsible for managing BB_Credit_CNT. BB_Credit is established during the fabric login. The sequence recipient (destination) port signals by sending a Receiver_Ready primitive signal to the transmitting port when it has free receive buffers for the incoming frames. Frame delivery is not confirmed.

Silicon Graphics fibre channel provides for service class 3. A connectionless service (“hopeful” or “datagram”), Class 3 is the most practical for connections requiring a lot of time. Class 3 has no end-to-end flow control or delivery acknowledgments. A datagram is the transfer of data without first establishing that data will be delivered. Timeouts are used to decide whether information is lost and retransmission required.

FC-3

FC-3 defines the common services provided by FC-PH to the ULPs. These services, which include striping, hunt groups, and multicast, are not yet implemented.

FC-4

FC-4, the upper-layer protocol (ULP) interface, defines the ULP-specific mapping protocols that provide interfaces between FC-PH and the ULPs. Silicon Graphics FC supports only the SCSI-3 packetized protocol for SCSI.

Fibre Channel Storage

Mass storage for the Silicon Graphics fibre channel option is provided by fibre channel arbitrated loop (FC-AL) disk drive modules. FC-AL disk modules have different interface cards from standard SCSI-2 disk modules, with 40-pin (instead of 80-pin) connectors. The disk modules run the SCSI-3 protocol and support Class 3 (connectionless, or datagram) service.

Up to ten disk modules can be contained in each fibre enclosure. Up to 11 Origin FibreVault enclosures can be contained in a fibre channel rack, for a maximum of 110 (non-RAID) fibre drives. A single fibre channel RAID loop supports a lesser number of drives. Technically, a maximum of five fibre channel RAID enclosures can fit in the rack. To provide room for expansion disk enclosures and other options, the recommended maximum is three enclosures in the rack.

Although each fibre channel option card can technically address up to 126 disk drives, there is no support for extending the fibre cable to another rack or cabinet. Therefore, the maximum amount of drives supported by a single fibre channel connection (loop) on a fibre board is equal to the maximum number (110) of drives installed in a single rack.

The FC-AL disk modules and the server chassis SCSI disk modules are not interchangeable and cannot be substituted for each other. Likewise, RAID and non-RAID fibre disk modules use different sectoring sizes and are not interchangeable.

This chapter includes the following main sections:

- “The FibreVault Enclosures” on page 44
- “FibreVault Components” on page 47
- “Fibre Channel RAID Enclosure Components” on page 52
- “Fibre Channel Rack” on page 58
- “Optional Standby Power Supply (SPS)” on page 59
- “Monitoring Storage System Status” on page 61
- “Shutting Down the Fibre Channel Rack and Disk Enclosures” on page 62

The FibreVault Enclosures

The non-RAID FibreVaults come in desktide (tower) and rackmount versions. Figure 3-1 shows the rear components on a rackmount FibreVault.

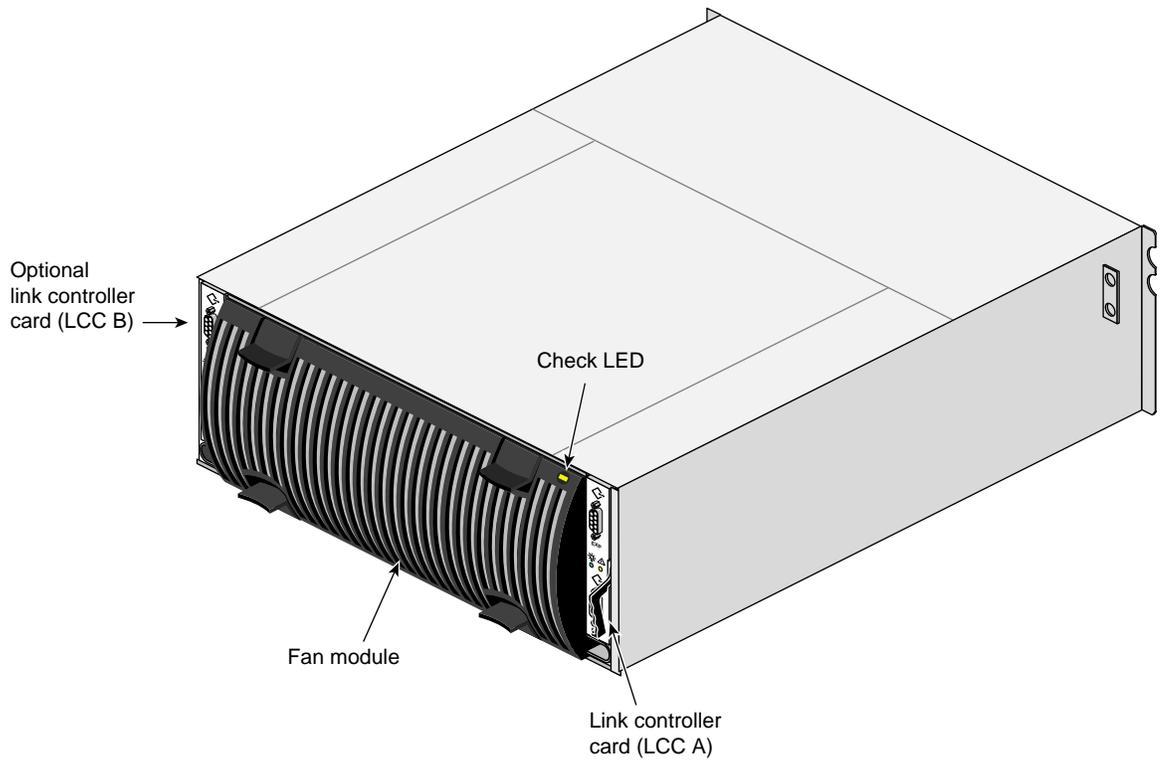


Figure 3-1 FibreVault Back View with Fan Module

Figure 3-2 shows the desktide FibreVault enclosure with its locking wheels and anti-tip “legs.” You can lock the front wheels by pushing the locking lever on each front wheel down. Lift the locks up to allow the system to roll.

Caution: The desktide tower’s anti-tip legs should always stay attached to the enclosure. The legs stabilize the tower unit and prevent it from tipping over.

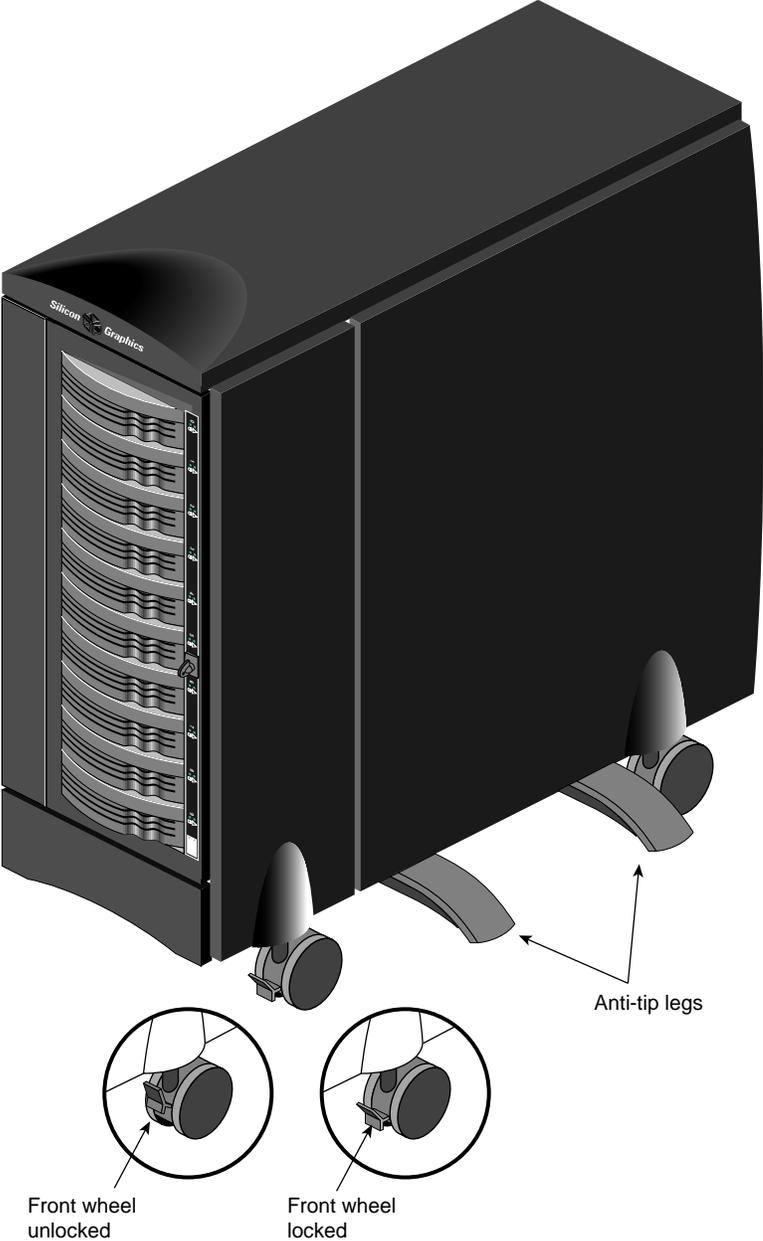


Figure 3-2 FibreVault Deskside (Tower) Unit

Figure 3-3 shows a rear view of the rackmount FibreVault components (fan module removed).

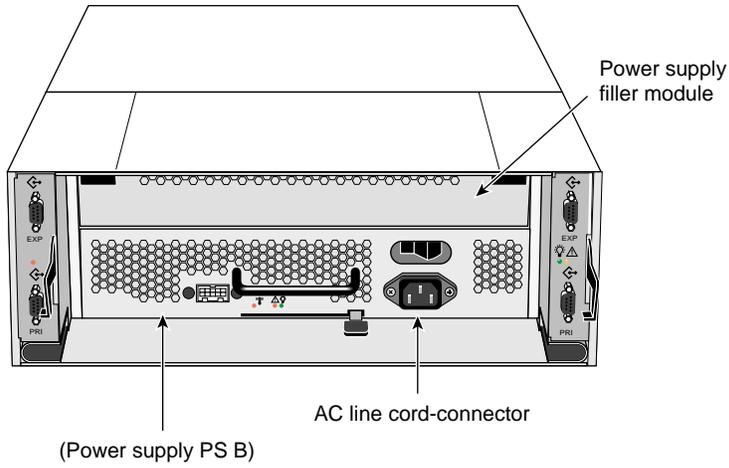


Figure 3-3 Rackmount FibreVault Enclosure (Fan Module Removed)

Figure 3-4 shows the FibreVault enclosure front panel.

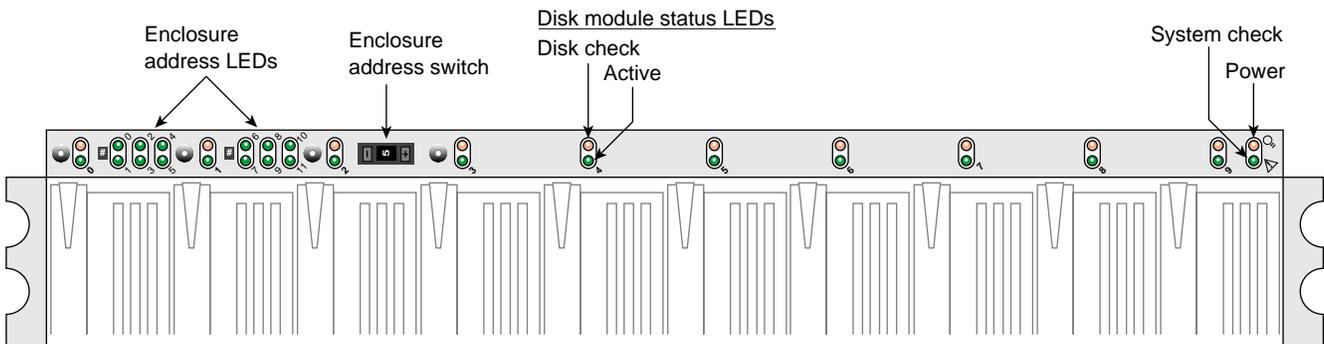


Figure 3-4 Origin FibreVault Enclosure Front Panel

FibreVault Components

This section explains more about the FibreVault enclosure's components:

- "Link Controller Card (LCC)" on page 47
- "Disk Module Fan Assembly" on page 49
- "The Enclosure Power Supplies" on page 50
- "The FibreVault Midplane" on page 52

The link controller card (LCC), disk cooling fan assembly, and enclosure power supply described in the next three sections are generic components used in both the FibreVault and fibre channel RAID expansion enclosures. The only difference between a FibreVault and a fibre channel RAID expansion enclosure is the sector size in the disk modules.

The listed components perform the same function in both types of enclosures and are completely interchangeable. The fibre channel RAID expansion enclosure is exactly the same mechanically as the FibreVault enclosure.

Note: Never exchange disk modules between a fibre channel RAID expansion enclosure and a FibreVault.

Link Controller Card (LCC)

The link controller card (LCC) controls data flow to and from the disk modules in FibreVault. The LCC also controls data flow between the FibreVault enclosure and the outside world (including additional linked enclosures) and monitors the enclosure environment.

The FibreVault enclosure comes standard with one LCC. An optional second LCC is required for high-availability or FailSafe operations. Note that special high-availability support software is also required. The LCC is visible at the back of the enclosure, as was shown in Figure 3-3. Figure 3-5 shows features of the LCC front panel.

Each LCC supports and controls one FC-AL. You can configure a FibreVault enclosure with one or two LCCs to provide one or two FC-ALs.

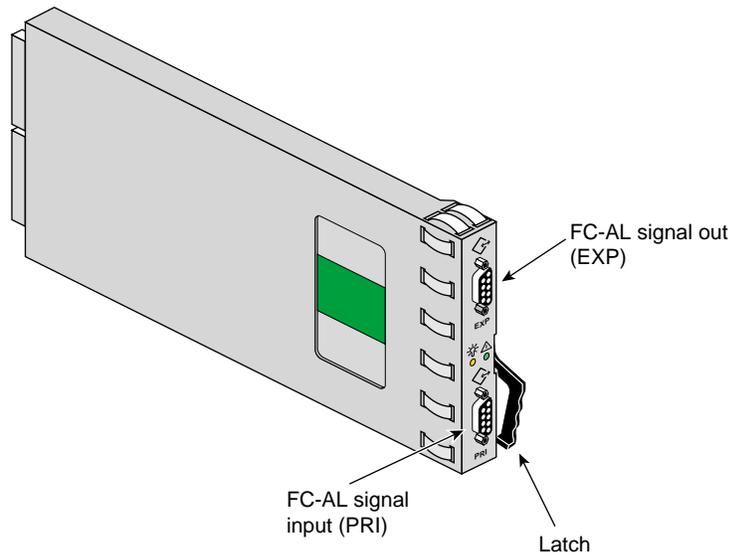


Figure 3-5 FibreVault Link Controller Card (LCC)

The LCCs in an enclosure are connected to host bus adapters in a server or to LCCs on other enclosures. The LCC receives and terminates the incoming FC-AL signal from either the server’s host bus adapter or from another enclosure. It passes the input signal to the disk drives in the enclosure, and then drives the output signal through the cabling to the next fibre enclosure in the loop. The FC-AL cabling is not explicitly configured as a loop (that is, a long return from the last fibre enclosure to the server), but instead as a full-duplex point-to-point connection, with the last fibre enclosure in the chain closing the loop on its LCC.

The fibre channel RAID enclosure receives data from the host through the system processor (SP) port and passes the loop through the expansion (**EXP**) port on its LCC to any fibre channel RAID expansion enclosure(s). The FibreVault LCC’s primary port (**PRI**) receives data from a host. Its expansion port (**EXP**) is used for sending data to the next in a series of linked FibreVaults.

The LCC status LEDs function as follows:

- active (green): The LCC is operating.
- check (amber): If the amber LED stays on, a fault occurred either in the LCC or an FC-AL connection.

LCC's have two physical methods of communicating status information or receiving control information from the host (or RAID processor) and target enclosure.

- A serial channel is duplexed over fiber cable in the RAID configuration. This permits communication with additional RAID enclosures with or without drives being present even when the loop is not being used.
- In non-RAID (FibreVault) configurations, the host communicates through the disk interface using the SFF 8067 protocol defined in the fibre channel standards.

The protocols that the LCC uses enable the host to poll fibre enclosure status and send commands that control the port bypass circuits on both LCCs, the enclosure, and the disk-module check lights. LCCs do not communicate with or control each other.

Disk Module Fan Assembly

Figure 3-6 shows features of the rear fan assembly used on the enclosures.

The fan assembly, which contains three fans, cools the entire FibreVault or fibre channel RAID expansion enclosure. Fibre channel RAID enclosures use an additional (SP) fan assembly that is located on the front of the fibre channel RAID enclosure.

Caution: To prevent thermal shutdown of the disk drive system, never remove the fan assembly for more than two minutes while the enclosure is powered on.

The fan assembly connects directly to both power supplies, either of which can power it. The fans operate at 9 volts during normal operation to minimize acoustic noise. To compensate for decreased air flow when a fan fails, the voltage increases to 12 volts, resulting in higher acoustic noise.

If the amber check LED is lit, a fan assembly fault has occurred. You can replace the fan assembly while the enclosure is powered on as long as the process is completed in less than two minutes.

If the fan assembly is removed for more than two minutes, the FC-AL interconnect system stays powered on, but the disk modules are automatically powered off. The disk modules power on again when a functioning drive fan assembly is installed in the system. "Replacing a Fan Assembly" on page 75 in Chapter 4 contains complete instructions for replacing a fan assembly.

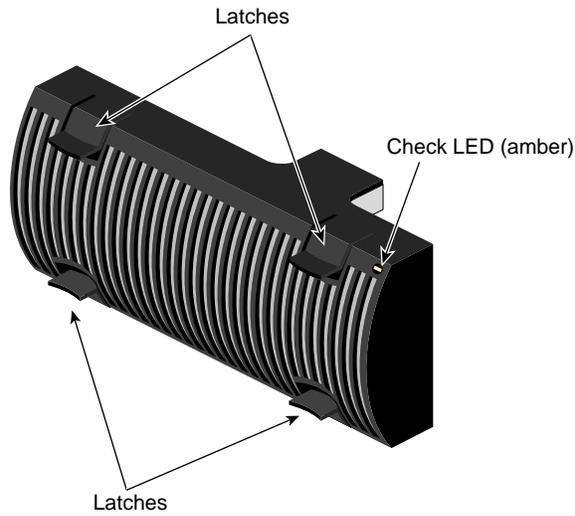


Figure 3-6 Fan Assembly Features

The Enclosure Power Supplies

Each fibre enclosure comes standard with one power supply and accommodates an optional second power supply. Figure 3-7 shows an example power supply.

The power supplies are located behind the fan assembly.

Note that the amber cooling check LED on the supply lights whenever the fan module is removed.

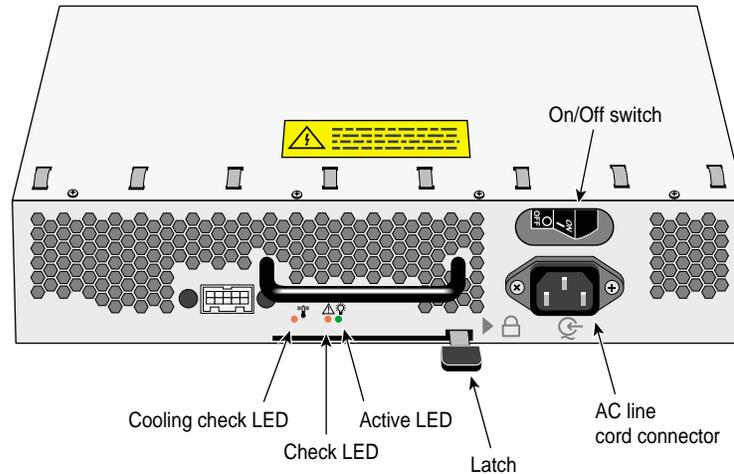


Figure 3-7 Power Supply Features

Each power supply is an autoranging, power-factor-corrected, multioutput offline converter. Each supply supports a fully configured (10-disk) fibre enclosure and shares current with the other supply if it is present.

The power supply status LEDs function as follows:

- Active (green): The power supply is operating.
- Power supply check (amber): If the LED stays on, a power supply fault has occurred.
- Cooling check (amber): If the LED blinks, a cooling fault has occurred.

If a cooling check fault continues for more than two minutes, the drive modules are powered off automatically.

The drive and LCC voltage lines have individual soft-start switches with short-circuit current-limit capability to protect the disk drives and link control cards if you install them while the vault is powered on.

Each power supply has its own on/off switch and its own power cord for connection to separate power sources, which is particularly useful for high-availability systems. Voltage is supplied independently to each disk module. One shorted-out disk module does not affect the other disk modules in the enclosure. If a disk module is added while the enclosure is operating, the voltage is automatically increased gradually to accommodate the new disk module.

The power supplies or their fillers are inverted with respect to each other in the chassis. The chassis is designed to prevent incorrect insertion of the power supplies. In a highly available configuration, you can add or remove the power supply while the enclosure is powered on.

The FibreVault Midplane

The FibreVault midplane features:

- 40-pin SCA-style unitized disk drive connector with blind mate and hot plug features
- Dual-port disk interface
- Slot-dependent disk addressing for each FibreVault enclosure in a rack
- Isolated DC power for each disk
- 3.3-volt support (not currently provided by power supply)
- Five disk LRCs in each direction

Fibre Channel RAID Enclosure Components

The fibre channel RAID enclosure shares a number of functionally similar components with the FibreVault enclosures. These components include the link controller card (LCC) and power supply.

Caution: Although the function and appearance of the LCCs and power supplies are similar between FibreVault and fibre channel RAID enclosures, they are not interchangeable.

These components are described in the section “FibreVault Components” on page 47.

The disk fan module is interchangeable with any of the enclosures.

This section concentrates on components that are unique to the fibre channel RAID enclosure.

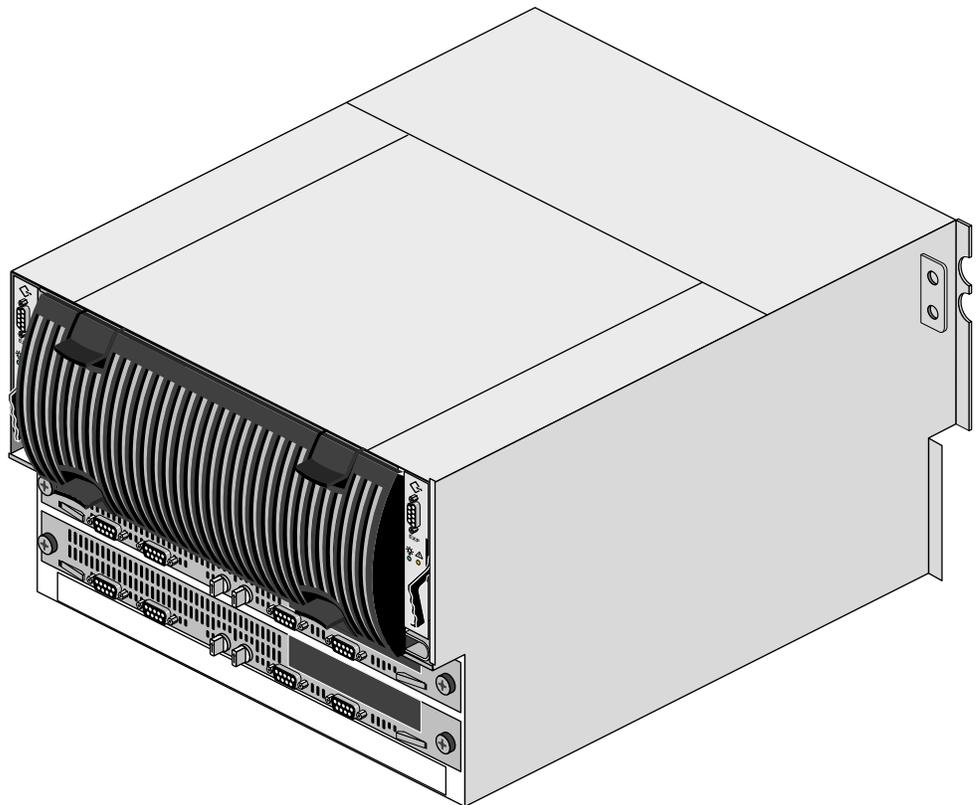


Figure 3-8 Fibre Channel RAID Disk Enclosure (Rear View)

Fibre Channel RAID Deskside Tower

Figure 3-9 shows the fibre channel RAID deskside tower system. It holds one fibre channel RAID disk enclosure see Figure 3-8, and one RAID expansion FibreVault enclosure, (see Figure 3-4). Each RAID tower unit requires a fibre channel connection. The fibre channel RAID and FibreVault RAID expansion enclosures mount vertically in the deskside tower.

The combined maximum capacity of the RAID deskside tower is 20 fibre channel RAID disks. Figure 3-9 shows the fibre channel RAID deskside unit.

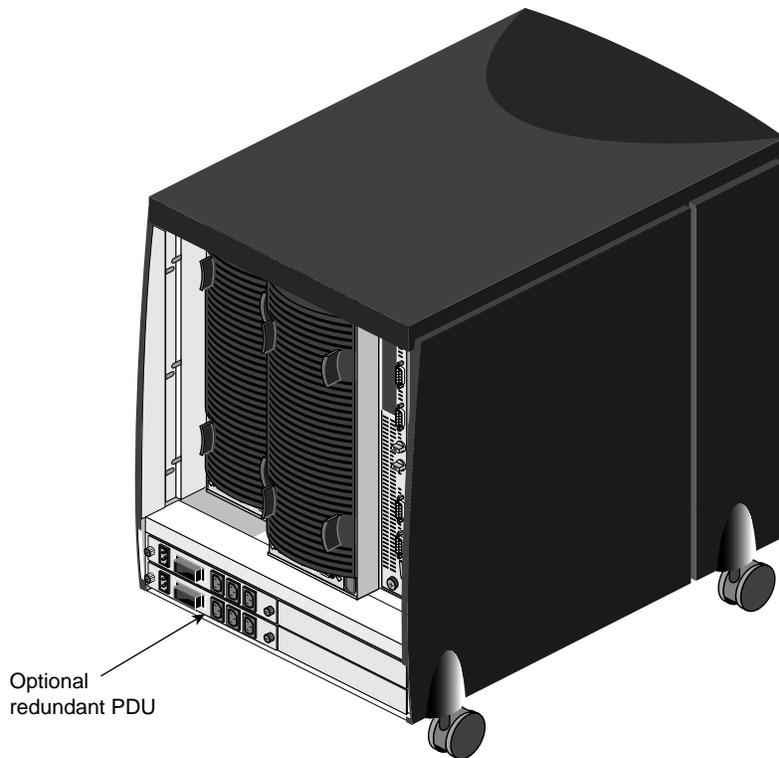


Figure 3-9 Deskside Fibre Channel RAID Tower (Rear View)

Storage Processor Board Assembly

Each fibre channel RAID enclosure has one or (optionally) two storage processor (SP) board assemblies. Figure 3-10 shows the fibre channel RAID enclosure's standard and optional storage processors (SPs). Figure 3-11 shows an example of the storage processor board (SP) and its components.

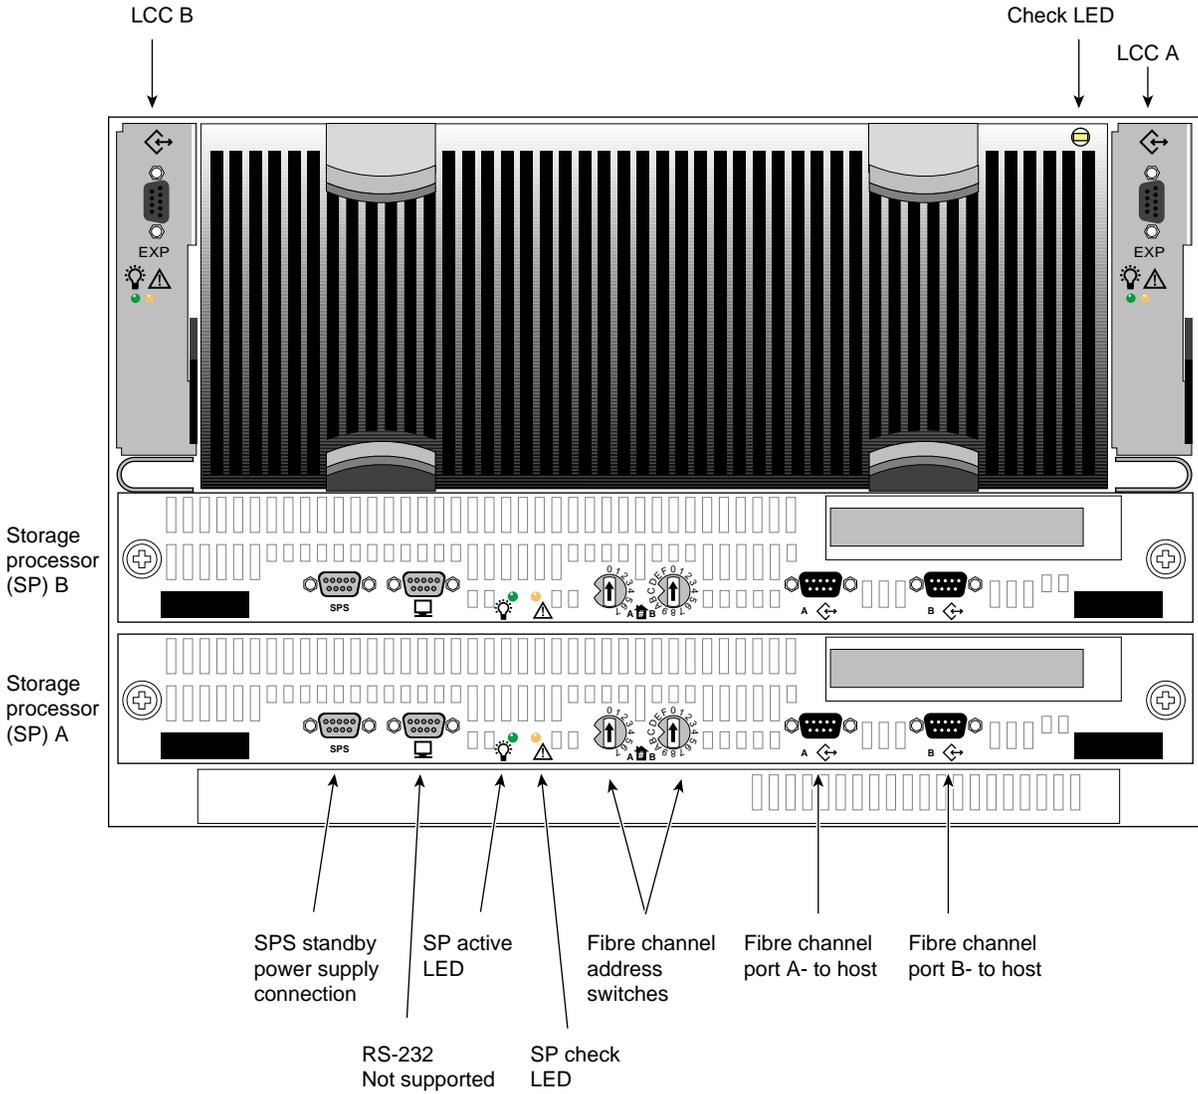


Figure 3-10 Storage Processor Connectors

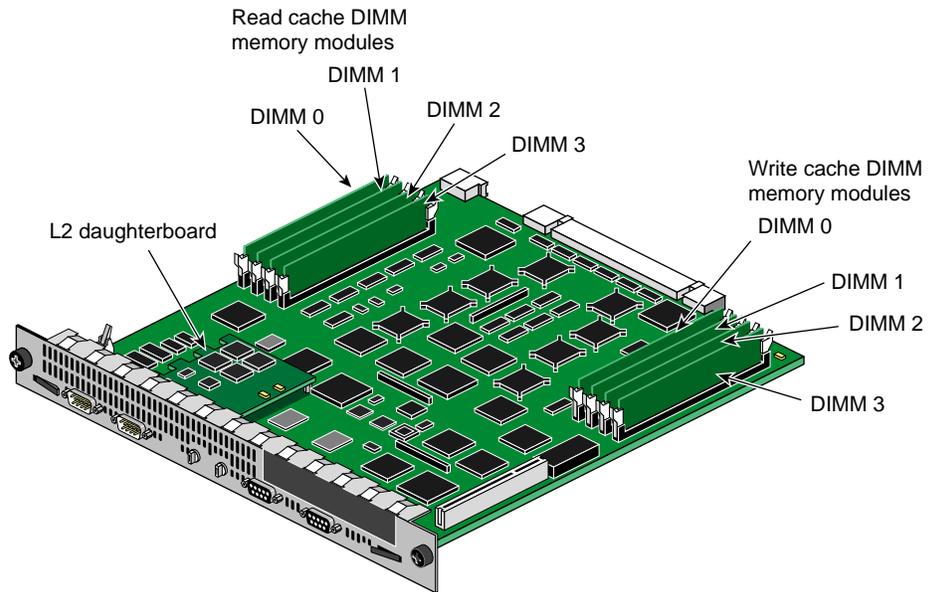


Figure 3-11 Storage Processor Board Components

Storage Processor Fan Module

The storage processor's fan module is located at the front of the fibre channel RAID enclosure. It is a completely different module from that used in the upper rear section of the enclosure to cool the disk assembly.

Figure 3-12 illustrates the fibre channel RAID storage processor (SP) fan module.

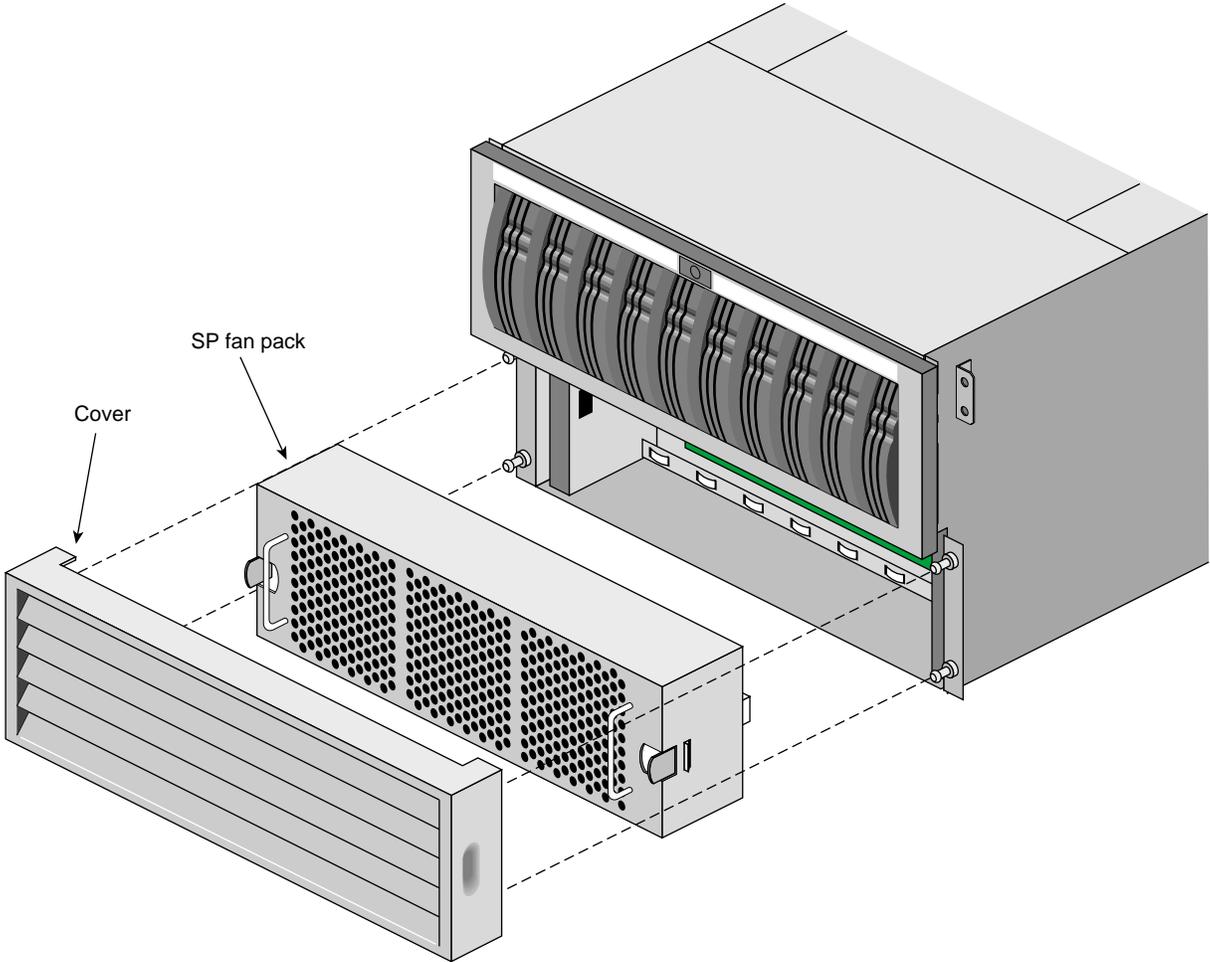


Figure 3-12 Storage Processor Fan Module

Fibre Channel Rack

The fibre channel rack accommodates up to 11 FibreVault (or five fibre channel RAID) enclosures. It has two power distribution units with two separate power inlets so that optional dual power supplies on the fibre enclosures can be connected to separate power sources.

Enclosures within the rack equipped with optional (external) MIAs can be cabled to a host up to 300 meters away. Up to 11 FibreVault enclosures within the rack can be daisy-chained to one port (loop) on a host's FC XIO board. Fibre channel RAID and fibre channel RAID expansion enclosures can also be daisy-chained within the rack.

A fibre channel rack fully loaded with 11 FibreVault enclosures, each with ten disk modules, weighs about 1100 lbs (about 500 kg).

Figure 3-13 diagrams a typical FC storage subsystem.

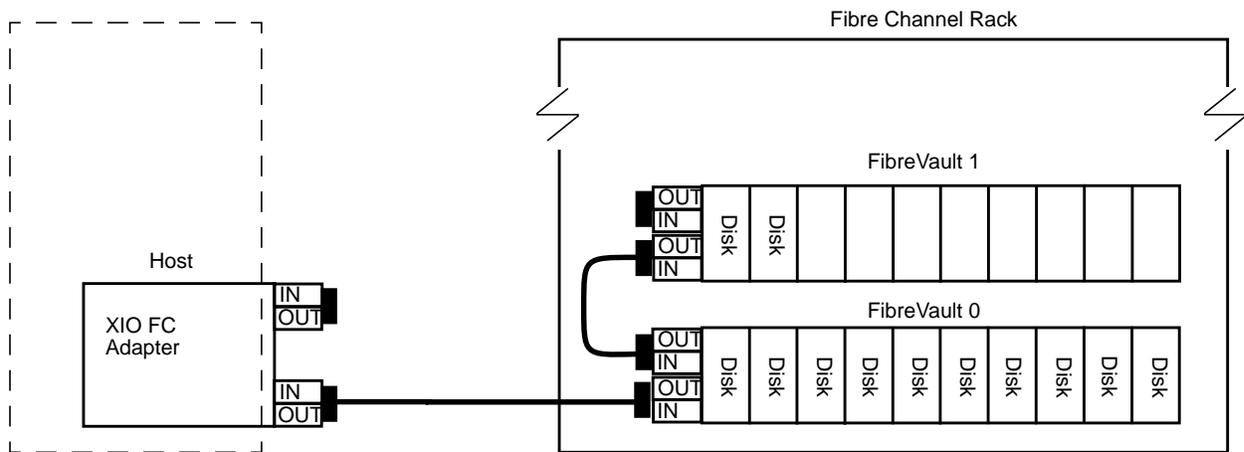


Figure 3-13 FC-AL Storage Subsystem Example Diagram

Hosts are connected to storage subsystems by full-duplex cabling, which can be copper or (optionally) optical. With full duplex cabling, each port provides both its input and output on a single connector. The last storage box in a chain automatically provides a loopback function (similar to connecting a SCSI terminator) which completes the loop. The LCC automatically shunts the open port if no signals from an additional connected enclosure are present.

Optional Standby Power Supply (SPS)

If you need higher availability (especially with RAID caching), you may use the standby power supply (SPS) option in the rack or desktside RAID systems. Figure 3-14 shows features of the SPS panel. Table 3-1 explains the SPS LEDs.

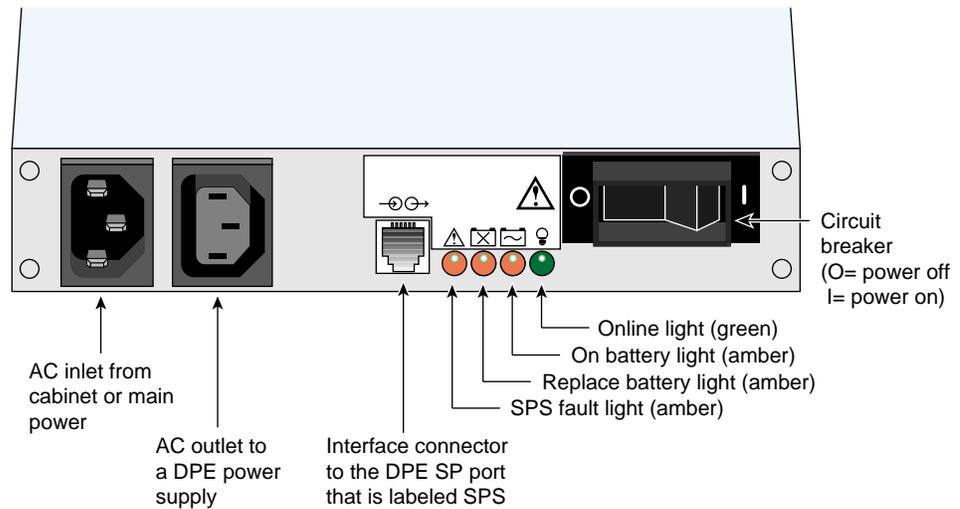


Figure 3-14 Fibre Channel RAID Standby Power Supply (SPS) Front Panel

Table 3-1 SPS LEDs

LED Type and Color	Function or Meaning
SPS internal fault (amber)	Lights when there is an internal SPS fault. Have the SPS replaced as soon as possible.
Replace battery (amber)	The SPS battery cannot support the expected load. The RAID SP flushes all cache data to disk and disables caching. Have the SPS replaced as soon as possible.
On battery (amber)	AC power has failed and the SPS is supplying AC power to the DPE. The DPE writes all cache data to disk and the SP event log records the incident. Restore AC power as soon as possible.
Power on (green)	When this LED is glowing steadily, the SPS is operating normally. Flashing means the SPS battery is being charged.

Note that the RAID enclosure's storage processor (SP) board cannot use write caching unless a fully charged SPS is connected. If the SP does not detect a functional SPS, it automatically disables write caching, reenabling it as soon as a functional SPS is detected. The SPS is always installed or replaced by a trained and approved field service technician. Figure 3-15 shows two SPS units connected to a RAID enclosure in a rack.

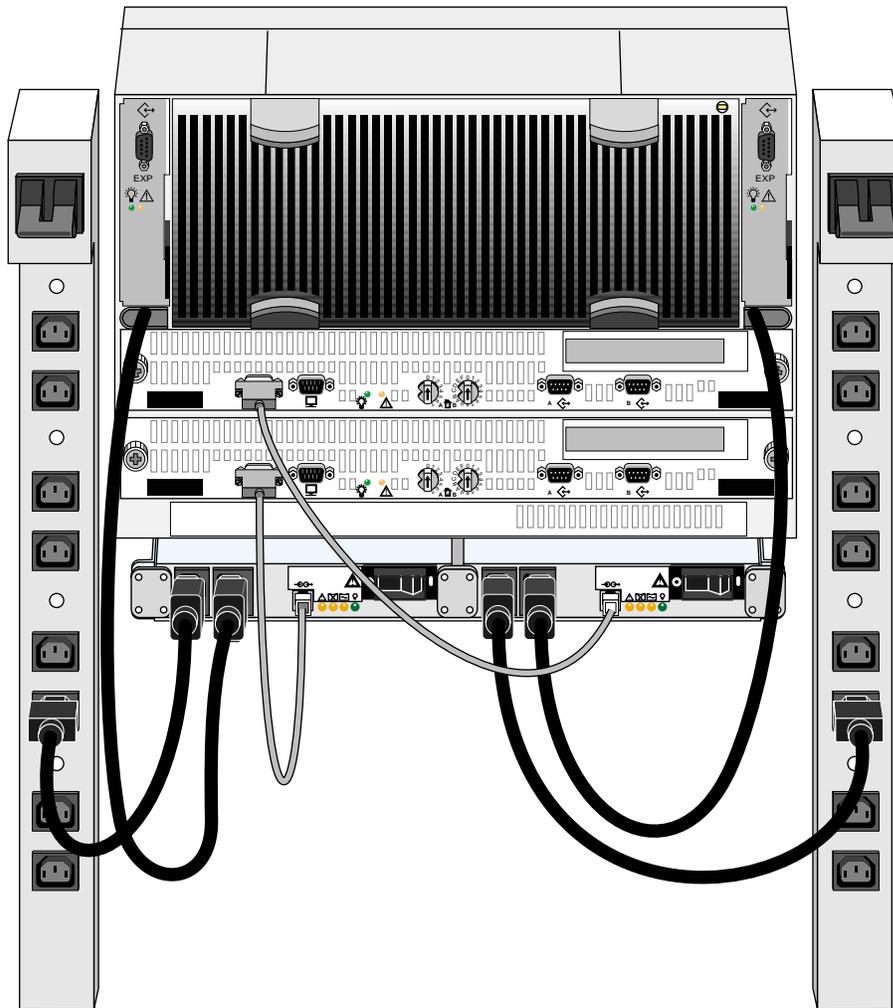


Figure 3-15 Two Rackmounted SPS Examples With Cabling (Rear View)

Monitoring Storage System Status

The FibreVault enclosure and its components are equipped with LEDs that indicate status. These LEDs are visible from the front and the back of the enclosures.

Table 3-2 summarizes LEDs visible from the front of the rack.

Table 3-2 LEDs Visible From the Front of the Fibre Enclosure

LED	Quantity	Color	Meaning
FibreVault active	1	Green	On when FibreVault is powered on.
FibreVault check	1	Amber	On when any fault condition exists. Flashing when a FibreVault cooling fault exists. When this fault persists for more than about two minutes, the FibreVault powers off the disk drives.
Drive active	1 per disk module slot	Green	Off when the disk module slot is empty or contains a filler. Flashing (mostly off) when the disk drive is powered on but is not spinning. Flashing (mostly on) when the disk drive is spinning.
Drive check	1 per disk module slot	Amber	On when the disk module is faulty or a port bypass is activated.
Enclosure address	12	Green	One is lit to indicate the enclosure address (see "Setting an Enclosure Address" on page 71 for details).

Note: If a FibreVault enclosure check LED is illuminated, inspect all other check LEDs to determine the faulty component. Have any faulty component replaced as soon as possible. If a component fails in an enclosure that is not highly available, it may be inoperable until the component is replaced. If a component fails in a highly available FibreVault enclosure, the enclosure's high availability is compromised.

Table 3-3 summarizes enclosure LEDs visible with the rear door of the rack open.

Table 3-3 LEDs Visible From the Rear of the Fibre Enclosure

LED	Quantity	Color	Meaning
LCC active	1 per LCC	Green	On when the LCC is powered on.
LCC check	1 per LCC	Amber	On when either the LCC or an FC-AL connection is faulty.
Power supply active	1 per power supply	Green	On when the power supply is operating.
Power supply check	1 per power supply	Amber	On when the power supply is faulty.
Cooling check	1 per power supply	Amber	Flashing either when multiple fans in the fan assembly are faulty or the assembly is removed. The enclosure powers off the disk drives when the fault persists for more than about two minutes.
Fan assembly check	1 per fan assembly	Yellow	Remains on when one fan in the fan assembly is faulty; flashes when two fans are faulty.

Chapter 4, “Fibre Channel Component Replacement,” contains complete instructions on replacing components.

Shutting Down the Fibre Channel Rack and Disk Enclosures

To power off the Origin fibre channel rack:

1. Use the host system’s administrative CLI or GUI to take all disk arrays or LUNs offline.
2. Open the rear door of the rack and push each of the PDU breaker switches down and into the Off position.
3. Unplug both of the rack’s power cables (see Figure 3-16) to cut off all power to the rack.

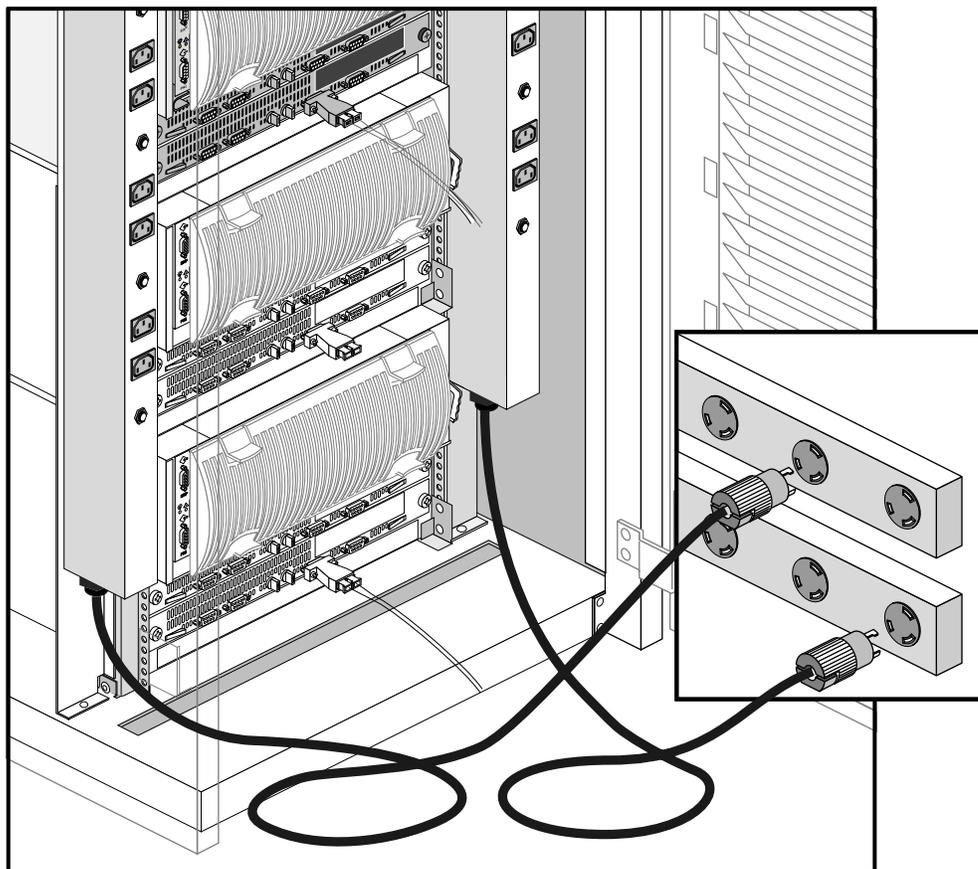


Figure 3-16 Disconnecting Rack Power Cables

To power off an individual fibre enclosure (while the rack is still powered on):

1. Use the host system's administrative CLI or GUI to take the target enclosure and its disks or (LUNs) offline.
2. Pull the disk fan module out of the back of the enclosure (see Figure 3-17).
3. Push the breaker switch on each power supply in the enclosure to 0 "Off".
4. Shut down the main AC lines to the enclosure by unplugging the power cord(s) that connect it to the rack's PDUs.

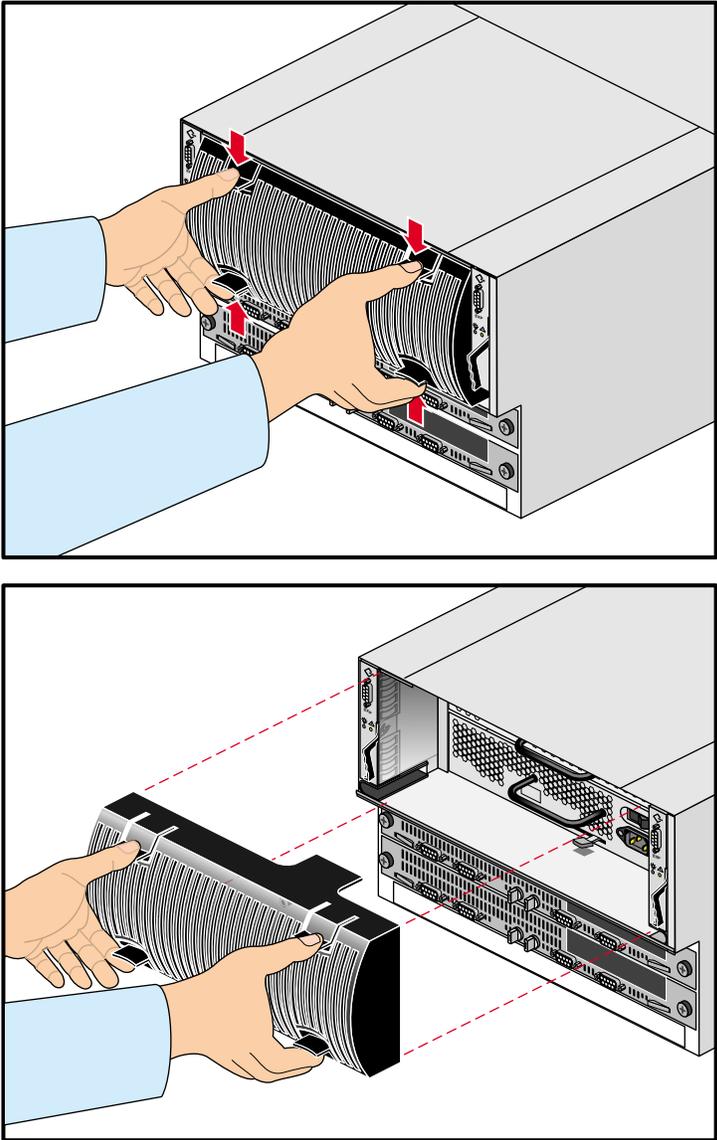


Figure 3-17 Removing an Enclosure Disk Fan Module Example

Fibre Channel Component Replacement

The FibreVault or fibre channel RAID enclosure is designed to remain powered on and hot-repairable; that is, you can replace some components without powering off the storage system or the host(s) to which it is attached. Each FibreVault or fibre channel RAID enclosure compartment should contain a component (drive, power supply, and so on) or filler panel to ensure EMI compliance and proper air flow over the components.

While the FibreVault or fibre channel RAID enclosure is powered on, you can remove a disk module, redundant link control card (LCC), or filler, but you should replace it as soon as possible with a replacement component or a filler. You can also remove the fan assembly while an enclosure is powered on, but you must replace it within two minutes to keep the disk drive modules powered on. If you do not replace the fan module within 2 minutes after removing it from an operating enclosure, the enclosure automatically powers off the drives. Do not remove a faulty component until you have a replacement component available. Contact your Silicon Graphics service or sales representative to order additional parts.

Caution: The FC-AL disk modules and the Ultra SCSI disk modules (used in Origin or Onyx2 chassis) are not interchangeable and cannot be substituted for each other. Fibre channel RAID and non-RAID disks are not interchangeable and cannot be substituted for each other.

This chapter includes the following main sections:

- “Adding or Replacing a Disk Module” on page 66
- “Setting an Enclosure Address” on page 71
- “Replacing a Fan Assembly” on page 75
- “Replacing or Adding a Power Supply in an Enclosure” on page 77
- “Replacing or Adding a Link Control Card (LCC)” on page 89

Adding or Replacing a Disk Module

Use only Silicon Graphics FC-AL disk modules. Always confirm the type of fibre drive (RAID or non-RAID) before you install it. Each drive module has an identifying sticker with its part number.

Table 4-1 provides descriptions and part numbers for the fibre drives.

Table 4-1 Fibre Channel Disk Drive Descriptions

Fibre Drive Type	Drive Capacity	Drive Speed	Sector Sizing	Part Number
Non-RAID	9.1 GB	7200 RPM	512-byte sectors	9470140
RAID	8.8 GB	7200 RPM	520-byte sectors	9470192
Non-RAID	9.1 GB	10,000 RPM	512-byte sectors	9470282
RAID	8.8 GB	10,000 RPM	520-byte sectors	9470281
Non-RAID	18 GB	7200 RPM	512-byte sectors	9470255
RAID	17.8 GB	7200 RPM	520-byte sectors	9470257

Note: Always consult with your service provider before replacing a disk with another that has a different capacity or speed rating (different part number).

Be aware of the following information when adding or replacing disks:

- RAID and non-RAID fibre drives are not compatible and may not be mixed in the same enclosure.
- A lower capacity RAID drive must not be used as a replacement disk or hot spare in an array of higher capacity disks. For example, you must never install an 8.8 GB RAID drive as a hot spare or replacement in an enclosure of 17.8 GB disks.
- Installing a lower speed (RPM) drive in an array of faster RPM disks causes the entire array bandwidth (read and write speed) to slow to the lower speed capacity.

Caution: The disk modules have no ESD shielding or protective case. Do not stack them on top of each other before installing them. Do not drop, jar, or handle the drives roughly. Use standard ESD precautions as listed in the instructions and always leave disks in an anti-static container until ready for installation.

Do not replace more than one disk module at a time unless the enclosure is powered off.

Special rules apply when you replace a fibre channel RAID disk module in either of the following groups:

- The first three disk modules in the first fibre channel RAID enclosure in a chain, or in the only enclosure in the array, contain licensed internal code (LIC). These disk modules, with disk IDs 00, 01, and 02, are known as *database disk modules*.

If all three are removed at the same time when the array is powered on, contact with the array is lost. If you remove them with the array powered off, label these as you remove them, because they must be reinstalled in their original positions.

- If the array uses write caching, the disk modules that the array uses for its cache vault are 00 through 08.

If a cache vault disk module fails, the array dumps its write cache image to the remaining modules in the vault. Then it writes all dirty (modified) pages to disk and disables write caching. Write caching remains disabled until a replacement disk module is inserted and the array rebuilds the LUN with the replacement module in it. The Write Cache State field in the Cache section of the SP Information window (see the section on SP cache information in Chapter 4 of the *Origin FibreVault and Fibre Channel RAID Administrator's Guide*) tells whether array write caching is enabled or disabled.

General Guidelines for Adding or Removing Disks

Before adding, removing, or replacing disk modules, note the following:

- Fill slots starting with the lowest address (0) and add disks until you reach the highest applicable address (left to right).
- In a fibre channel RAID enclosure, drives 0, 1, and 2 contain licensed internal code. They should be replaced one at a time.
- If you want to replace all the disk drive modules in a RAID LUN, you must have your service provider update the licensed internal code and PROM code. This procedure requires a service call.
- All fibre channel RAID disk enclosures and their expansion enclosures attached to an FC-AL should have at least five disk modules installed.
- If you must configure a non-RAID FibreVault with only one drive, it must be either the last enclosure on the loop or the only enclosure on the loop.

- The non-RAID FibreVault enclosure must always have a disk module installed in either slot 0 or 1. Access to enclosure-specific environmental information depends on at least one of these slots being populated with a disk.
- A FibreVault or fibre channel RAID enclosure cannot remain connected to an FC-AL if it has no drives installed.
- Note that a RAID 3 LUN can use only one hot spare in case of disk failure. If a second hot spare is available and a second disk module in the RAID 3 LUN fails, the LUN does not use the second hot spare.

Confirming a Failed Disk Module

In a fibre channel RAID enclosure, removing the wrong disk module can introduce an additional fault that shuts down the LUN containing the failed drive. Before removing a disk module, check the module's amber check or fault light. If it is illuminated, or if both LEDs are off, the host is bypassing it, indicating a failure.

To check further on the status of a disk module:

- Use the `fccli GETSTATUS;` command for non-RAID (FibreVault) disks.
- For a RAID disk module, use the command `ssmcli -d device getdisk -state` or the GUI.

For full details on these commands, see the *Origin FibreVault and Fibre Channel RAID Administrator's Guide*.

Physically Replacing a Disk

To add or replace a disk module, follow these steps:

1. Be sure you understand the rules and limits for replacing fibre channel RAID and non-RAID disks. See "General Guidelines for Adding or Removing Disks" on page 67 and "Confirming a Failed Disk Module" on page 68.
2. Have ready the proper replacement disk module.
3. Back up any information on the disk(s) you are replacing as needed.
4. Use one of the following methods to take the disk offline and prepare for installation:

- Tell the host that you are removing a (non-RAID) disk module and take it off line using the `fcccli REMOVE` option. The *Origin FibreVault and Fibre Channel RAID Administrator's Guide* has details for using the line commands.
 - Shutdown the host and fibre channel enclosure systems.
 - If a failed disk is part of a RAID LUN, you might need to unbind the LUN before the replacement under certain circumstances. See the *Origin FibreVault and Fibre Channel RAID Administrator's Guide*.
5. Open the FibreVault or fibre channel RAID enclosure's front door by inserting the key in the door's latch and turning the key to the unlocked icon position.
 6. Remove the key. Push the door's latch, and pull the door down until it is fully open.
If you accidentally pull the door down with so much force that it becomes unhinged, snap it back on the hinges.
 7. Attach the clip of the ESD wristband to the FibreVault or fibre channel RAID enclosure chassis and put the wristband around your wrist with the metal button against your skin.
Caution: The enclosure has no mechanism to prevent you from removing a disk module while the drive is spinning. Do not remove a module unless its green active light is off and its amber check light is on (or the entire array may shutdown).
 8. Locate the disk or filler module to be removed from the target slot. This module's yellow LED is lit (or both LEDs are off) to show that the host is bypassing it. The green activity LEDs of the other disk modules in the loop blink at a regular interval as part of a background loop verification process.
Note: Do not remove a disk module or filler unless you have its replacement ready to install. Do not remove more than one disk module at a time.
 9. Grasp the disk module's handle so that your thumb is on its latch at the top of the handle.
 10. Push the latch and slowly pull the module from the slot, supporting it with your other hand, see Figure 4-1. The disk module's active and check LEDs turn off. If power is on and if the disk you removed was bound in a RAID LUN, wait for the active LEDs of the other disk modules in the loop to resume blinking at a regular interval, which indicates that the driver has completed any necessary loop discovery.

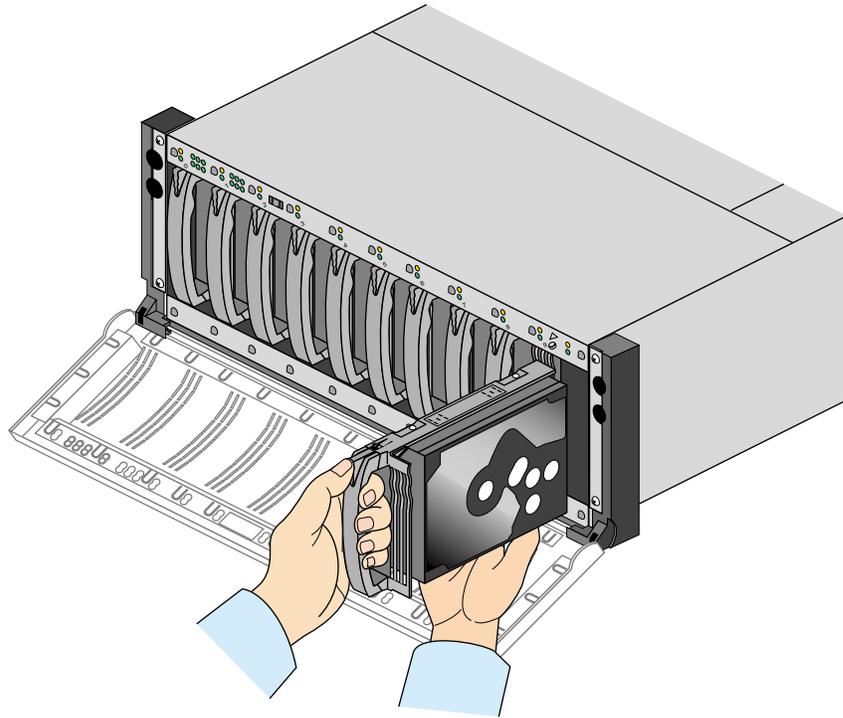


Figure 4-1 The Disk Module

11. Grasp the handle of the new disk module and align the module with the guides in the empty slot. The module's thumb latch should be at the top.
12. Slowly push the disk module into the slot until it clicks into place. The disk module's active LED flashes and then remains on (if power is on).

If power is on and you have replaced a RAID disk module, the disk module's activity LED blinks; it may also flash rapidly.

When the replacement RAID disk module functions normally as part of the LUN, its green LED blinks at the same interval as the rest of the disk modules on the loop. At this point, you can replace another RAID disk module if necessary.

Note: If the fault LED stays on after you insert the disk module, remove it and reinsert it. If it remains on, use another replacement disk module.

If you are replacing a failed non-RAID disk module, its amber LED remains lit until you use the `fccli INSERT` command as explained in Step 14.

13. Remove and store the ESD wristband when all replacement procedures are complete.
14. When you have replaced or added the disk module, use one of the following procedures to bring the disk online.
 - Power the host and the fibre channel enclosure back on if they were powered off during the installation procedure.
 - Use the `fccli INSERT` option in the command line interface to inform the host (that owns the loop) of its presence, see the *Origin FibreVault and Fibre Channel RAID Administrator's Guide*. The fault LED remains lit until the drive is online.
15. Push the fibre enclosure's front door up until it latches into place. Note that the maximum number of times a disk module can be inserted and removed is 500. Exceeding this limit may damage the SCA connector on the disk module or the connector on the enclosure midplane.

Setting an Enclosure Address

The drives read their FC-AL physical address only at power-on or when the drive is reset. As a result, each enclosure address is set when installed in the Fibre rack while the power is off; do not change the address while the enclosure's power is on.

Set the enclosure address for each FibreVault enclosure by using a pencil or ballpoint pen to move the down or up switch.

The LED lights up for the number you set. FibreVault enclosures are numbered 0-10; use only those numbers. Figure 4-2 shows the address switch. When you are finished setting the FibreVault address, push the enclosure's front door closed until it latches into place. If desired, lock the enclosure's front door.

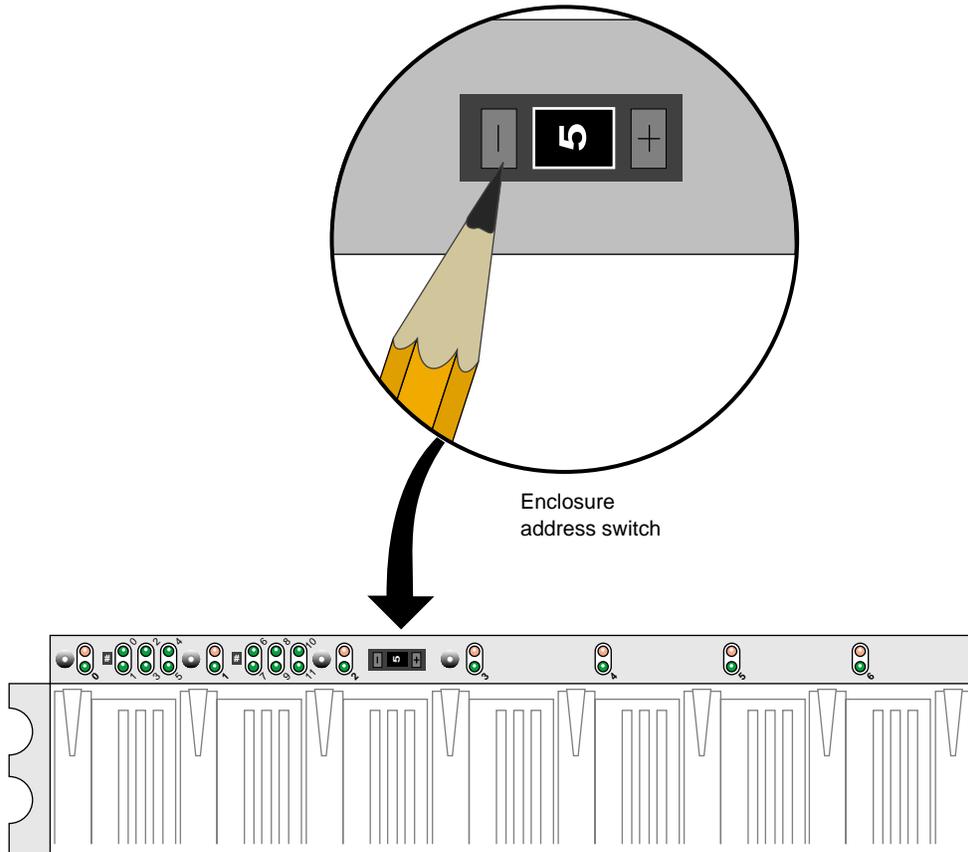


Figure 4-2 Setting the FibreVault ID

Valid address ranges for fibre channel RAID (see Figure 4-3) are 0 through 109 (decimal). This equals 0 through 6D hexadecimal (0 on both switches through 6 on the left switch and D on the right switch). Each SP must have a unique address set using the switches.

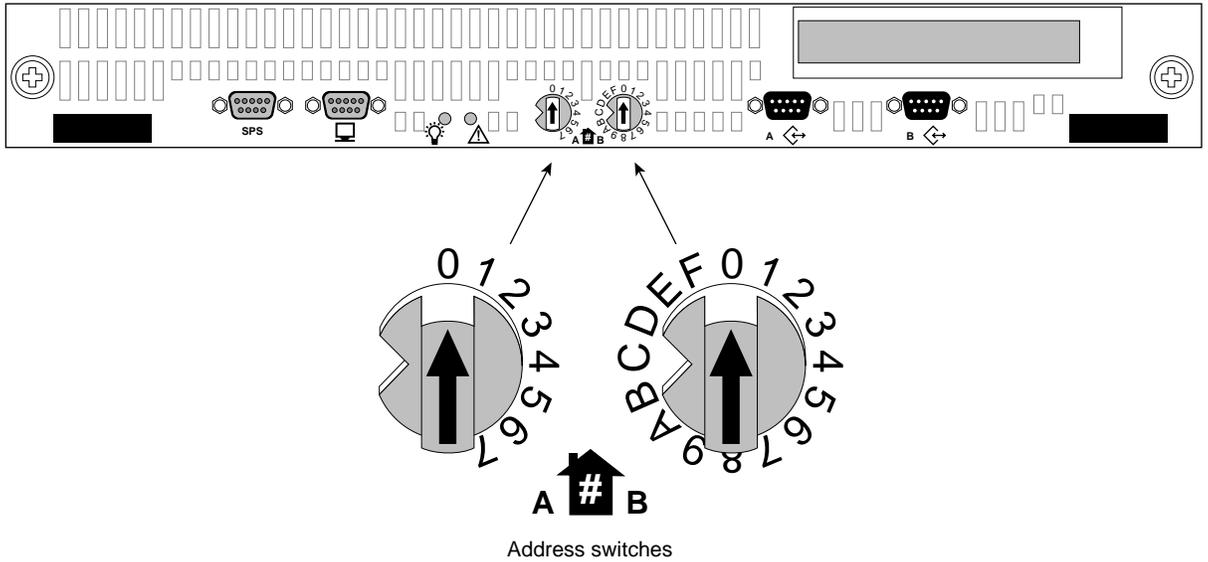


Figure 4-3 Fibre Channel RAID Enclosure Channel Address Switches

The fibre channel RAID enclosure address the loop sees is determined and set on the system processor (SP) board (see Figure 4-4 for the location of each SP). Each SP must have a unique fibre channel front-end address. Table 4-2 shows the maximum disk numbering for FibreVault (fibre channel RAID is lower).

Table 4-2 FibreVault or Fibre Channel RAID Disk Slot Numbering

Fibre Enclosure ID	Disk Slot Number
0	0-9
1	10-19
2	20-29
3	30-39
...	...
9	90-99
10	100-109

The disk slots in a FibreVault or fibre channel RAID enclosure are numbered 0-9. Disk modules are always numbered based on the ID of the FibreVault or fibre channel RAID enclosure.

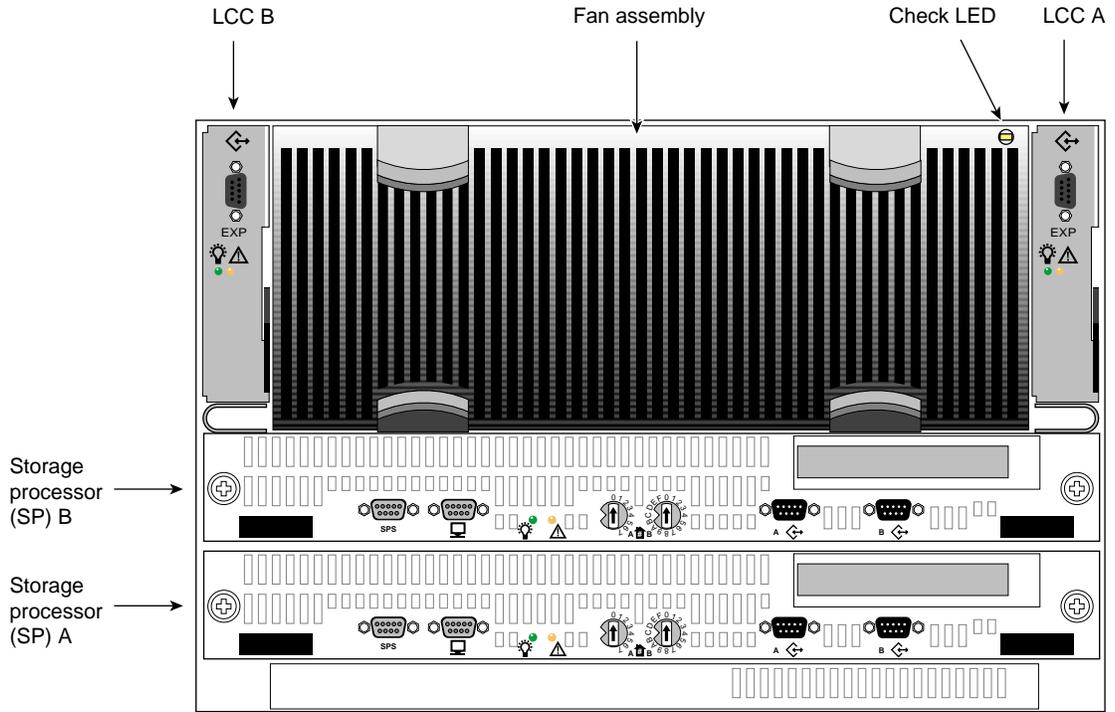


Figure 4-4 Fibre Channel RAID Enclosure Rear View

Replacing a Fan Assembly

If you are replacing the fan assembly in an operating storage system, you have a total of two minutes in which to complete the process. Do not remove a fan assembly unless you have its replacement ready to install in the enclosure.

Note: An operating drive assembly automatically shuts down after 2 minutes if the fan assembly is not replaced, or if more than one fan fails. If the disk modules do shut down under these circumstances, the green disk module LEDs and the green (active) SP LEDs remain illuminated. Even though the disk modules and SP(s) are shut down, they are not faulted. The disk modules and SP(s) automatically power back on when a functional fan assembly is installed.

Follow these steps:

1. Have the replacement fan assembly at hand, unwrapped and ready to insert.
2. Locate the enclosure with the faulty fan assembly; the fan assembly check LED is on.
3. Grasp the fan assembly's handles (squeeze latches) on one side and pull them inward towards the middle of the fan assembly until you feel them loosen.
4. While keeping a thumb on each of the already loose latches, grasp the set of squeeze latches on the opposite side of the assembly with your fingers and pull them in toward the middle of the assembly and outward, see Figure 4-5.

Caution: Do not leave the fan assembly out of the fibre enclosure for more than two minutes if it is running.

Once the fan assembly is disconnected from the power supply, its check LED turns off. The enclosure's front panel check LED and the cooling check LED on each power supply start flashing.

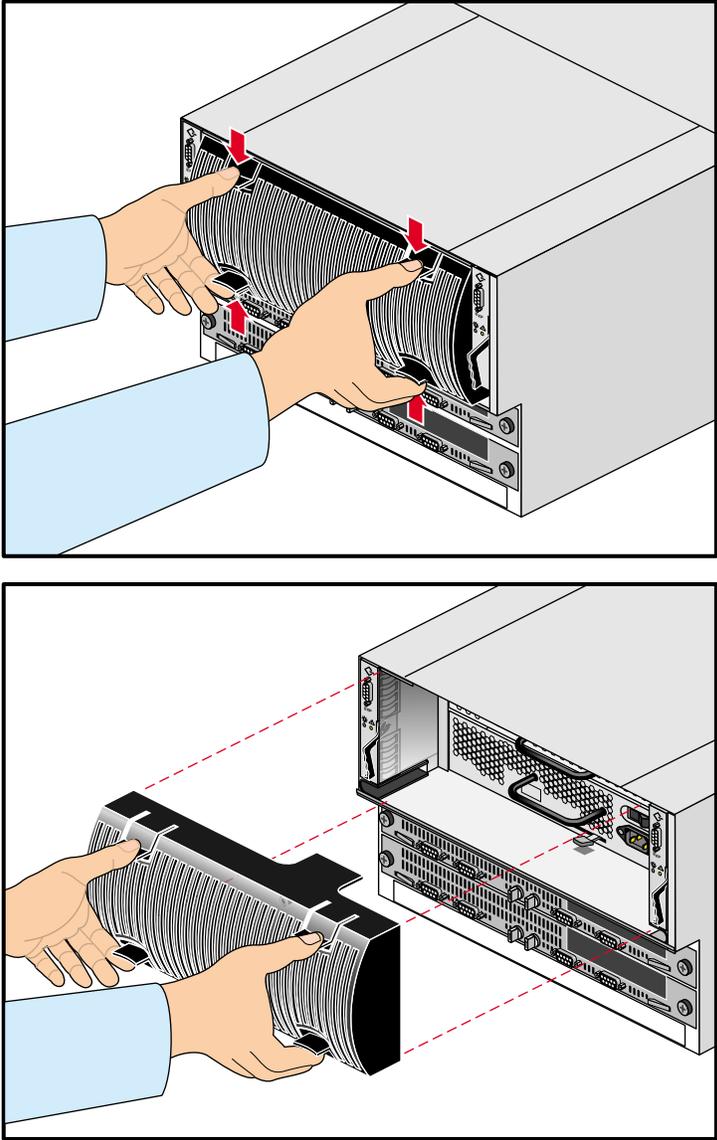


Figure 4-5 Removing the Disk Module Fan Assembly

5. Orient the new fan assembly. Although the fan assembly can attach to the enclosure with either side up, for a consistent image with multiple FibreVault or fibre channel RAID enclosures, install it with the status light in the upper right corner, as shown in Figure 4-6.
6. Squeezing the fan assembly latches toward each other, align and push the assembly forward into the chassis. The fan assembly snaps onto the power supply.

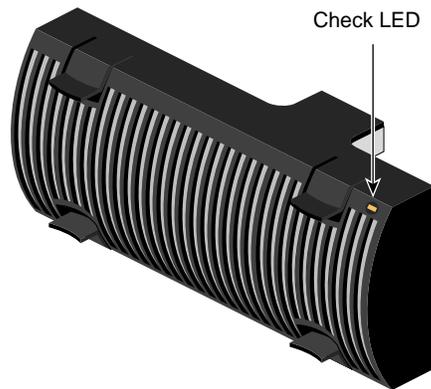


Figure 4-6 Fan Assembly Status LED Orientation

Replacing or Adding a Power Supply in an Enclosure

The power supply modules are located behind the fan unit in the Origin FibreVault or fibre channel RAID enclosure, as was shown in Figure 4-5. This procedure's steps are only fully applicable when the vault has dual (redundant) power supplies. If there is only one supply installed, the FibreVault or fibre channel RAID enclosure shuts down if the supply fails.

This section contains the following subsections:

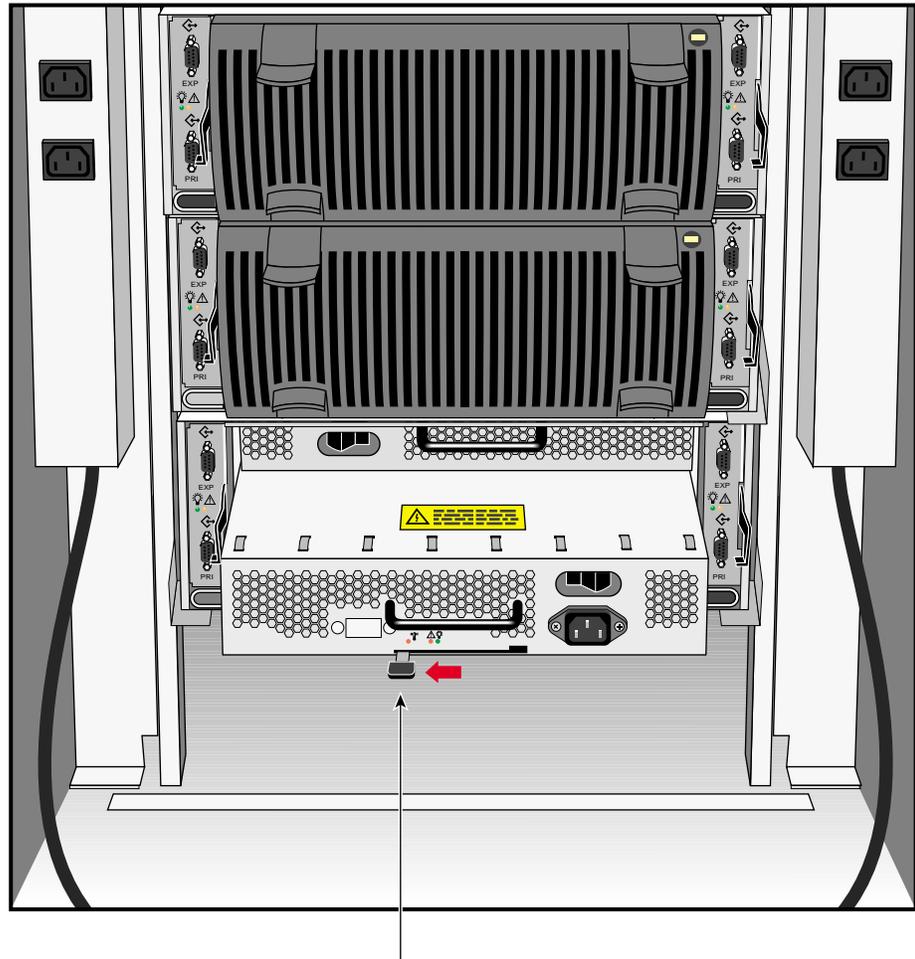
- "Preparing the New Power Supply" on page 78
- "Removing a FibreVault Power Supply or Filler" on page 80
- "Inserting the New FibreVault Power Supply" on page 83
- "Removing and Replacing a Fibre Channel RAID Power Supply" on page 84

Do not remove a power supply or filler unless you have its replacement ready to install.

Preparing the New Power Supply

Follow these steps:

1. Have the new (replacement or additional) power supply at hand, unwrapped and ready to insert. Have space available to receive the fan unit and the power supply you are replacing.
2. Make sure the lever on the new power supply is in the unlock position, as shown in Figure 4-7.



Latch
(shown in unlocked position)

Figure 4-7 Unlatching a New Power Supply

Removing a FibreVault Power Supply or Filler

If you are replacing the power supply in an operating storage system, you have a total of two minutes in which to complete the process. This is the maximum amount of time allowed for the fan assembly to be out before an automatic shutdown of the enclosure is executed.

If two minutes is inadequate time for you to make the exchange, consider shutting down the enclosure first. See “Shutting Down the Fibre Channel Rack and Disk Enclosures” in Chapter 3.

1. Identify the FibreVault or fibre channel RAID enclosure containing the power supply with its check light on. The light is visible through the fan assembly.
2. Attach the clip of the ESD wristband to the enclosure chassis and put the wristband around your wrist with the metal button against your skin.
3. Grasp the fan assembly by the handles on either side and pull it out to expose the power supply (bay), as shown in Figure 4-5.

Caution: Do not leave the fan assembly out of the enclosure for more than two minutes or the vault goes through an automatic thermal shutdown.

4. Turn off the faulty power supply by pressing the on/off switch. The power supply's active and check LEDs turn off.
5. Unplug the AC line cord from the faulty power supply (see Figure 4-8).

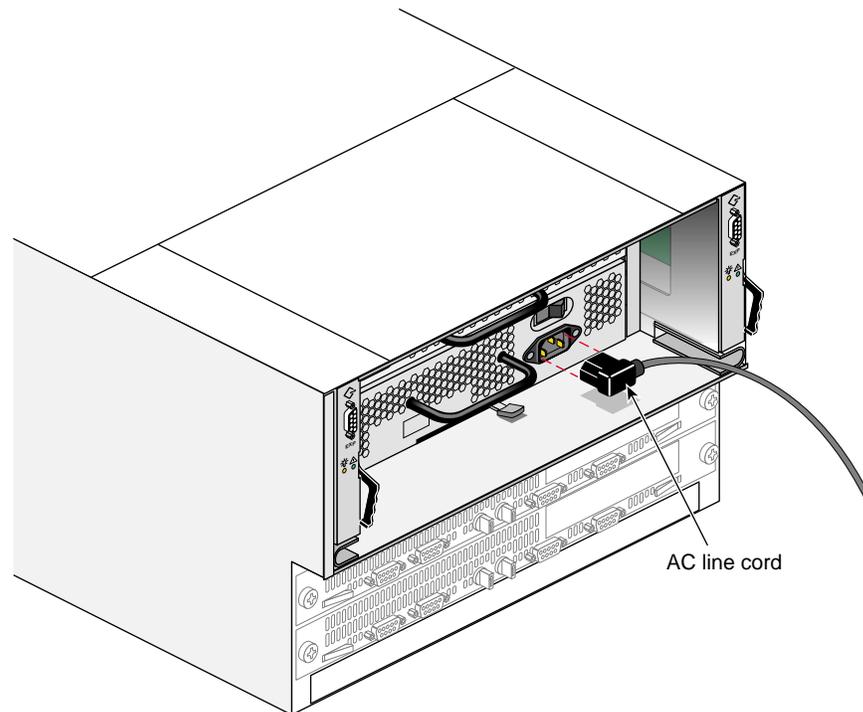


Figure 4-8 Power Supply Power Cord

6. Remove the faulty power supply (see Figure 4-9):
 - For the top power supply, put your right palm against the right side of the power supply handle; with your thumb, push the power supply's latch up and then right to the unlock position.
 - For the bottom power supply, put your left palm against the left side of the power supply handle; with your thumb, push the power supply's latch down and then left to the unlock position.

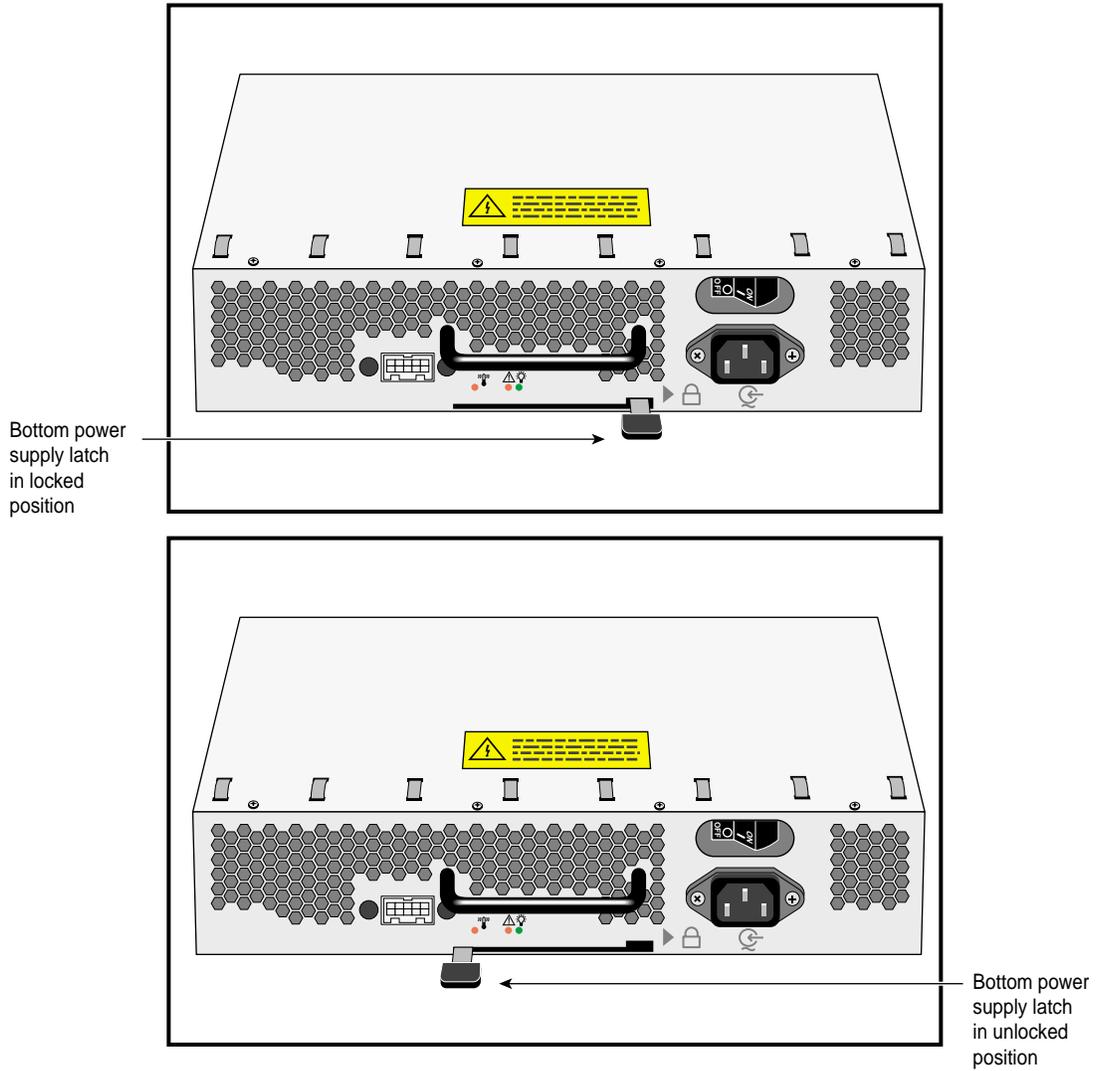


Figure 4-9 Unlatching a Power Supply

7. Pulling the handle, slide the filler or power supply out of the bay, supporting it with your other hand, as shown in Figure 4-10.

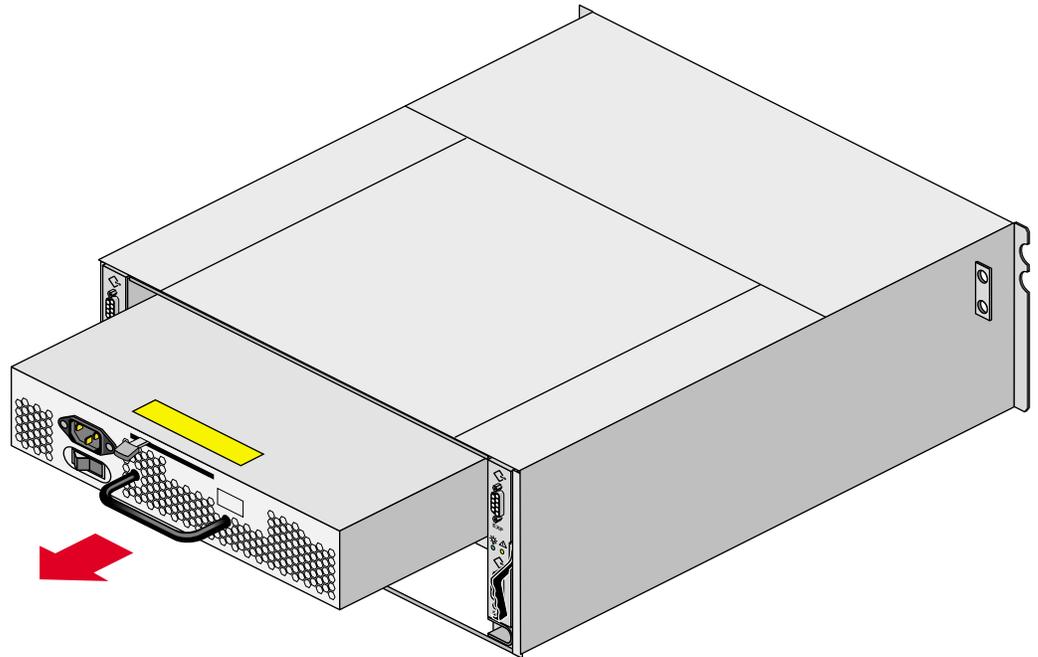


Figure 4-10 Removing a FibreVault Power Supply

Inserting the New FibreVault Power Supply

To insert a new power supply into a FibreVault, follow these steps:

1. Make sure the lever on the new power supply is in the unlock position, as was shown in Figure 4-7.
2. Orient the new power supply, as shown in Figure 4-9, and slide it into the bay.
3. Move the lever to the lock position, as shown in Figure 4-9.
4. Plug the AC line cord into the new power supply, as shown in Figure 4-8.
5. Turn on the power supply by pressing the on/off switch.

The power supply's active LED turns on.

6. Reattach the fan assembly by squeezing the latches toward each other and pushing the assembly into the chassis.

Do not leave the fan assembly out of the fibre enclosure for more than two minutes if you want the enclosure to keep running.

You can install the fan assembly in either orientation. However, for a consistent image with multiple FibreVault or fibre channel RAID enclosures, install it with the status light in the upper right corner as shown in Figure 4-6.

Removing and Replacing a Fibre Channel RAID Power Supply

To remove, replace, or add a power supply to the fibre channel RAID enclosure, you must first remove the disk fan module (see “Replacing a Fan Assembly” on page 75). As noted in the fan replacement instructions, if you leave the fan assembly out of the fibre enclosure for more than two minutes, the disks shut down automatically. In the case of the fibre channel RAID enclosure, the system processor board(s) shut down also. Unless you are experienced at power supply removal or replacement, it is likely that the procedure will require more than two minutes to complete.

Handle the power supply with care and use all standard ESD precautions. Use the following steps if you are removing and replacing an existing power supply:

1. After removing the fan module, turn off the supply’s power by pressing the rocker switch to the 0 position.
2. Unplug and remove the supply’s AC line connector (see Figure 4-11).
3. Unlatch the top power supply by moving the handle up and to the far right.
4. Unlatch the bottom power supply by moving the handle down and to the left.
5. Pull the supply straight out using the handle and both hands to steady it (see Figure 4-12).

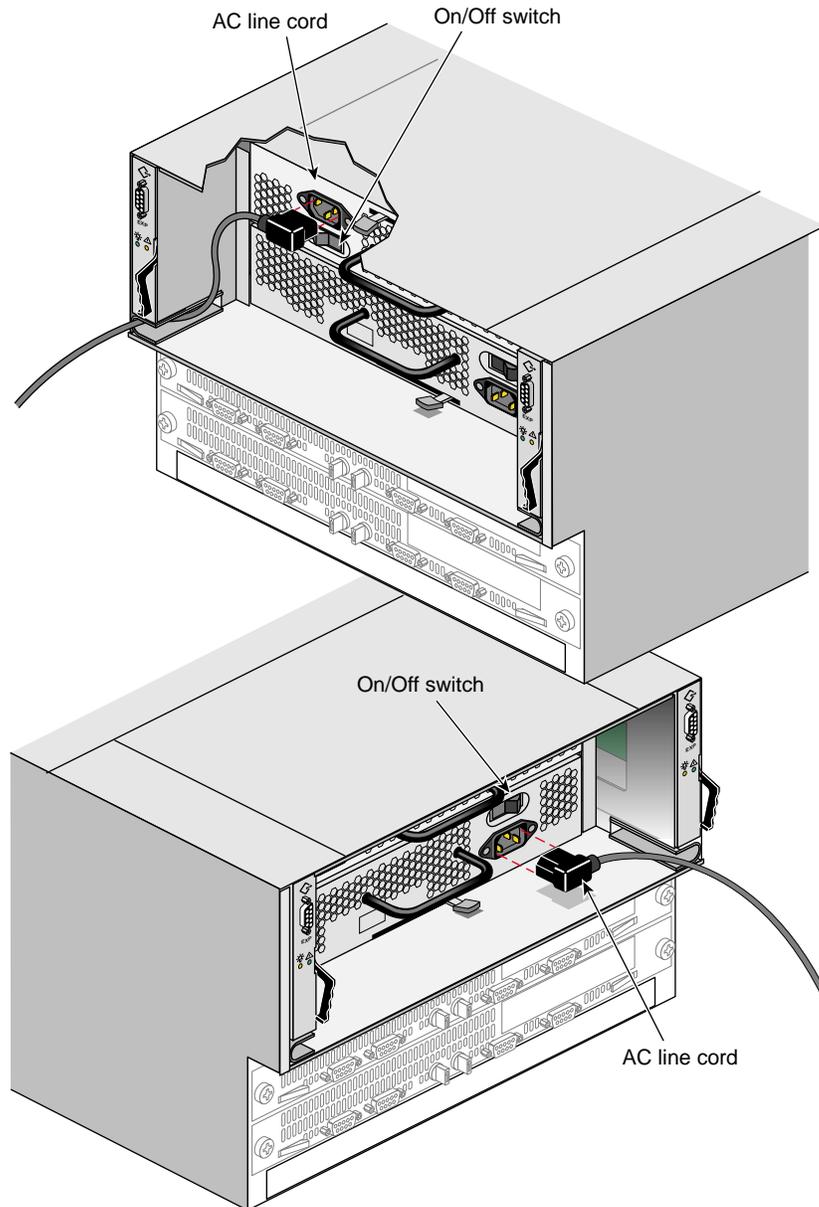


Figure 4-11 Turn Off Power and Unplug Cord from the Supply

6. After confirming the replacement supply's latch is in the open position, install the new supply by aligning it and sliding it straight in.
Note: The orientation of the bottom supply is inverted (upside down) from that of the top.
7. Push the supply into place until you see the latching lever move toward the latched position (see Figure 4-13).
8. Push the latch all the way into the locked position.
9. Plug the AC connector into the supply and route the cord along the power supply, then under the LCC position so that the cord exits the chassis without blocking the installation of the fan module.
10. Set the supply's rocker switch to the On position.
11. Install the fan module, see "Replacing a Fan Assembly" on page 75.

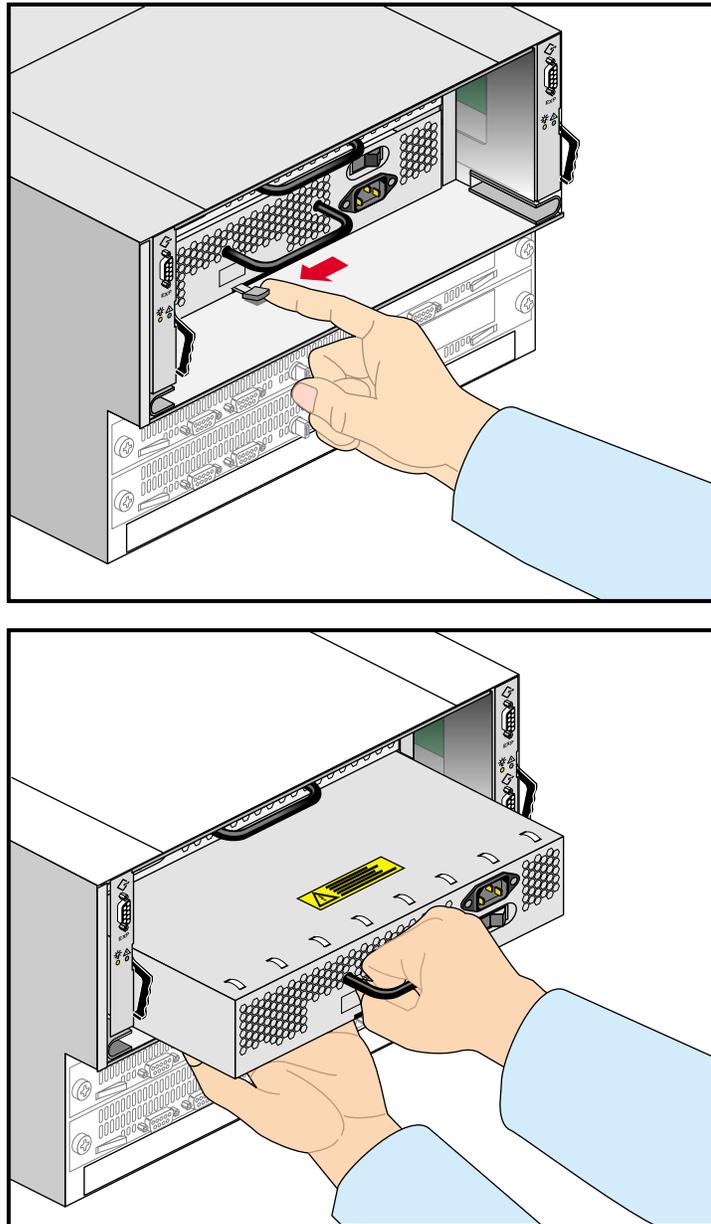


Figure 4-12 Removing an Existing Fibre Channel RAID Power Supply or Filler Panel

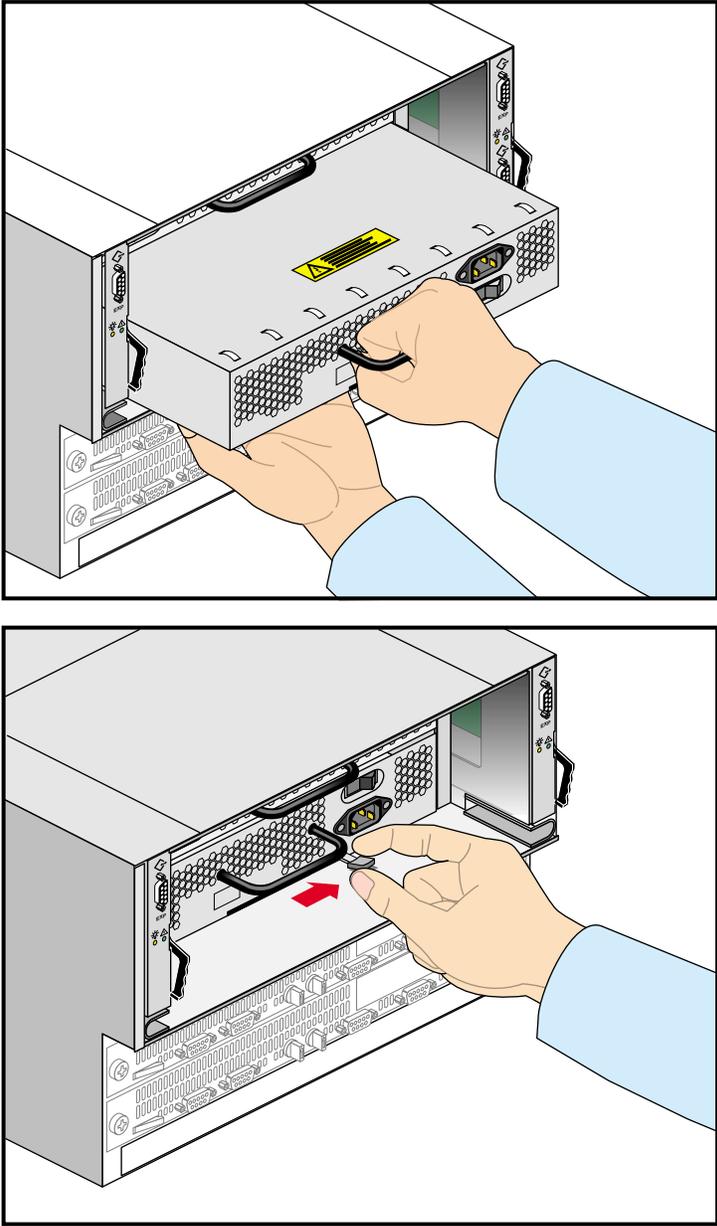


Figure 4-13 Installing the Bottom Fibre Channel RAID Power Supply

Replacing or Adding a Link Control Card (LCC)

This section explains removing and replacing an LCC or filler panel.

Note: Do not remove an LCC or filler unless you have its replacement ready to install.

Follow these steps to remove an existing LCC or filler panel and install a new LCC:

1. Have ready the new or replacement LCC.
Caution: When handling an LCC, be careful not to touch any of its electronic circuits or components. Remove the LCC from its packaging only when ready to install it.
2. Locate the FibreVault or fibre channel RAID enclosure containing the LCC with its check LED illuminated.
3. Attach the clip of the ESD wristband to the fibre enclosure chassis and put the wristband around your wrist with the metal button against your skin.
4. Unplug the cable(s) (as applicable) connected to the faulty LCC, then pull up on the latch of the LCC or filler panel.
5. Grasp the LCC and slowly pull it out of the chassis, as shown in Figure 4-14. The LCC's active and check LEDs turn off if the enclosure is powered on.
6. Wait ten seconds to allow the fibre channel system FLARE code time to recognize that the LCC was removed.
7. Align the new LCC with the guides in the empty bay.
8. Slide the LCC into the slot, as shown in Figure 4-15.
9. Push down on the LCC's latch. The LCC's active LED illuminates if the enclosure is powered on.
10. Attach the LCC cables.
11. Remove and store the ESD wristband.

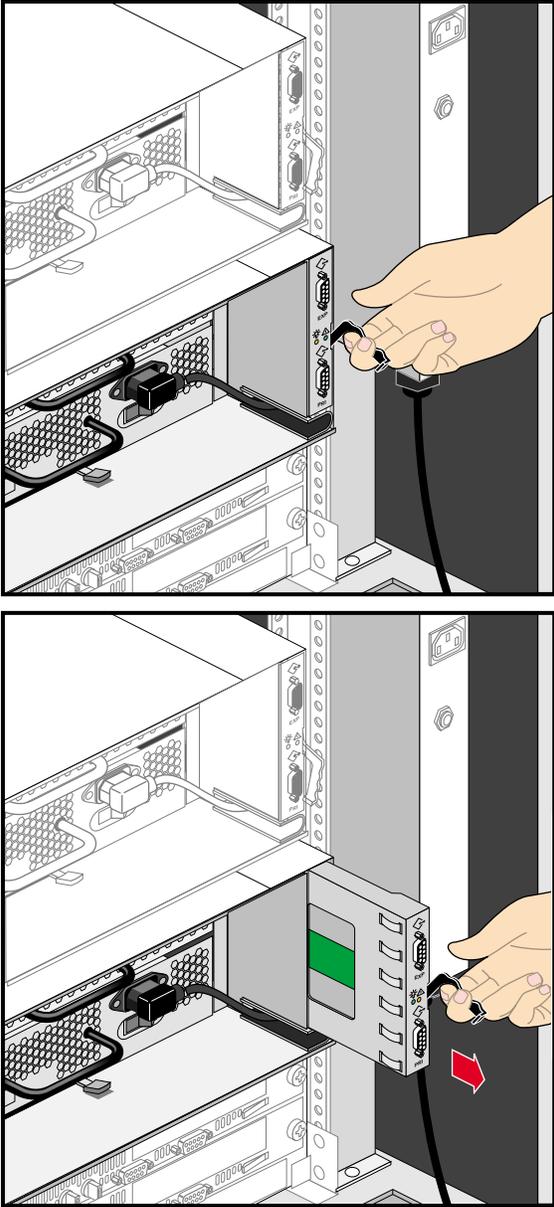


Figure 4-14 LCC Removal Example

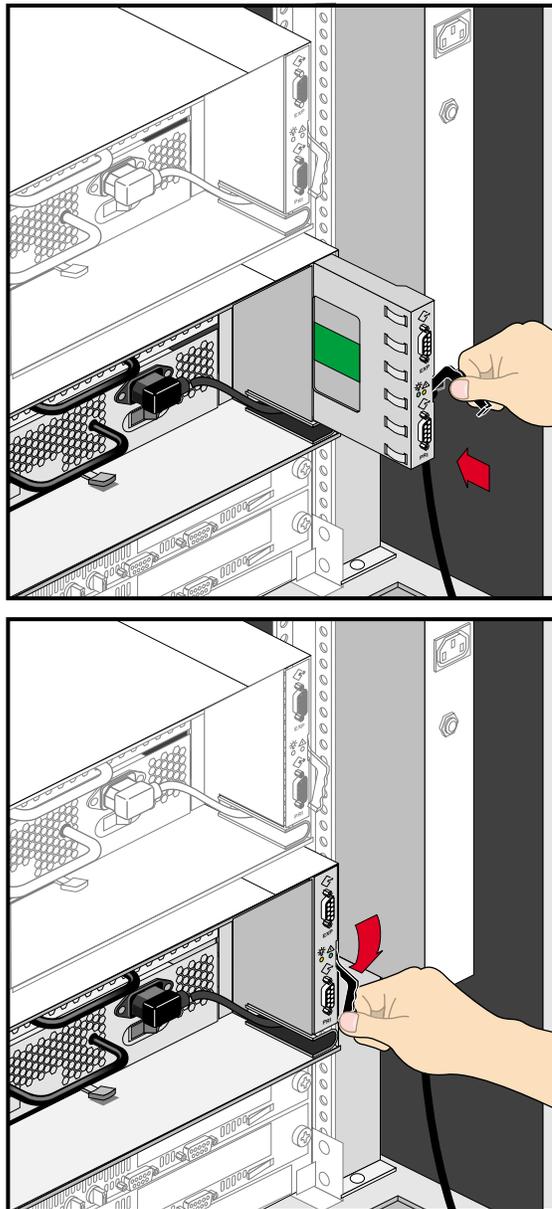


Figure 4-15 LCC Installation Example

Configuring a Fibre Channel System

The information in this chapter presumes that your fibre channel option(s) and any fibre channel storage has been set up and initially configured by a qualified and trained installer.

XIO Interface Configurations

The maximum supported number of XIO fibre channel interfaces in an Origin2000 multi-rack system is 44. Technically, you may configure up to 12 fibre channel XIO boards in a single module. Note that this may not always be practical as other XIO boards may be needed in the module. These maximums are a function of software support and do not reflect the number of XIO fibre boards that will physically fit in a system. The OCTANE workstation supports up to three optional fibre XIO boards.

Future product releases may support a different number of boards; consult your product release notes.

Table 5-1 Configuration Maximums

Maximums	Maximum Value
Fibre channel XIO boards in an Origin2000 multi-rack system	44
Initiators on an FC-AL loop	1 standard (multiple initiators only with optional FailSafe configurations)
Maximum disks on each FC-AL loop	110 non-RAID or 100 RAID
RAID storage systems on an FC-AL loop	1 RAID enclosure plus 9 RAID expansion enclosures

Getting System Status

You can use three methods to obtain system status information from a host controlling a fibre disk array:

- Use the command line interfaces (also known as the *fccli* and *ssmcli*) to get fibre system status
- Use the graphical user interface (also known as the *ssmgui*) to get RAID array status
- Use the IRIX command line to obtain hardware graph or hardware inventory information

These methods are explained in the *Origin FibreVault and Fibre Channel RAID Administrator's Guide*.

Technical Specifications

This appendix summarizes technical information and operational limits for the fibre products. Topics covered are:

- “Enclosure Height Specifications” on page 95
- “Site Requirements for the Fibre Channel Rack and FibreVault Deskside” on page 96
- “Technical Specifications for the FibreVault Enclosure” on page 97
- “Technical Specifications for the FibreVault Deskside (Tower)” on page 98
- “Technical Specifications for the Fibre Channel RAID Enclosure” on page 99
- “Technical Specifications for the Fibre Channel RAID Deskside Tower” on page 100
- “Technical Specifications for the Fibre Channel Rack” on page 102
- “Cabling Interfaces” on page 103

Enclosure Height Specifications

Table A-1 gives the heights of the various enclosures that install in the fibre channel rack, where 1U = 1.75 inches (4.45 cm).

Table A-1 Enclosure Heights

Component	Number of Us	Inches	Centimeters
Fibre channel rack (internal)	39 U	68.25	173.35
FibreVault or fibre channel RAID expansion	3.5 U	6.125	15.56
Optional standby power supply (SPS)	1 U	1.75	4.45

Table A-1 (continued) Enclosure Heights

Component	Number of Us	Inches	Centimeters
Fibre channel RAID enclosure	6.5 U	11.375	28.9
RAID enclosure with standby power supply (SPS) in rack	7.5 U	13.125	33.33

Caution: The operating limits for temperature and humidity must not be exceeded inside the closed rack in which a fibre enclosure is mounted. Mounting fibre equipment in a rack directly above or below an enclosure does not restrict air flow to an enclosure because air flows through it from front to back.

Site Requirements for the Fibre Channel Rack and FibreVault Deskside

Before installing the fibre channel rack, make sure the site meets the power, thermal, and cabling requirements for the fibre rack and its enclosures. See Table A-6 for specific fibre channel rack information.

FibreVault deskside tower systems have essentially the same power requirements and technical specifications as the FibreVault enclosures in a rack system. Use the technical information in Table A-2 for all questions regarding the deskside FibreVault.

Enclosure Power Requirements

To determine an enclosure’s power requirements, always use the power rating that is on the label on the chassis as the final authority. This rating is the maximum power required for a fully loaded chassis.

The input current, power (VA), and dissipation per enclosure are based on the maximum capability of the power supplies and cooling system to provide internal regulated power. Typical values may be less, depending on the number and manufacturer of disk drives.

A failure of one of the two power supplies in the enclosure results in the remaining supply and cord supporting the full load. The FC rackmount cabinet's AC power distribution can handle these values for the number of enclosures that you will interconnect. The site must also have main branch AC distribution that can handle these values.

Enclosure Thermal Requirements

The ambient temperature specification is measured at the enclosure's front door inlet. The site must have air conditioning of the correct size and placement to maintain the specified ambient temperature range. The air conditioning must be able to handle the BTU requirements of the enclosure(s).

Technical Specifications for the FibreVault Enclosure

Table A-2 lists specifications for the FibreVault enclosure. Note that these specifications are virtually identical for the deskside (stand-alone) FibreVault enclosure.

Table A-2 FibreVault Enclosure Technical Specifications

Classification	Specification	Value
AC and current	AC line voltage	100 to 240 VAC single phase (47 to 63 Hz)
	AC line current	4.0 amps maximum at 100 VAC
	Input power	392 watts maximum (305 watts average)
	Inrush current	50.0 amps maximum for 0.5 line cycles (per supply)
	AC circuit breaker	8 amp (thermal)
	AC inlet type	IEC320-C14 appliance coupler
	Drive module power	25 watts per bay maximum
Operational limits	Ambient temperatures	10°C to 40°C (50°F to 104°F)
	Temperature gradient	10°C per hour (50°F per hour)
	Relative humidity	20% to 80% noncondensing

Table A-2 (continued) FibreVault Enclosure Technical Specifications

Classification	Specification	Value
Non-operating limits	Max ambient temperatures at higher elevations	40°C (104°F) at 2,438 m (8,000 ft) 37°C (98.6°F) at 3,050 m (10,000 ft)
	Heat dissipation	1411 x 10 ³ J/hr (1,340 BTU/hr) maximum
	Ambient temps	-40°C to 65°C (-40°F to 149°F)
	Temp gradient	24°C/hr (43.2°F/hr)
Dimensions	Relative humidity	10% to 90% noncondensing
	Elevation	7,625 m (25,000 ft)
	Height	15.4 cm (6.06 in.) 3.5U (NEMA units)
	Width	44.5 cm (17.5 in.)
Weights	Depth	63.27 cm (24.9 in.) front door to back of fan pack
	Minimum weight	26.4 kg (58 lbs) without redundant components
	Maximum weight	35.5 kg (78 lbs) in high-availability configuration

Technical Specifications for the FibreVault Deskside (Tower)

Table A-3 lists specifications for the stand-alone deskside tower version of the FibreVault. All environmental operating parameters are the same as those for the rackmounted FibreVault. Height, weight, and dimensions are the primary differences between the two products. Most other specifications are the same.

Table A-3 Deskside FibreVault Enclosure Technical Specifications

Classification	Specification	Value
AC and current	AC line voltage	100 to 240 VAC single phase (47 to 63 Hz)
	AC circuit breaker	8 amps (thermal)
	AC inlet type	IEC320-C14 appliance coupler (3 outlets)
Dimensions	Tower height	67.95 cm (26.75 in.)

Table A-3 (continued) Deskside FibreVault Enclosure Technical Specifications

Classification	Specification	Value
	Tower width without anti-tip legs	26.67 cm (10.5 in.)
	Depth	74.7 cm (29.5 in) front door to back of unit
Weights	Maximum weight	60.0 kg (132 lbs) fully configured
	Minimum weight	50.1 kg (112 lbs) minimum configuration

Technical Specifications for the Fibre Channel RAID Enclosure

Table A-4 lists specifications for the fibre channel RAID enclosure.

Table A-4 Fibre Channel RAID Enclosure Technical Specifications

Classification	Specification	Value/Type
AC and current	AC line voltage	100 to 240 VAC single phase (47 to 63 Hz)
	AC line current	8.0 amps maximum at 100 volts
	Input power	700 watts maximum (560 watts average)
	Inrush current	50.0 amps maximum for 0.5 line cycle (per supply)
	AC circuit breaker	12 amps (thermal) in each power supply
	AC inlet	IEC320-C14 appliance coupler
	Drive module power	25 watts per bay maximum
Operational limits	Ambient temperatures	10°C to 40°C (50°F to 104°F)
	Temperature gradient	10°C per hour (50°F per hour)
	Relative humidity	20% to 80% noncondensing
	Max operational ambient temperatures at higher elevations	40°C (104°F) at 2,438 m (8,000 ft) 37°C (98.6°F) at 3,050 m (10,000 ft)
	Heat dissipation	2520 x 10 ³ J/hr (2,390 BTU/hr) maximum

Table A-4 (continued) Fibre Channel RAID Enclosure Technical Specifications

Classification	Specification	Value/Type
Non-operating limits	Ambient temperature	-40°C to 65°C (-40°F to 149°F)
	Temperature gradient	24°C/hr (43.2°F/hr)
	Relative humidity	10% to 90% noncondensing
	Elevation	7,625 m (25,000 ft)
Dimensions	Height	28.59 cm (11.25 in.) 6.5U (NEMA units)
	Width	44.5 cm (17.5 in.)
	Depth	70.02 cm (27.57 in.) front door to drive fan pack
Weights	Minimum weight	41.4 kg (91 lbs)
	Maximum weight	52 kg (114 lbs)

Technical Specifications for the Fibre Channel RAID Deskside Tower

The deskside fibre channel RAID tower houses one fibre channel RAID (DPE) and one RAID expansion FibreVault (DAE). It may contain as few as five or as many as 20 fibre disk drives. All figures cited are maximum estimates unless otherwise specified below in Table A-5.

Table A-5 Fibre Channel RAID Deskside Tower Technical Specifications

Classification	Specification	Value/Type
AC and current	AC line voltage	100 to 240 VAC single phase (47 to 63 Hz)
	AC line current	12.0 amps maximum at 100 volts
	Input power	1,200 watts maximum (865 watts average)
	Inrush current	80.0 amps maximum for 0.5 line cycle
	AC circuit breaker	15 amps (thermal) at deskside AC inlet (PDU)
	AC inlet	IEC320-C14 appliance coupler
	Drive module power	25 watts per bay maximum

Table A-5 (continued) Fibre Channel RAID Deskside Tower Technical Specifications

Classification	Specification	Value/Type
Operational limits	Ambient temperatures	10°C to 40°C (50°F to 104°F)
	Temperature gradient	10°C per hour (50°F per hour)
	Relative humidity	20% to 80% noncondensing
	Max operational ambient temperatures at higher elevations	40°C (104°F) at 2,438 m (8,000 ft) 37°C (98.6°F) at 3,050 m (10,000 ft)
	Heat dissipation	3,920 x 10 ³ J/hr (3,710 BTU/hr) maximum
Non-operating limits	Ambient temperature	-40°C to 65°C (-40°F to 149°F)
	Temperature gradient	24°C/hr (43.2°F/hr)
	Relative humidity	10% to 90% noncondensing
	Elevation	7,625 m (25,000 ft)
Dimensions	Height	68 cm (26.8 in.)
	Width	52.1 cm (20.6 in.)
	Depth	74.7 cm (27.57 in.) front door to drive fan pack
	Min. service clearances	30.3 cm (1 ft.) front, 60.6 cm (2 ft.) rear
Weights	Minimum weight	76.5 kg (168.6 lbs) with all disks, LCCs, SPSs, and power supplies removed
	Maximum weight	152 kg (335 lbs) maximum configuration

Technical Specifications for the Fibre Channel Rack

Table A-6 lists specifications for the fibre channel rack.

Table A-6 Fibre Channel Rack Specifications

Classification	Specification	Value
AC power requirements	Cabinet voltage	200 VAC to 240 VAC, single-phase 47 Hz to 63 Hz
	Current draw per enclosure @200VAC	2.3 amps per FibreVault or RAID expansion 4.1 amps per Fibre Channel RAID enclosure
	Power consumption use per enclosure	400 watts max (305 watts average) per FibreVault 700 watts max (560 watts average) per RAID
Connector	Type	Domestic: L6-30P International: IEC Plug 309 2-P
Operating limits	Ambient temp.	10°C to 40°C (50°F to 104°F)
	Relative humidity	20 to 80% noncondensing
	Max ambient temperatures at higher elevations	40°C (104°F) at 2,438 m (8,000 ft) 37°C (98.6°F) at 3,050 m (10,000 ft)
	Heat dissipation	3,168 x 10 ³ J/hr (3,000 BTU/hr) max per chassis assembly
	Shock	3 g @ 11 ms
Nonoperating limits (shipping and storing)	Vibration	0.25 g peak @ 5 Hz to 500 Hz
	Ambient temperature	-40°C to 65°C (-40°F to 149°F)
	Gradient, maximum	24°C/hr (43.2°F/hr)
	Relative humidity	10% to 90% noncondensing
Physical	Elevation	7,625 m (25,010 ft)
	Cabinet dimensions	Height: 192 cm (75.3 in.)
		Width: 65 cm (25.5 in.)

Table A-6 (continued) Fibre Channel Rack Specifications

Classification	Specification	Value
		Depth: 87 cm (34.25 in.)
	Approx. max. wt.	Approx. 500 kg (1,100 lbs):
Service clearance	Front access	91.5 cm (36 in)
	Rear and side access	76.2 cm (30 in)
Buses	External host fibre bus	Fibre channel (up to 1-Gb-per-second transfer rate)

Cabling Interfaces

Because the native interface for the FC-AL disk drives is copper, the Silicon Graphics fibre channel option uses a copper interface with a female DB9 connector.

To support longer distances and provide a fallback in case of grounding and EMI issues with copper, an optional fiber optic interface using a media interface adapter is available.

The optical interface uses a shortwave laser (CD-ROM laser) with a wavelength of 780 nm. Silicon Graphics supports cable lengths of 25 to 300 meters (from the host to the fibre enclosure).

This interface is also be used with 62.5 μm optical cable and the total length supported is 300 meters (about 60% compared to the 50 μm cable). Silicon Graphics fibre products support the 62.5 μm optical cable and a maximum distance of 300 meters.

Many buildings are already wired with 62.5 μm multimode optical cables.

Caution: To maintain acceptable quality and signal integrity, use only Silicon Graphics supplied copper and optical cables with your fibre enclosures.

Table A-7 summarizes characteristics for the copper twinax and fiber optic cables.

Table A-7 Cable Characteristics

Characteristic	Copper Twinax	Fibre Optic
Variant	100-TW-EL-S full duplex (quad conductor)	62.5 μm Type: 780 nm
Range	10 to 25 meters	up to 300 meters
Transmitter	PECL	100-M5-SL/N-I
Medium	150 Ohm	0.002-0.5
Connector	Male DB9	SC duplex connector
Maximum length	Host system to enclosure: 25 m. Enclosure-to-enclosure: 0.3 m or 0.5 m	Host system to enclosure: 300 m. Fiber optic cable is not used for enclosure-to-enclosure links.

The copper twinax cable is a fully shielded, full-duplex, balanced cable capable of supporting distances of 25 meters at the 1.0625-Gbaud transfer rate. Distances greater than 10 meters require equalization in the cable connector.

A copper cable connecting a fibre enclosure and the host system is either 10 or 25 meters. Cables longer than 10 meters must be equalized. All copper fibre channel cables shipped with Silicon Graphics products are equalized. FibreVault in an Origin2000 or Onyx2 rack always uses a 3-meter copper cable. See the sections “Fibre Channel Board Logic” on page 13 and “Copper or Optical Media Interface Adapter” on page 16 in Chapter 1 for more information on cabling.

Fibre enclosures are connected (linked) by 0.3-meter copper cables (0.5-meter cables are used in the Origin rack implementation). They must never be longer than 10 meters. Fibre enclosure interconnections should maintain LCC consistency. One FC loop should interconnect all and only the LCC As, and the other FC loop should interconnect all and only LCC Bs.

The connector for both ends of the twinax cable is a male DB9 connector. The shield on each end of the cable connects directly to the connector shield, providing a DC ground on both ends. With a system and a fibre enclosure on separate building grounds, a potential for grounding problems exists when this cabling is connected with a copper cable utilizing a DC ground shield on both ends. In this case, use the optional fibre optical cabling to link host and enclosure.

There are a number of fibre channel interface cables available from Silicon Graphics. Table A-8 lists information regarding some of these optional components.

Table A-8 Silicon Graphics Fibre Interface Cables

Cable Use	Cable and Connector Type	Cable Length	Silicon Graphics Part #
LCC-to-LCC	Copper with DB9	0.3 m (11 inches)	9470156
LCC-to-LCC	Copper with DB9	0.5 m (19.7 inches)	Shipped with FibreVault in Origin kit only
Origin2000 or Onyx2 host to internal FibreVault enclosure	Copper with DB9	3 meters (10 feet)	018-0742-00x
XIO or PCI (host) to fibre enclosure	Copper with DB9	10 meters (32.8 feet)	018-0570-00x
XIO or PCI (host) to fibre enclosure	Copper with DB9	25 meters (82 feet)	018-0571-00x
XIO or PCI (host) to fibre enclosure	Optical cable	25 meters (82 feet)	018-0656-201
XIO or PCI (host) to fibre enclosure	Optical cable	100 meters (328 feet)	018-0656-301
XIO or PCI (host) to fibre enclosure	Optical cable	300 meters (984 feet)	018-0656-401

Glossary

Active

The state of a Sequence Initiator until all the Data frames for the Sequence have been transmitted. The state of a Sequence Recipient until all the Data frames for the Sequence have been received.

address identifier

An address value used to identify source (S_ID) or destination (D_ID) of a frame.

alias address identifier

One or more address identifiers which may be recognized by an N_Port in addition to its N_Port Identifier. An alias address identifier is Fabric unique and may be common to multiple N_Ports.

attenuation

The transmission medium power loss expressed in units of dB.

average power

The optical power measured using an average reading power meter when the fibre channel is transmitting a specified code sequence as defined in the test procedure.

C/E

Copper cable-to-electrical converter module. A C/E module can be used in conjunction with an MIA module option to provide a common physical footprint so that the copper or optical options do not need to be factory configured. See *MIA*.

Class 1, 2, 3

Classes of fibre channel connection:

- Class 1 provides a dedicated connection between 2 N_Ports. Class 1 provides end-to-end flow control, buffer-to-buffer flow control, and frame acknowledgments, with in-order-of-frame delivery guaranteed.
- Class 2 is a frame-switched, connectionless service allowing bandwidth to be shared by multiplexing frames from multiple sources onto the same channel. Class 2 provides end-to-end flow control, buffer-to-buffer flow control, and frame acknowledgments, but in-order-of-frame delivery is not guaranteed.
- Class 3 is similar to class 2 except there is no end-to-end flow control or delivery acknowledgments.

equalization

In copper cables, equalization can be used to compensate for high-frequency losses by attenuating lower frequencies to match the loss of higher frequencies. Equalization is applied at the receiving end of a signal, and is typically tuned to the cable length and data rate. Cables that have been equalized should not be connected to other cables.

exchange

An exchange is composed of one or more nonconcurrent sequences for a single operation. The exchanges can be unidirectional or bidirectional between two N_Ports. Within a single exchange, only one sequence can be active at any one time, but sequences of different exchanges can be concurrently active.

Fast 20 SCSI

The Fast 20 SCSI parallel interface, which can burst data at 40 MB/sec. Also known as Ultra SCSI and Hyper SCSI.

fabric

A fibre channel fabric is implemented by switches connecting F_Ports. A nonblocking switch that allows full connectivity between ports. An FC fabric allows different classes of service (Classes 1, 2, 3).

FC-AL

Fibre Channel Arbitrated Loop, a low-cost implementation of fibre channel for connecting disk drives in a loop.

FC-PH

Fibre channel physical (hardware) interface, which includes FC-0, FC-1, and FC-2.

FC-0

Fibre channel physical interface, that is, media type, electrical and optical characteristics, connectors, cables, and so on.

FC-1

Defines the fibre channel transmission protocol, including 8b/10b encoding, error detection.

FC-2

Defines the fibre channel framing protocol.

FC-3

Defines the fibre channel common services, for example, striping, hunt groups, multicast.

FC-4

Defines the fibre channel upper-layer protocol, that is, SCSI-3.

frame

The basic building block of an FC connection, a frame contains the information to be transmitted, the address of the source and destination ports, and link control information. Frames are of two types, Data frames and Link_control frames. Data frames can be used as Link_Data frames and Device_Data frames. Link_control frames are classified as Acknowledge (ACK) and Link_Response (Busy and Reject) frames.

Fibre Channel

Fibre channel is a general term used to refer to the fibre channel interface standard as a whole, and includes FC-PH, FC-AL, and FC 0-4. The word Fibre is used as a generic term which can indicate copper or optical implementations of fibre channel products.

hub

A hub provides a physical star topology while maintaining a logical loop that provides minimal disruption to the loop in the event of a host failure or during reconfiguration.

LRC

Loop Redundancy Circuit. See *PBE*.

MIA

Media interface adapter module (also referred to as a media shifter, optical-to-electrical converter, or O/E). An MIA module converts photons to electrons in one direction, and converts electrons to photons in the other direction.

Origin Fibre Channel RAID

A fibre channel enclosure with an intelligent storage processor (SP) controller board that supports RAID operations. Available as a rackmount enclosure or deskside (tower).

Origin FibreVault

A fibre channel chassis a disks only (disk array enclosure) with no controllers. It is not a RAID enclosure. Available as a standalone deskside (tower) unit or as a chassis mountable in the fibre channel rack or the Origin and Onyx2 rack.

PBE

Port Bypass Enable (also known as Loop Redundancy Circuit (LRC)). A high-speed multiplexer which is used to bypass a failed or not present disk drive in a disk enclosure. A PBE typically redrives but does not retime the signals, which adds a small amount of jitter. PBEs are controlled by a port bypass enable signal, which is normally provided by a disk drive (or the lack of a disk drive) or an environmental interface.

SCA 40

The fibre channel SCA (Single Connector Attach) connector is a 40 pin connector that is used for connecting disk drives directly to a backplane. This connector features a ground mate first connection for hot plug support, and blind mate features. All power, ground, dual interface, addressing, and option support pins are provided in one connector. This connector is similar to the 80 pin SCSI SCA connector.

SCSI 3

SCSI 3 refers to the new SCSI architecture, which includes Parallel SCSI, SSA, or fibre channel.

sequence

Transmission between fibre channel buffers. A sequence is formed by a set of one or more related frames transmitted unidirectionally from one N_Port to an other. Each frame within a sequence is uniquely numbered with a sequence count. Error recovery, which is controlled by an upper protocol layer, is usually performed at sequence boundaries.

SSA

Serial Storage Architecture. SSA is an emerging serial storage interface promoted by IBM which is a 20 MB/sec native interface, soon to be doubled to 40 MB/sec.

striping

An FC-3 common service, striping is multiplying bandwidth using multiple N_Ports in parallel to transmit a single information unit across multiple links. Not yet implemented.

_Port: N_Port, L_Port, F_Port, NL_Port, FL_Port

Ports for

- N_Port: can be connected to a node only
- L_Port: can be connected to a loop only
- F_Port: can be connected to a fabric only
- NL_Port: can be connected to either a node or a loop
- FL_Port: can be connected to either a fabric or a loop

XIO bus

The XIO bus consists of two unidirectional 400 MHz source synchronous 8-bit (400 MB), or 16 bit (800 MB) buses. Although the Bridge ASIC can be connected to a 8- or 16-bit Crosstalk bus, the bus itself uses only 8 bits.

Index

C

cabling, 103-105
channels defined, 2-3
conventions, xv

D

data
 reconstruction, 7
 redundancy, 7
 transfer in fibre channel, 36
datagram service class, 41
disk module, 18-20, 65
 adding, replacing, 66-71
 and SCSI-2 disk module, 43
 non-RAID, 43, 66
 number of pins, 43
 part number, 66
 RAID, 43, 66
 slot numbering, 73-74
 stacking, 20
 swapping, 9
 Ultra SCSI, 65

E

enclosure
 disk guidelines, 68
enclosure address, setting, 71
exchange in FC-2, 40

F

fan module, 53
 removing, 64
 replacing, 75
 See also FibreRAID storage processor board
 assembly fan module
FC-0, 37
FC-1, 38
FC-2, 38-41
FC-3, 42
FC-4, 42
FC-AL, 3-4, 33-35
fibre channel
 administration, 29-30
 applications, 5-9
 arbitrated loop, 3-4
 architecture, 31-42
 configuring the system, 93
 drivers, 30

- enclosure
 - heights, 95
 - power requirements, 96
 - site requirements, 96-97
 - thermal requirements, 97
 - host interface option boards, 9
 - interface, maximum, 93
 - layers, 37-42
 - option board, 10-17
 - block diagram, 15
 - chassis grounding, 17
 - interface diagram, 10
 - logic, 13-15
 - port, 3, 33-35
 - RAID
 - agent, 29
 - RAID implementation, 5
 - standard, 1-2
 - storage, 18, 43-63
 - fibre channel RAID, 52
 - deskside specifications, 101
 - fan module, 53
 - FibreRack, 24-25, 58
 - connectivity diagram, 58
 - height, 95
 - shutting down, 62-63
 - specifications, 102-103
 - weight, 58
 - FibreRAID, 21-24
 - height, 95
 - key, 22
 - power supply
 - removing, replacing, 84-88
 - specifications, 99-100
 - storage processor board assembly, 54-57
 - connectors, 55
 - fan module, 56-57
 - FibreVault, 21-22, 44-52
 - deskside, 44-45
 - specifications, 98-99
 - fan assembly, 49
 - height, 95
 - in FibreRACK, 43
 - key, 22
 - midplane, 52
 - power supply, 50-51
 - filler, removing, 80-83
 - inserting, 83-84
 - removing, 80-83
 - rackmount, 46
 - specifications, 97-98
 - flow control in FC-2, 41
 - frame in FC-2, 40
- H**
- high availability and RAID, 6
 - hot spare, 7
 - and RAID 3, 68
- I**
- IRIS FailSafe, 6
- K**
- key for FibreVault, FibreRAID, 22

L

LCC, 6, 47-49
 adding, replacing, 89-91
LEDs, 61-62
link control card. *See* LCC

M

media interface adapter, 103
media interface adapter. *See* MIA
MIA, 24

N

networks defined, 2-3

O

optical cable, 103
ordered set in FC-2, 39

P

PCI fibre board, 12
physical address. *See* enclosure address
power supply
 adding, replacing, 77
 preparing new, 78-79
 See also FibreVault power supply, FibreRAID
 power supply
 standby, height, 95
protocol in FC-2, 41

R

RAID
 high availability, 5
 hot spare, 8
 parallel access array, 8
RAID-3, 7
RAID support levels, 8

S

SCSI-3 protocol, 43
sequence in FC-2, 40
service class 3, 41
specifications, 95-105
SPS
 front panel, 59
 LEDs, 59
ssmagent, 29
system status, 61-62, 94

T

troubleshooting, 61-62

U

Ultra SCSI disk module, 65

X

XIO fibre controller boards, 18

Tell Us About This Manual

As a user of Silicon Graphics products, you can help us to better understand your needs and to improve the quality of our documentation.

Any information that you provide will be useful. Here is a list of suggested topics:

- General impression of the document
- Omission of material that you expected to find
- Technical errors
- Relevance of the material to the job you had to do
- Quality of the printing and binding

Please send the title and part number of the document with your comments. The part number for this document is 007-3428-004.

Thank you!

Three Ways to Reach Us

- To send your comments by **electronic mail**, use either of these addresses:
 - On the Internet: techpubs@sgi.com
 - For UUCP mail (through any backbone site): *[your_site]!sgi!techpubs*
- To **fax** your comments (or annotated copies of manual pages), use this fax number: 650-932-0801
- To send your comments by **traditional mail**, use this address:

Technical Publications
Silicon Graphics, Inc.
2011 North Shoreline Boulevard, M/S 535
Mountain View, California 94043-1389