Origin™ FibreVault™ and Fibre Channel RAID Installation Instructions

Document Number 108-0154-001

Contributors

Written by Carolyn Curtis and Mark Schwenden
Illustrated by Cheri Brown, Dan Young, and Carolyn Curtis
Production by Linda Rae Sande
Engineering contributions by Sammy Wilborn, Paul Stolle, Henry Ortiz, Ron Naminski, James Bergman, Frank Demcak, Jeremy Higdon, and Paul Tsien

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

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 his own expense.

Attention

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

Origin FibreVault and Fibre Channel RAID Installation Instructions
Document Number 108-0154-001

Silicon Graphics, Inc.
Mountain View, California

Contents

Figures

Tables

About This Guide

The FibreVault™ enclosure storage system provides a compact, high-capacity, high-availability source of disk storage for Silicon Graphics® Origin2000™ servers and Onyx2™ supercomputing graphics workstations.

The fibre channel storage system uses disk technology that allows you or the customer to replace a disk while the system and enclosure(s) continue to run. This type of "hot swapping" must be done only under controlled circumstances.

For maximum storage, the Fibre Channel Rack storage system offers room for up to 110 non-RAID disks. The enclosures in a rack can be connected to one or more host system fibre XIO boards separately or in combination (loop). Fibre Channel RAID products are supported in this release.

A command-line interface is used to administer and monitor the non-RAID fibre disk (also known as JBOD) arrays. A JBOD ("just a bunch of disks") array has no RAID abilities or restrictions.

In addition to a command-line interface, a graphical user interface is also available for administering Fibre Channel RAID disk arrays.

Note: This document is intended for use by Silicon Graphics trained and authorized service personnel only. The information is structured for an installer with knowledge of the Origin2000 and Onyx2 product lines. For detailed installation information on the host systems, see the *Origin2000 and Onyx2 Deskside and Rackmount Installation Instructions*.

Installers interested in additional information regarding fibre channel technology should read the *Origin FibreVault and Fibre Channel RAID Owner's Guide* (P/N 007-3428-xxx).

Structure of This Document

This guide contains the following chapters:

- Chapter 1, "Fibre Channel Component Overview," illustrates and describes all field replaceable (FRU) fibre channel components.
- Chapter 2, "Installing a Fibre Channel Rack," describes how to set up a Fibre Channel Rack, how to add or replace enclosures in a rack, and how to replace other rack field replaceable components.

- Chapter 3, "Fibre Channel XIO Board Installation in the Host System," provides the information necessary for installing the XIO board in an Origin2000 or Onyx2 host system.
- Chapter 4, "Deskside FibreVault Disk Tower," gives an overview of the deskside non-RAID disk tower. Basic unpacking and installation information is covered.
- Chapter 5, "Fibre Channel Component Replacement," provides removal and replacement procedures for the basic hardware components that are field replaceable units (FRUs).
- Appendix A, "Technical Specifications," summarizes technical and site installation information for the fibre channel disk modules, enclosures, and rack.
- Appendix B, "Fibre Channel Command-Line Administration," documents the basics of each command line interface (CLI). There is a CLI for non-RAID (JBOD) disks and a separate CLI for RAID disk arrays.
- Appendix C, "Graphical User Interface Administration," documents the graphical user interface (GUI) used to administer and monitor RAID disk arrays only.

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
`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) on-screen menu items and references in text to document section titles
[]	(Brackets) surround optional syntax statement arguments

Chapter 1

Fibre Channel Component Overview

This chapter introduces, summarizes, and discusses the components used when installing, repairing, or replacing a fibre channel storage system or its subsystems.

Note: Before installing or significantly upgrading a fibre channel storage system, make sure that the installation site meets the air conditioning, AC power requirements, and operating limits listed in Appendix A.

1.1 Field-Replaceable Components and Cables

In the Silicon Graphics fibre storage implementation, the two basic types of cable used are copper (which is standard) and optical (which is optional). These cables are further divided into two fibre channel arbitrated loop (FC-AL) functions in a system.

- Host-to-fibre enclosure cabling (copper or optical)
- Enclosure-to-enclosure cabling (always copper, always short)

The cables that link enclosures together are always copper and are short. The optional host-to-enclosure optical cables are longer (minimum of 25 meters). Table 1-1 provides some examples. Note that a single FC-AL supports enclosures in a single rack. Do not extend the loop to another rack (even when fewer than 110 drives are installed).

Additional basic cabling guidelines include:

- Always connect the first FC-AL to port A on System Processor A (SP A) on a Fibre Channel RAID enclosure.
- Never disconnect a cable from an operating enclosure.
- Never connect a cable from LCC A to LCC B.

1.1.1 Cable Labels

Each fibre channel XIO board ships with at least one 10-meter copper cable, and label kit for identifying fibre channel connection cables.

Cable labels identify which controller on which XIO fibre board is supporting a particular enclosure, fibre rack, or loop.

A number of optional fibre interface cables are available from Silicon Graphics. Table 1-1 lists information regarding these components. If parts are missing, or if incorrect parts are included in the kit, log a call with the Technical Assistance Center (TAC).

1.1.2 Power Cord Information

The Fibre Channel Rack power cable destination kit is DK-FC1-002. All fibre enclosures in the rack use the same three-prong connector to plug to the rack's power distribution units (PDUs).

Table 1-1 lists marketing codes for power cord destination kits for the deskside chassis. Always consult your local site electrician for proper electrical connection and wiring requirements. See Appendix A for technical specifications.

Note: All Fibre Channel Racks come equipped with two single-phase cables and L6-30P connectors.

1.1.3 Field-Replaceable Units List

Table 1-1 lists field-replaceable units (FRUs) that trained service personnel are qualified to replace.

Table 1-1 Fibre Channel Marketing Codes and FRU Part Numbers

Module	Marketing Code	Part Number
Fibre Channel Rack (empty)	P-F-RACK	9470157
FC deskside (tower) enclosure with 10 drives (non-RAID)	P-F-B10X9-DS	9470128
FC deskside (tower) enclosure (empty)	P-F-Vault-DS	9470126
Additional FibreVault or vault expansion LCC card	P-F-LCC	9470144
Fibre Channel RAID LCC card	P-F-RAID-LCC	9470145
Additional power supply for FibreVault (w/cord)	P-F-PWR	9470143
Fibre Channel RAID power supply (w/cord)	P-F-RAID-PWR	9470153
Two MIA modules	P-F-OE-KIT	9980952
FC box rackmount kit (includes FibreVault and RAID rails)	P-F-RM-KIT	
0.3-m FC copper cable with 9-pin DIN	X-F-COP-0.3M	9470156
10-m FC copper cable with 9-pin DIN	X-F-COP-10M	018-0570-00x
25-m FC copper cable with 9-pin DIN	X-F-COP-25M	018-0571-00x
25-m FC optical cable (uses P-F-OE-KIT)	X-F-OPT-25M	018-0656-201
100-m FC optical cable (uses P-F-OE-KIT)	X-F-OPT-100M	018-0656-301
300-m FC optical cable (uses P-F-OE-KIT)	X-F-OPT-300M	018-0656-401

Table 1-1 (continued) Fibre Channel Marketing Codes and FRU Part Numbers

Module	Marketing Code	Part Number
Destination kit for standalone peripherals (110V)	DK-PT1-001	
Destination kit for Origin Fibre Channel Rack (220VAC/1P)	DK-TA2-001	
FC disk module (9.1 GB), non-RAID (512 bytes/sector)	P-F-9	9470140
Rackmount FibreVault enclosure with 10 drives	P-F-B10X9-R	9470125
Fibre Channel RAID (8.8 GB) disk module (520 bytes/sector)	P-F-RAID-1X9	9470138
Set of 5 Fibre Channel RAID (8.8 GB) disk modules	P-F-RAID-5X9	9470142
Rackmounted Fibre Channel RAID enclosure with 5 RAID disks	P-F-RAID-B5X9-R	9470134
Rackmounted Fibre Channel RAID expansion with 5 RAID disks	P-F-RAID-E5X9-R	9470136
Additional RAID system processor (SP) controller board	P-F-RAID-CNTRL	9470146
SP cache (8 MB single DIMM), 4 required	P-F-RAID-C8	9470150
Rack PDU AC Power Cord Intl (220V)	DK-FC1-002	
Destination kit peripherals, 220V Thailand	DK-P4-002	9350033
Destination kit peripherals, 240V USA	DK-P4-003	9350027
Destination kit peripherals, 220V Germany	DK-P4-004	9350018
Destination kit peripherals, 220V UK	DK-P4-005	9350017
Destination kit peripherals, 220V Australia	DK-P4-007	9350032
Destination kit peripherals, 220V Switzerland	DK-P4-008	9350026
Destination kit peripherals, 240V Italy	DK-P4-009	9350025
Empty disk sled assembly	P-F-SLED	9470174

1.2 Fibre Channel Enclosures

Silicon Graphics fibre storage expansion systems use enclosures and disk modules that are for fibre applications only. Do not attempt to install enclosures or disk modules into other Silicon Graphics mass-storage products.

This section provides a general introduction to the enclosures that support the Silicon Graphics fibre disk storage options.

1.2.1 FibreVault Enclosure

Figure 1-1 shows the rackmount FibreVault (also known as the DAE).

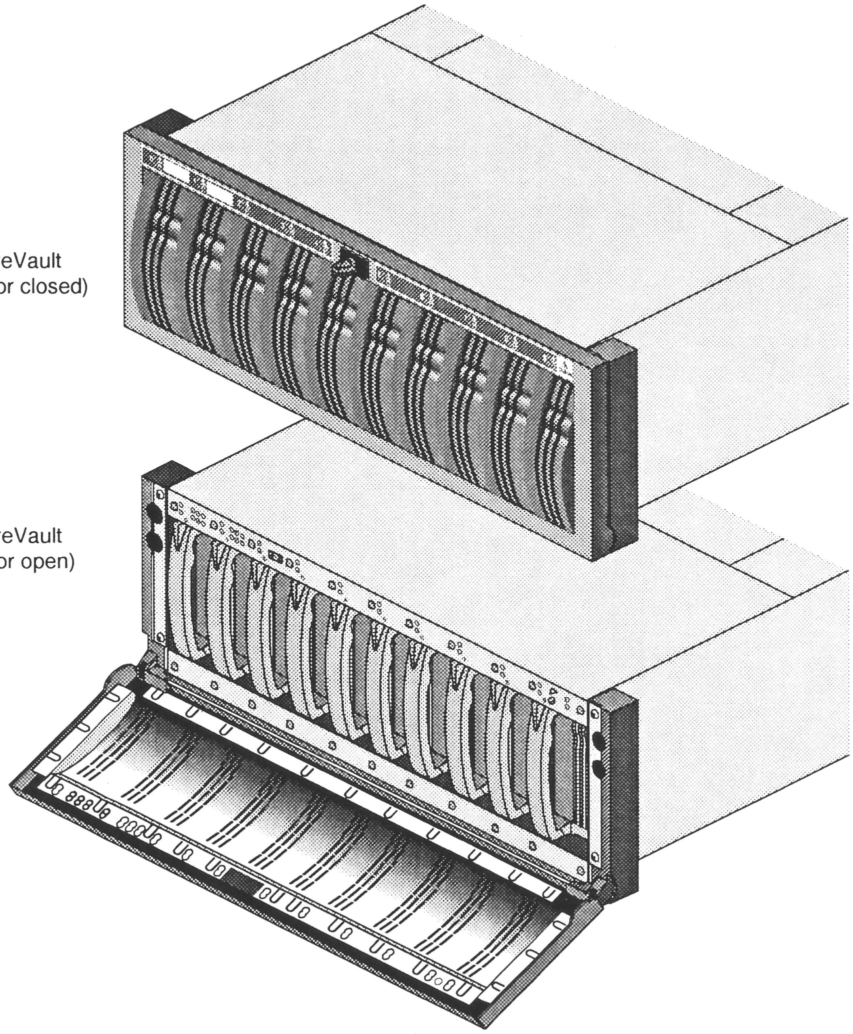

Figure 1-1 FibreVault (Front View)

Note that the Fibre Channel RAID expansion enclosure looks the same as the FibreVault enclosure. The difference is that the Fibre Channel RAID expansion uses RAID formatted drives and is connected to and controlled by a storage processor in a Fibre Channel RAID enclosure. Sectors on non-RAID drives are 512 bytes and RAID disks use 520-byte sectors. The two types of drives are *not* interchangeable. You must use the Fibre Channel RAID expansion enclosure only with Fibre Channel RAID storage systems and disk drive modules. Each drive module has an identifying sticker with its part number:

- P/N 9470138 for (8.8 GB) RAID drives (520-byte sectors)
- P/N 9470140 for (9.1 GB) non-RAID drives (512-byte sectors)

If you are unsure whether an enclosure is a FibreVault or Fibre Channel RAID expansion, shut it down, pull a drive module(s) and check the part number(s).

Figure 1-2 shows both sides of an example fibre disk module and carrier that would install in a fibre enclosure.

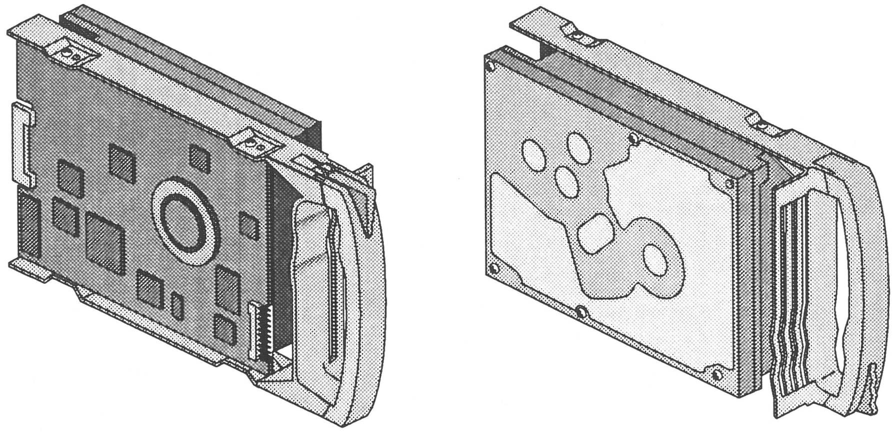

Figure 1-2 Fibre Channel Disk Module

The rackmount FibreVault consists of

- the rackmount enclosure
- one to ten fibre channel disk modules
- one or two power supplies (auto-ranging)
- one or two link control cards (LCCs) with two 9-pin connectors
- one fan module containing three fans

  Note: The enclosure will continue to operate with one fan offline. In this case, replace the fan module as soon as possible.

1.2.2 Deskside FibreVault

The deskside (tower) version of the FibreVault enclosure (see Figure 1-3) is a standalone, vertically oriented non-RAID disk enclosure that does not require rack mounting. General installation and configuration information for this product is covered in Chapter 4, "Deskside FibreVault Disk Tower."

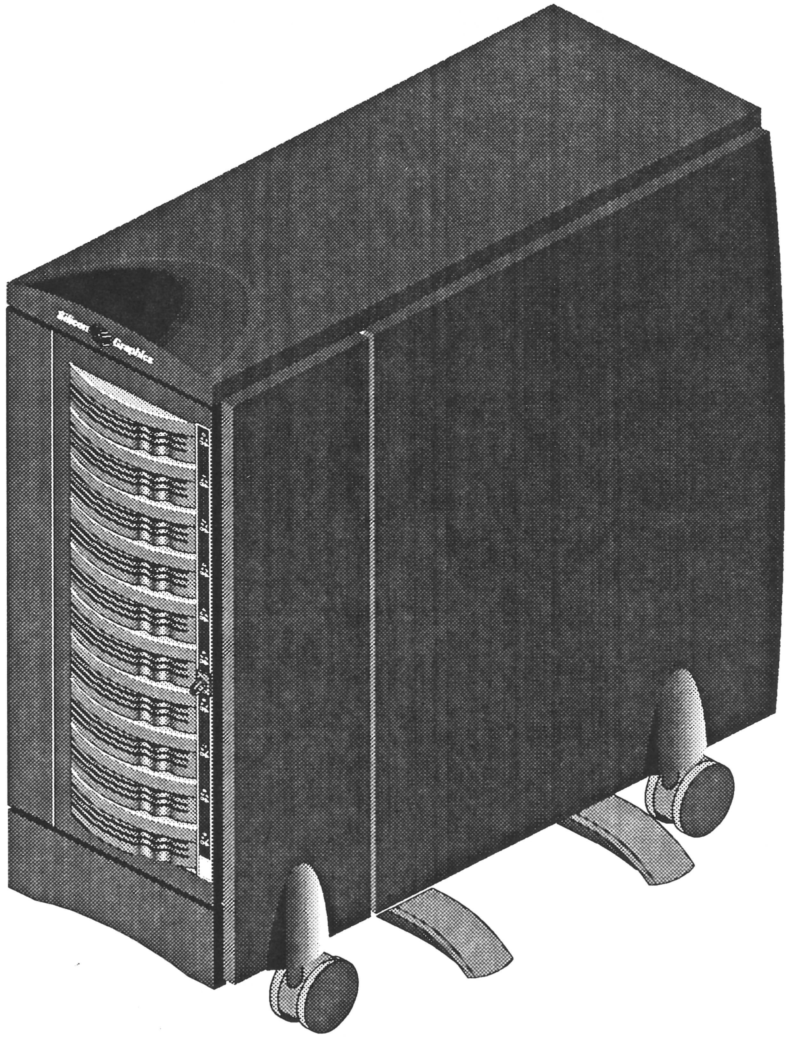

Figure 1-3 Deskside FibreVault (Tower)

The deskside FibreVault consists of the following major components:

- the front door
- the main enclosure (mounted on wheels and using anti-tip side legs)
- up to ten fibre disk modules
- one or two link control cards (LCCs)
- one or two power supplies
- one drive fan pack
- one power distribution unit

1.2.3 Fibre Channel RAID Enclosure

The Fibre Channel RAID enclosure (also known a the DPE) shares a number of functionally similar components with the FibreVault enclosure, including the link controller card (LCC) and power supply. Note that while the function and appearance of the LCCs and power supplies is similar between FibreVault and Fibre Channel RAID, they are not interchangeable.

Figure 1-4 shows a front view of the rackmountable Fibre Channel RAID enclosure.

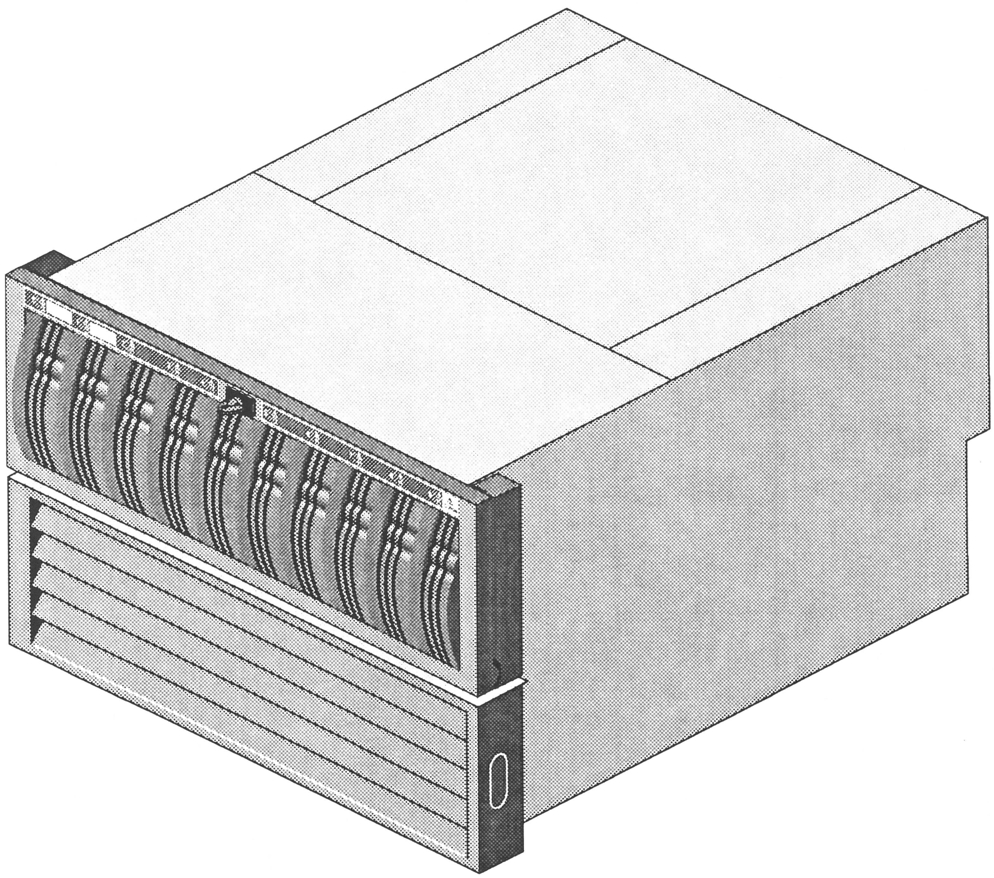

Figure 1-4 Fibre Channel RAID Rackmount Enclosure (Front View)

The Fibre Channel RAID enclosure consists of the following major components:

- the rackmount enclosure
- one to ten fibre channel disk modules
- one or two power supplies (auto-ranging)
- one or two link control cards (LCCs) with one 9-pin (**EXP**) connector on each board
- one or two storage processor (SP) RAID controller boards

- one disk fan module containing three fans (this component is interchangeable with the fan module in the FibreVault enclosure)
- one SP fan module containing three fans

1.2.4 Fibre Channel Rack

The fibre channel rack (Fibre Channel Rack) holds up to 11 FibreVault disk enclosures. Note that while it is possible to fit up to 5 Fibre Channel RAID enclosures in the rack, a "practical" configuration would be 3 units per Fibre Channel Rack. Installing a maximum of 3 Fibre Channel RAID enclosures provides room for RAID FibreVault expansion enclosures or future installation of other options.

The Fibre Channel Rack uses two power distribution units to supply the fibre enclosures with power. The rack has the (high-availability) feature of two separate 200-240 VAC power cables to supply independent power to each PDU.

You may mix the types of enclosures in the rack as long as each type is supported by a separate fibre cable (loop).

Note: Never link RAID and non-RAID enclosures together to the same XIO host connector (loop).

Figure 1-5 shows a front external view of the fibre channel rack.

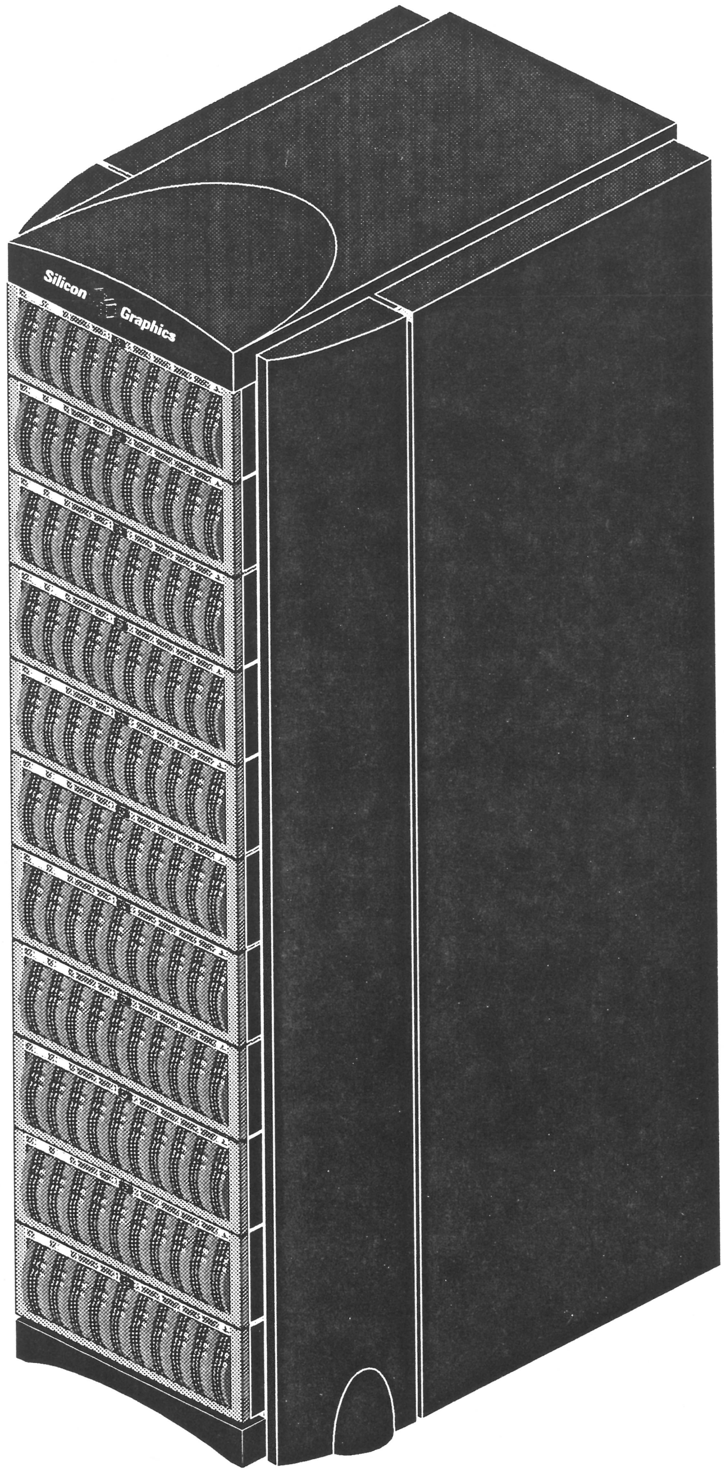

Figure 1-5 Fibre Channel Rack

Chapter 2

Installing a Fibre Channel Rack

Customers can perform basic installation and upgrade procedures such as adding or replacing disks.

The fibre channel installation and upgrade procedures covered in this chapter are performed by SSEs only:

- "Setting Up a Fibre Channel Rack,"
- "Adding Disk Enclosures to a Fibre Channel Rack,"
- "Configuring a Fibre Channel Enclosure,"
- "Replacing the Power Distribution Unit in a Fibre Channel Rack."

2.1 Setting Up a Fibre Channel Rack

Warning: **The fully loaded fibre channel rack weighs approximately 1100 lbs (500 kg). Its weight is distributed within a relatively small footprint of 6 ft² (0.56 m²), with a potential maximum floor loading of 183 lbs/ft² (893 kg/m²). Some sites have floor loadings as low as 50 lbs/ft² (244 kg/m²). Check that the flooring, especially raised flooring, can properly support the distribution of the rack's weight. Particular care must be given to installations with more than one Fibre Channel Rack placed close together.**

At least two persons are required to unpack and move a Fibre Channel Rack. The Fibre Channel Rack is prone to tipping on even a slight grade. Do not undertake this procedure unless another trained and qualified installer is available to help.

For unbolting the fibre channel rack, you need a 9/16ths-inch (14.275 mm) nut driver; the ratcheting type is recommended. Use a 3-inch (7.62 cm) extension if available. An adjustable wrench may also be used if the exact size (9/16ths-inch nut driver) is not available.

You should also have an adjustable wrench available for raising and lowering the rack's screw-down feet.

To unpack and set up a fibre channel rack, follow these steps:

1. Remove the cardboard top cover and sides, then verify that all parts are present and in good condition. If they are not, notify the shipper of the damaged or missing parts.

   Notice the ramp attached to the pallet. You'll use this ramp to move the rack off the pallet, as shown in Figure 2-1.

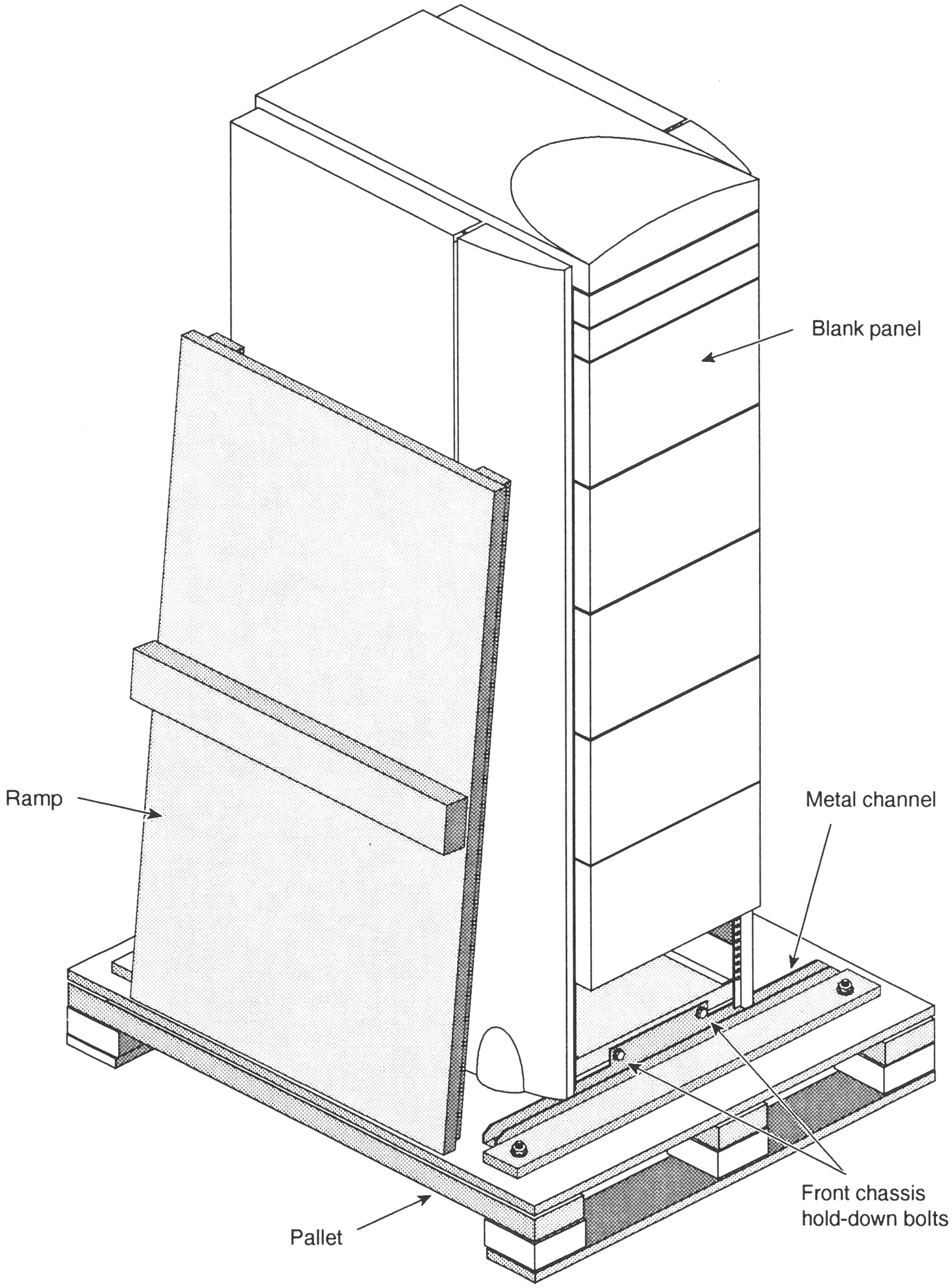

Figure 2-1 Ramp and Pallet (Shown With Plastic Removed)

2. When possible, leave the rack bolted to the pallet and use a pallet jack to move it as near the installation location as possible.

   Note: The rack can be very difficult to move on rugs; it is much easier to move it while it is still bolted to the pallet. A fully loaded rack system weighs approximately 1100 pounds (500 kg). The rack moves fairly easily on a hard floor surface.

3. Remove the plastic wrap and fold down the pallet ramp. The ramp with pallet folded down takes up a space approximately 7.5 feet (2.29 m) by 4 feet (1.22 m); you will also need another 3.5 feet (1.1 m) clear at the end of the ramp to maneuver the rack.
4. Open or remove the rear door of the rack (see Section 2.1.1) and undo the two 9/16ths-inch bolts that connect the back of the rack to the pallet-to-chassis retaining bracket. You should not need to remove the bolts that hold the bracket to the pallet.
5. Remove any filler panels that block your access to the front bolts that connect the rack and pallet-to-chassis retaining bracket, then undo the two 9/16ths-inch bolts attached to the front of the rack (see Figure 2-1).
6. Remove the four 9/16ths-inch bolts that attach the bracket to the pallet.
7. Use an adjustable wrench to loosen each of the four screw-down "feet" directly adjacent to each caster. These feet must be raised to properly roll the rack.
8. With one person in front and one in back, turn the front of the rack toward the ramp, (see Figure 2-2).

   Note: Handle the rack carefully as the side panels scratch or mar easily.

9. Position the ramp at the side of the pallet with one edge resting securely on the two feet of the pallet. Note the front casters rotate completely (the rear casters are fixed).
10. With help from your partner, carefully move the rack off the pallet using the ramp.

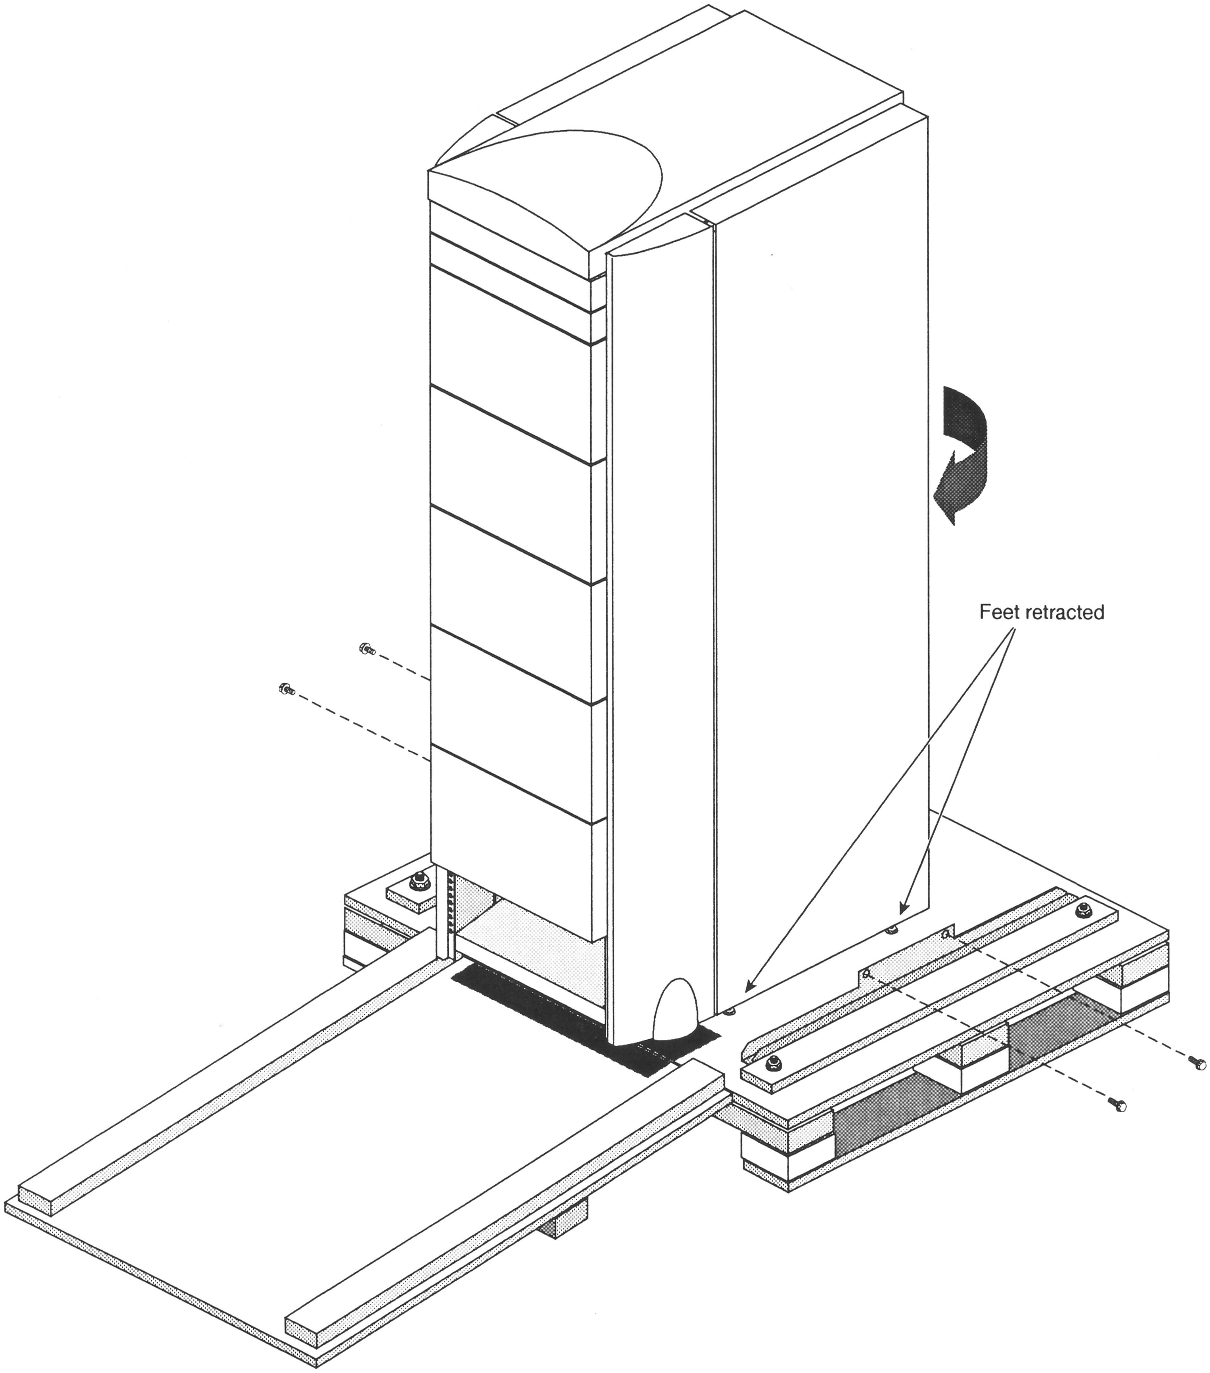

Figure 2-2 Rolling the Rack off the Pallet

2.1.1 Opening the Rear Door of a Fibre Channel Rack

To check chassis assemblies in the rack, or install enclosures, you must open the rear door.

To open the rear door of the rack, grasp the rear handle and tug the door straight outward, as shown in Figure 2-3.

Note: If you remove the rear door entirely for any procedure, rehanging the rear door might require the help of a second person because of the difficulty of lining up the door's three pinlocks.

Figure 2-4 shows the Fibre Channel Rack with its rear door open.

Caution: Proper service access requires that the installation location of the Fibre Channel Rack should have 36 inches (91.5 cm) of direct front access and 30 inches (76.2 cm) of rear space.

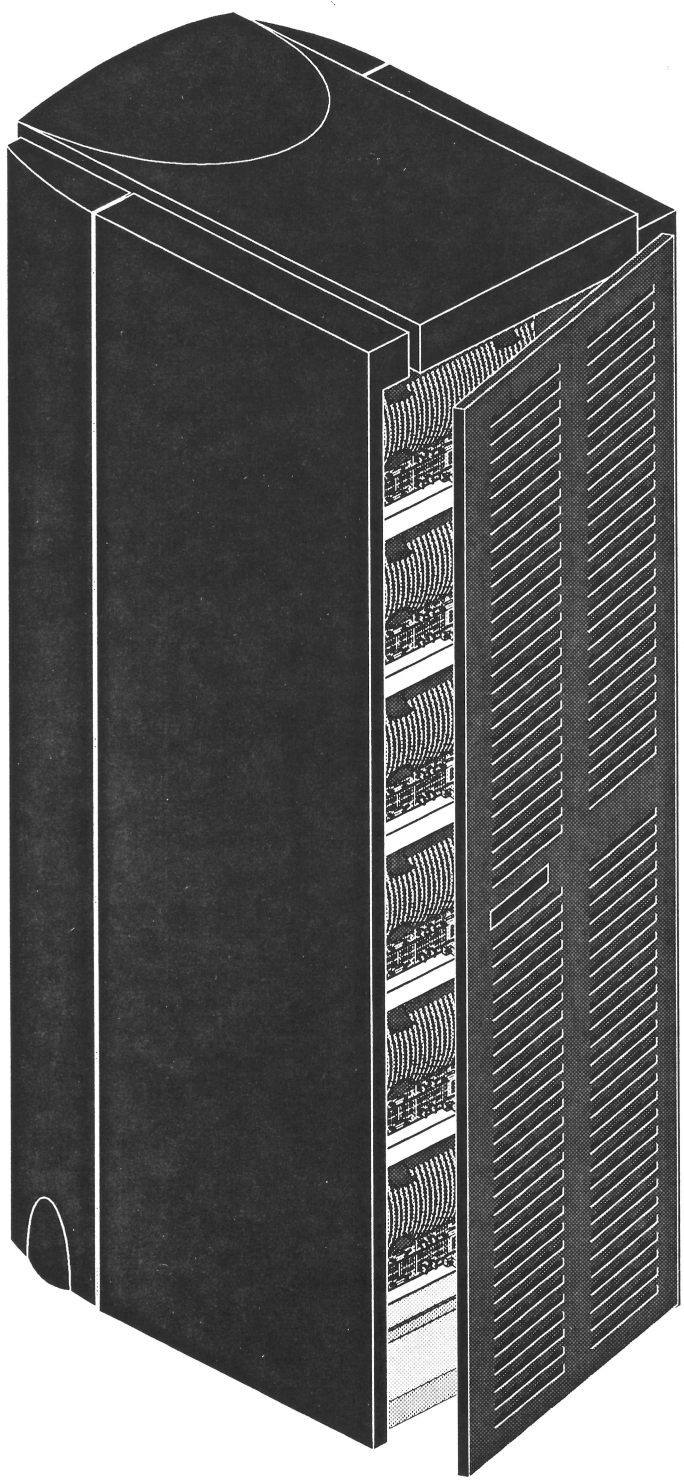

Figure 2-3 Opening the Fibre Channel Rack Rear Door

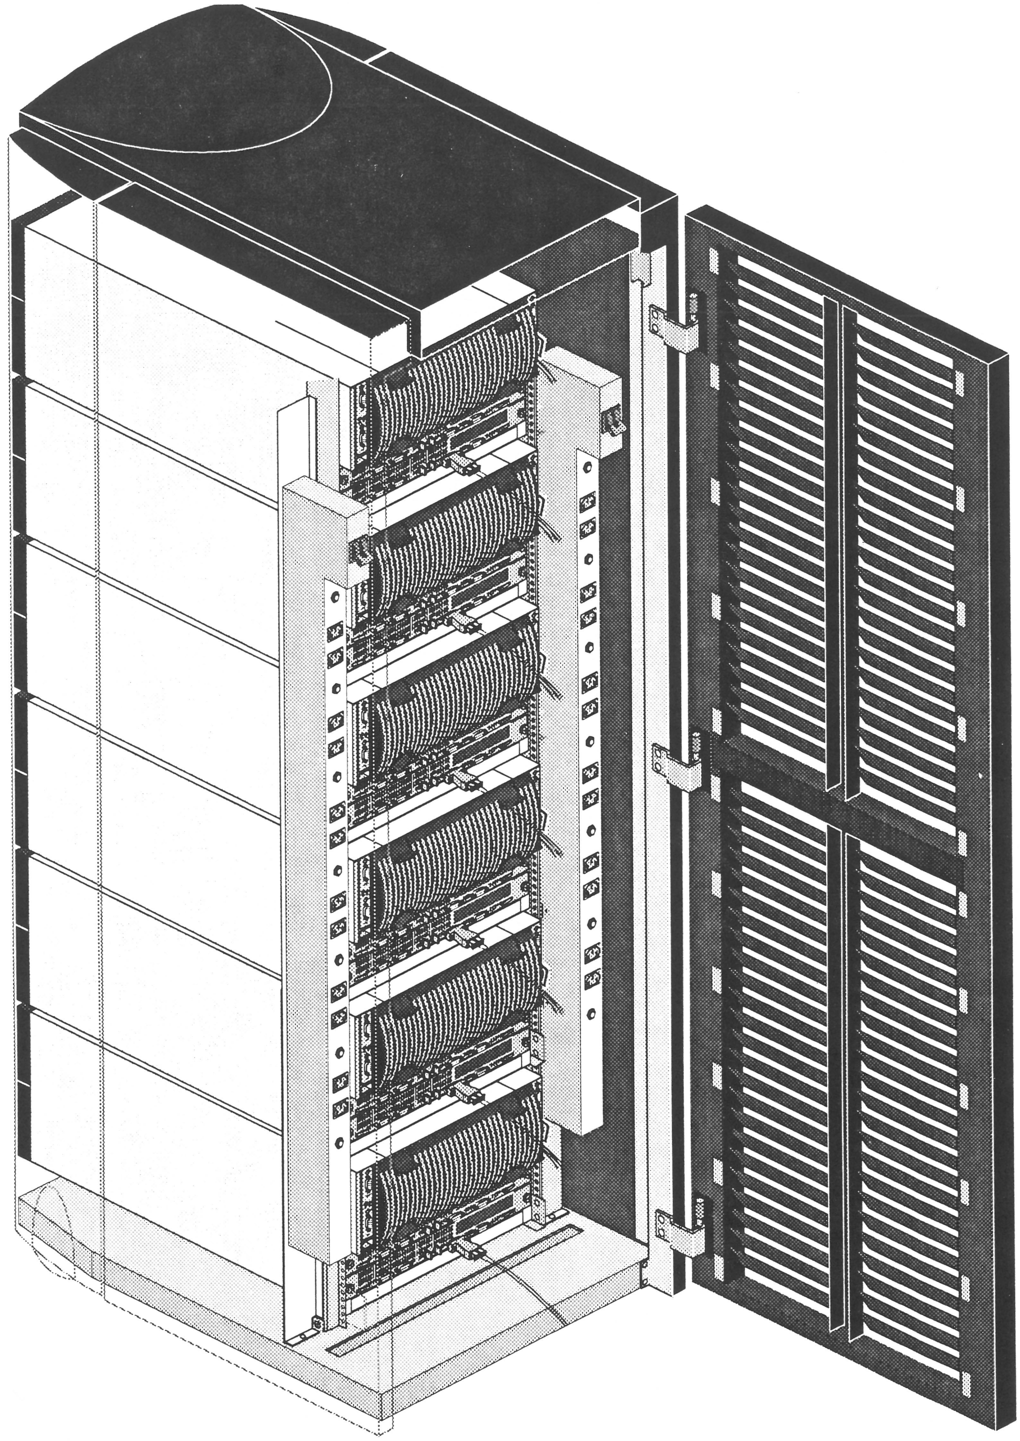

Figure 2-4 Fibre Channel Rack With Rear Door Open

2.2 Adding Disk Enclosures to a Fibre Channel Rack

To add disk enclosures to a Fibre Channel Rack:

- install the FibreVault or Fibre Channel RAID enclosures in the rack
- set the enclosure address
- connect the disk enclosure power cord(s)
- cable the disk enclosures together

These instructions do not cover the mounting hardware and procedures used to install or remove enclosures in an Onyx2 or Origin2000 rack.

Warning: **Two persons are required to install disk enclosures in a Fibre Channel Rack. Do not undertake this procedure unless another trained and qualified installer is available.**

2.3 Installing an Enclosure in the Fibre Channel Rack

Installing an enclosure is completed in three basic procedures:

- installation preparation
- positioning and installing mounting hardware
- installing and securing the enclosure
- setting the enclosure address
- connecting the power cord

2.3.1 Enclosure Installation Preparation

Prepare to install an additional FibreVault or Fibre Channel RAID enclosure in a fibre channel rack with these steps:

1. Unpack the enclosure unit and verify that all parts are present. If they are not, notify the shipper of the missing parts. Note that enclosures ship with 0, 5, 9, or 10 disks installed. Examine the disk bays carefully as the air-flow filler panels look similar to a disk module.
2. You can open the front door of each enclosure you are installing by unlocking the main latch and pushing the release. Temporarily removing the rear fan module also provides an extra hand hold location for lifting the enclosure.
3. To reduce the weight of the enclosure for insertion into the rack, you may install it without disk modules. This reduces the total weight by approximately 2 lbs. (0.9 kg) per disk module. Note that in a Fibre Channel RAID enclosure, drives 0, 1, and 2 contain a licensed internal code and *must* be reinstalled in their original positions.

4. Power off the rack before installation of a new fibre disk enclosure.
   - Stop all drives in the Fibre Channel Rack from the host system.
   - Set the rack's main circuit breaker switches to the off position (located at the top of both PDUs).
5. Starting at the bottom of the rack, remove the filler panel at the first open location. Always fill the rack with FibreVault or Fibre Channel RAID enclosures starting at the bottom, not at the top. If you are installing a 6.5 U Fibre Channel RAID enclosure, you need to remove at least two filler panels.
6. Using a #2 Phillips head screwdriver, remove the plastic and metal retaining tongue assembly from the vertical rail on each side of the rack (see Figure 2-5 for an example).

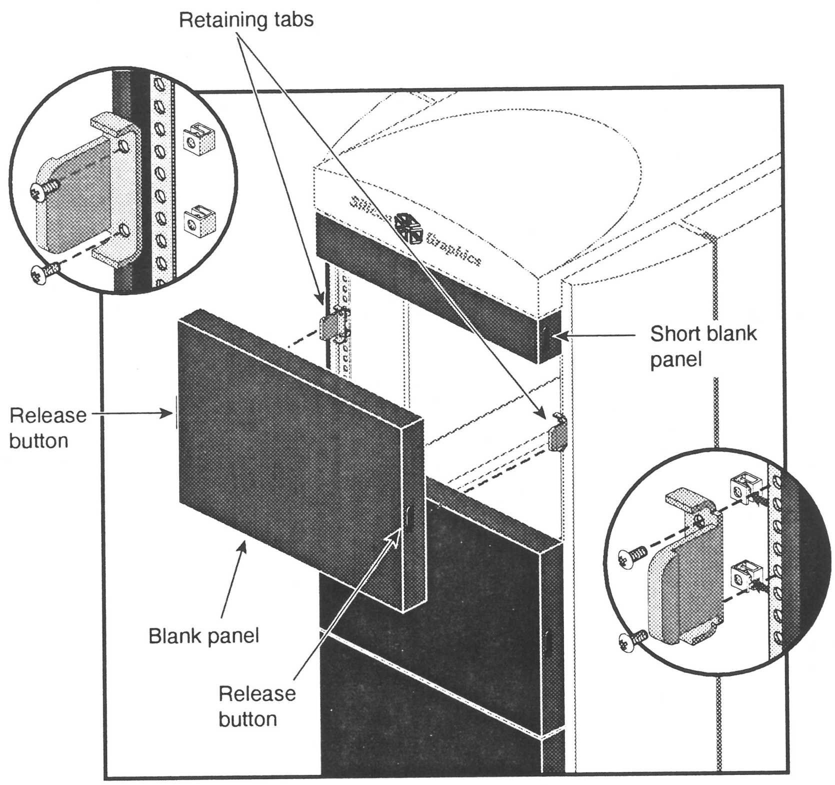

Figure 2-5 Fibre Channel Rack Filler Panel Retaining Tabs

7. Determine the appropriate mounting location for a new fibre channel enclosure using the following tips.
   - Cut a piece of sheet plastic or foam to exactly 6.5 U (11.375 inches, or 28.89 cm) or 3.5 U (6.125 inches, or 15.56 cm). Use it to determine the exact position the additional enclosure will occupy.
   - 1U (1.75 inches, or 4.45 cm) is required for installation of a RAID standby power supply (battery backup unit).

Note: When installing a Fibre Channel RAID enclosure you may want to leave a 1 U (1.75 inches, or 4.45 cm) space directly underneath the unit. This allows room for the optional installation of a standby power supply. Reserve this space if needed for future applications.

When you install the maximum number of enclosures in a rack, be sure the lowest mounting rail is in the lowest position. The bottom of the enclosure aligns with the center of the lowest rail hole when the rail is installed correctly.

2.3.2 Horizontal Shipping Brace

If you are installing an enclosure(s) into an already installed Fibre Channel Rack, skip this section and go on to "Mounting Hole Measurements" on page 2-12.

If you are installing a new Fibre Channel Rack, you will need to remove the horizontal shipping brace that comes with each new rack.

The brace is installed in the lower middle section of the front vertical attachment rails. The brace is installed to keep the vertical rails from flexing outward during shipment and handling.

Undo the four nuts and remove the horizontal shipping brace prior to enclosure installation, see Figure 2-6.

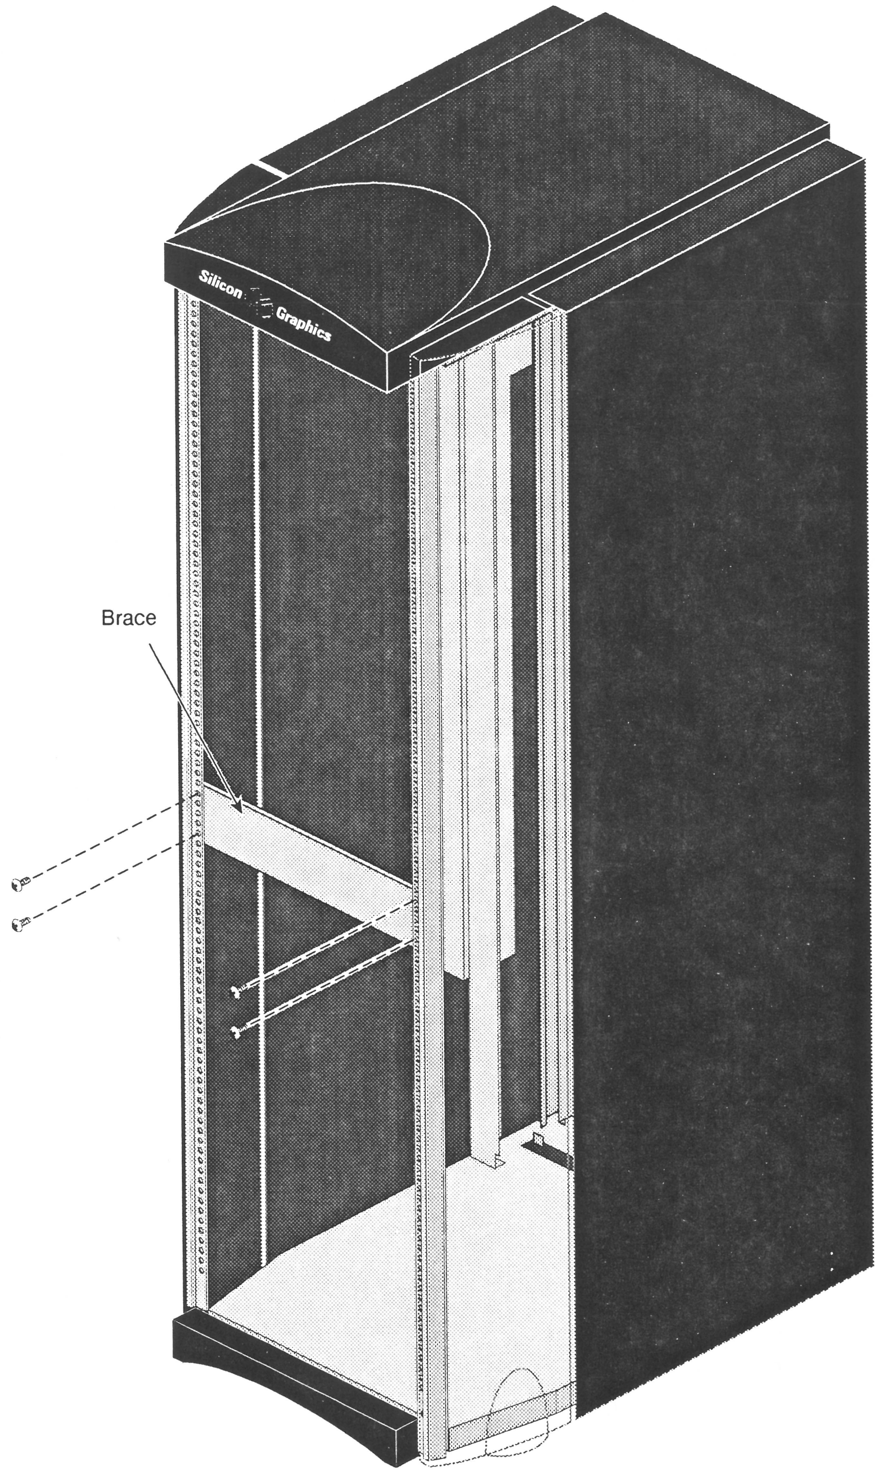

Figure 2-6 Removing the Horizontal Shipping Brace

2.3.3 Mounting Hole Measurements

The mounting holes in the Fibre Channel Rack are pre-drilled at distances of 0.5 inch (1.27 cm) and 0.625 inch (1.59 cm). As you count upward, the center of the lowest hole is 0.5 inch (1.27 cm) from the center of the next hole up. The next hole's center is 0.625 inch (1.59 cm) from the last hole's center, as is the third hole. The sequence of holes is then repeated, see Figure 2-7.

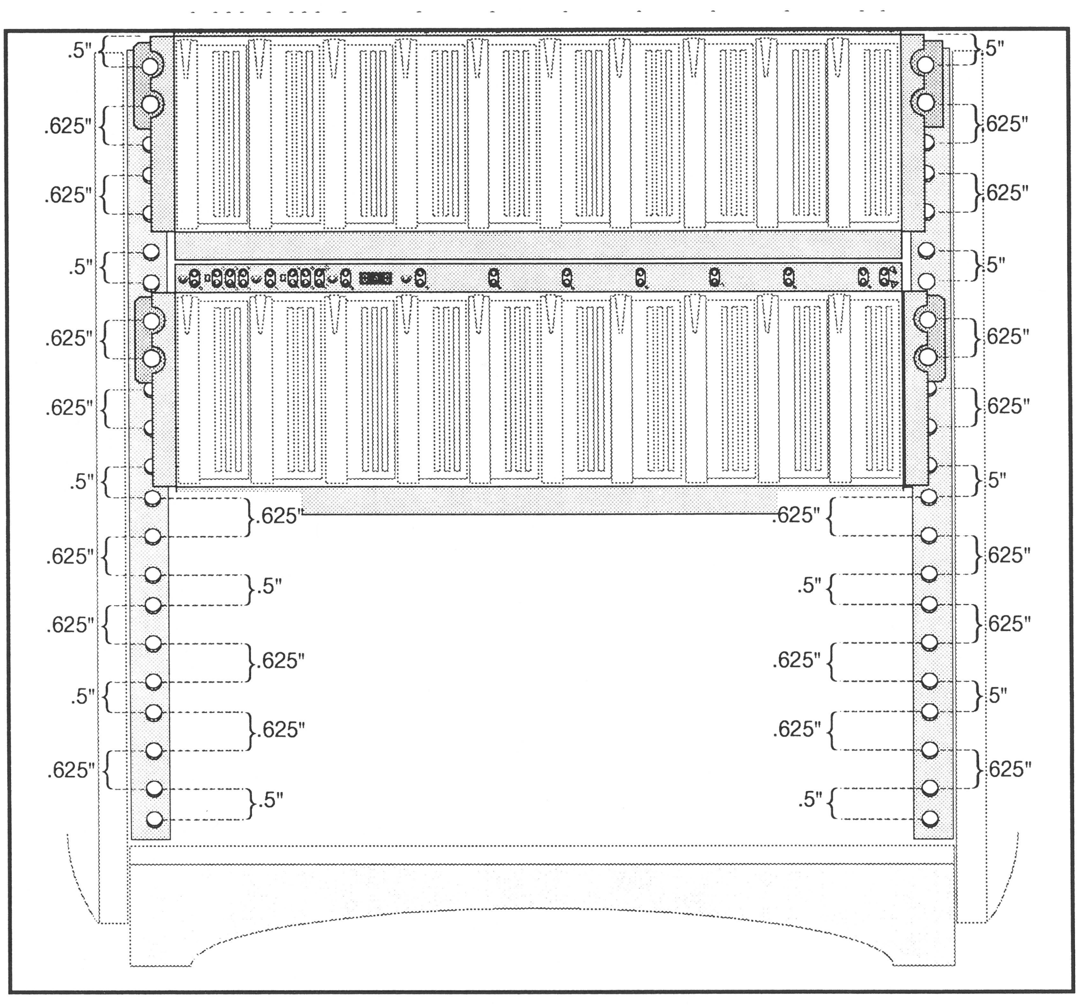

Figure 2-7 Mounting Hole Measurements

Table 2-1 provides height and hole alignment information for the Fibre Channel RAID, standby power supplies (SPS), and FibreVault products.

Table 2-1 Fibre Enclosure Height Requirements

Device (Enclosure Type)	Height	Number of Channel Holes (Starting U-Aligned)
Fibre Channel RAID enclosure (DPE)	6.5 U (11.375 in./28.89 cm)	19
Standby power supply (SPS)	1 U (1.75 in./4.45 cm)	03
FibreVault enclosure (DAE) or Fibre Channel RAID expansion enclosure	3.5 U (6.125 in./15.56 cm)	10

2.3.4 Installing the Rack's Mounting Rails

Two types of mounting rails are used with fibre enclosures. Both types are described in this section.

Installation of the mounting rails should be calculated carefully so that the Fibre Channel Rack holds the maximum number of fibre enclosures. See the information in Section 2.3.3 for measuring details.

After you select the proper mounting location (keeping in mind the measurements in Section 2.3.3), identify the proper rail hardware and prepare it for installation:

- Each FibreVault or Fibre Channel RAID expansion enclosure mounting rail uses two screws at the front and two screws plus two "clip nuts" at the rear. The rails have a distinctive "shark fin" configuration at the rear mounting end.
- The Fibre Channel RAID enclosures use a uniform (straight) rail design. The front of the mounting rail uses the same two screws as the FibreVault rails. The rear of the rails mount to the vertical channels with two nut and washer assemblies that go on threaded studs. The nuts are removed or tightened with a 5/16ths-inch (0.31 inch, or 7.9 mm) wrench or a small adjustable wrench. Be sure not to overtighten the nuts that attach the studs at the rear rails as they *will break off*.

It is generally recommended that you mount the front rail retaining screws first. See Figure 2-8 for an example.

Install the mounting rails using the following steps:

1. For easier access you can remove the cabinet rear side panels if they are in place (this is not necessary to complete the procedure).
2. Mark the lowest available U-aligned screw position on the cabinet front channels.

   If you are installing an enclosure in the first (bottom) location, the bottom of the rail should be just above the center of the lowest hole.
3. Position the front of the first mounting rail so that the screw holes are aligned.
4. Secure the two front mounting rail screws, see Figure 2-8.

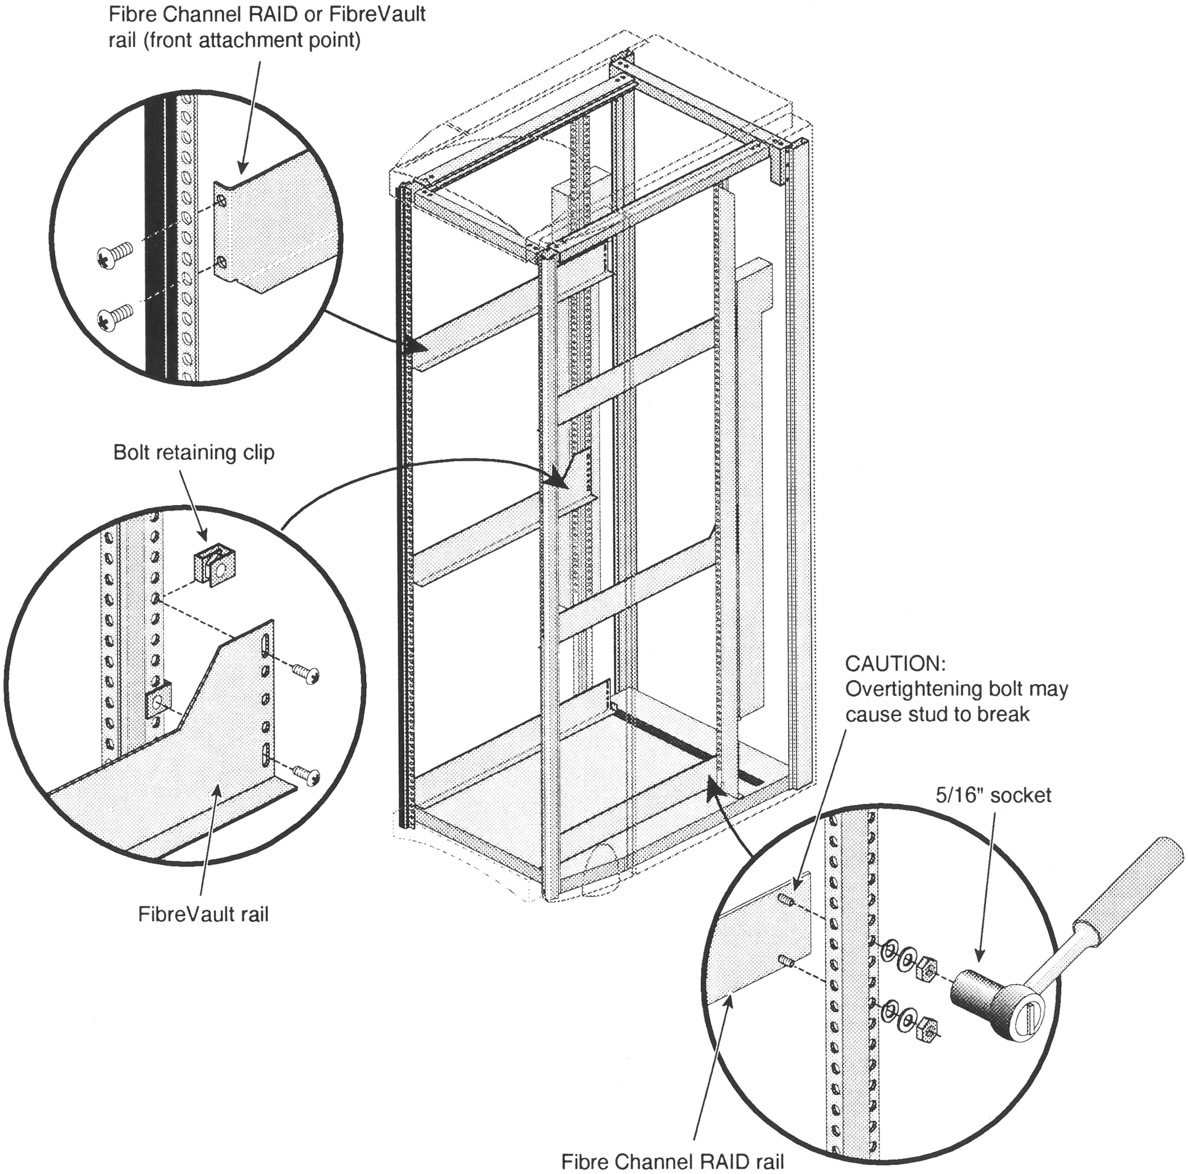

Figure 2-8 Mounting the Rails in the Fibre Channel Rack

5. Confirm that the two mounting holes for the rear of the slide rail are properly aligned.
6. Secure the rear of the mounting rail using one of the two following procedures:
   - For Fibre Channel RAID enclosures, slip the two attached screws at the rear of the rail through the proper holes, then place a washer, lock washer, and nut over each. Tighten each assembly carefully using a 5/16ths-inch (7.9 mm) box wrench or small adjustable wrench. Do not overtighten.
   - For FibreVault (or Fibre Channel RAID expansion) enclosures, slip a clip nut over each of the aligned rear holes and attach the rail with the screws provided.
7. Repeat the process with the next rail.

2.3.5 Installing a 1/2-U Filler Panel

In some rack configurations you may find that you have a 1/2-U gap between the top of an enclosure and the rack's other blank filler panels. In this case you may wish to use a 1/2-U filler panel that attaches to the top of an enclosure.

Figure 2-9 shows an example of the 1/2-U filler panel installed on a FibreVault enclosure.

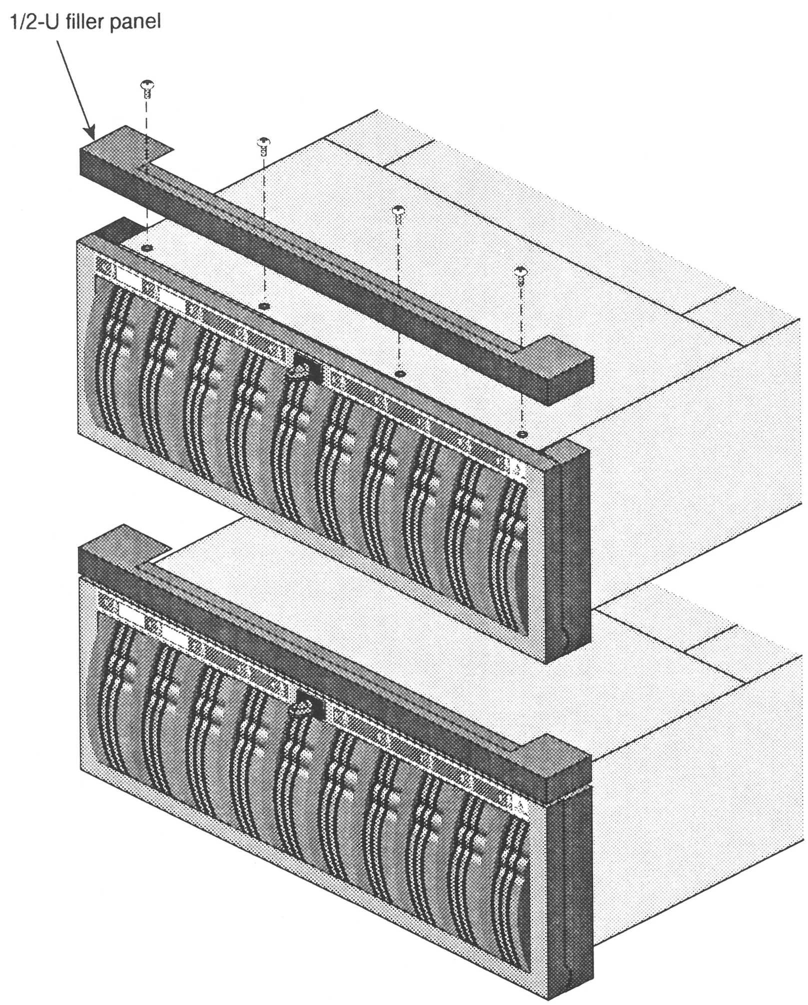

Figure 2-9 1/2-U Front Filler Panel for the Rack

2.3.6 Installing the Enclosure

Use the following information to install the fibre enclosure(s) into the rack.

Caution: The enclosures are heavy. The Fibre Channel RAID enclosure can weigh up to 110 lb. (49.6 kg) fully loaded. Never remove, replace, or install an enclosure alone.

1. Remove the rear fan module from the enclosure (on Fibre Channel RAID remove the front panel and fan module also) see "Replacing a Fan Assembly" in Chapter 5 if you need instructions. Removing the fan makes the enclosure much easier to grasp and maneuver.
2. Working with another installer, lift the enclosure you want to install and engage the rails with the bottom-rear of the enclosure. Be careful not to damage the LCC handles or scrape the new enclosure against the rack or other enclosures.

   Note: When installing an enclosure in the middle of several others (as when replacing an existing unit) you must tilt the front of the enclosure downward approximately 45 degrees before engaging the bottom sheet metal with the mounting rails.
3. Slide the enclosure into the rack as shown in the example in Figure 2-10.
4. Secure the enclosure in the rack using one retaining screw per rail.
   - The Fibre Channel RAID enclosure (Figure 2-10) secures in the rack at the rear of each rail.
   - On the front of the FibreVault, note that there are two openings for screws on each side. Only one aligns properly with a screw hole. Insert a clip nut in the proper position on each of the rack's vertical mounting bars and secure the enclosure.
5. Install any disk modules that were removed earlier.
6. Replace any fan modules removed. Note that you may want to leave the rear fan module(s) out until you have installed the enclosure power cords. See Section 2.3.9, "Connecting the Enclosure Power Cord(s)" for more information.
7. Repeat steps 1 through 6 for any additional enclosures to be installed.

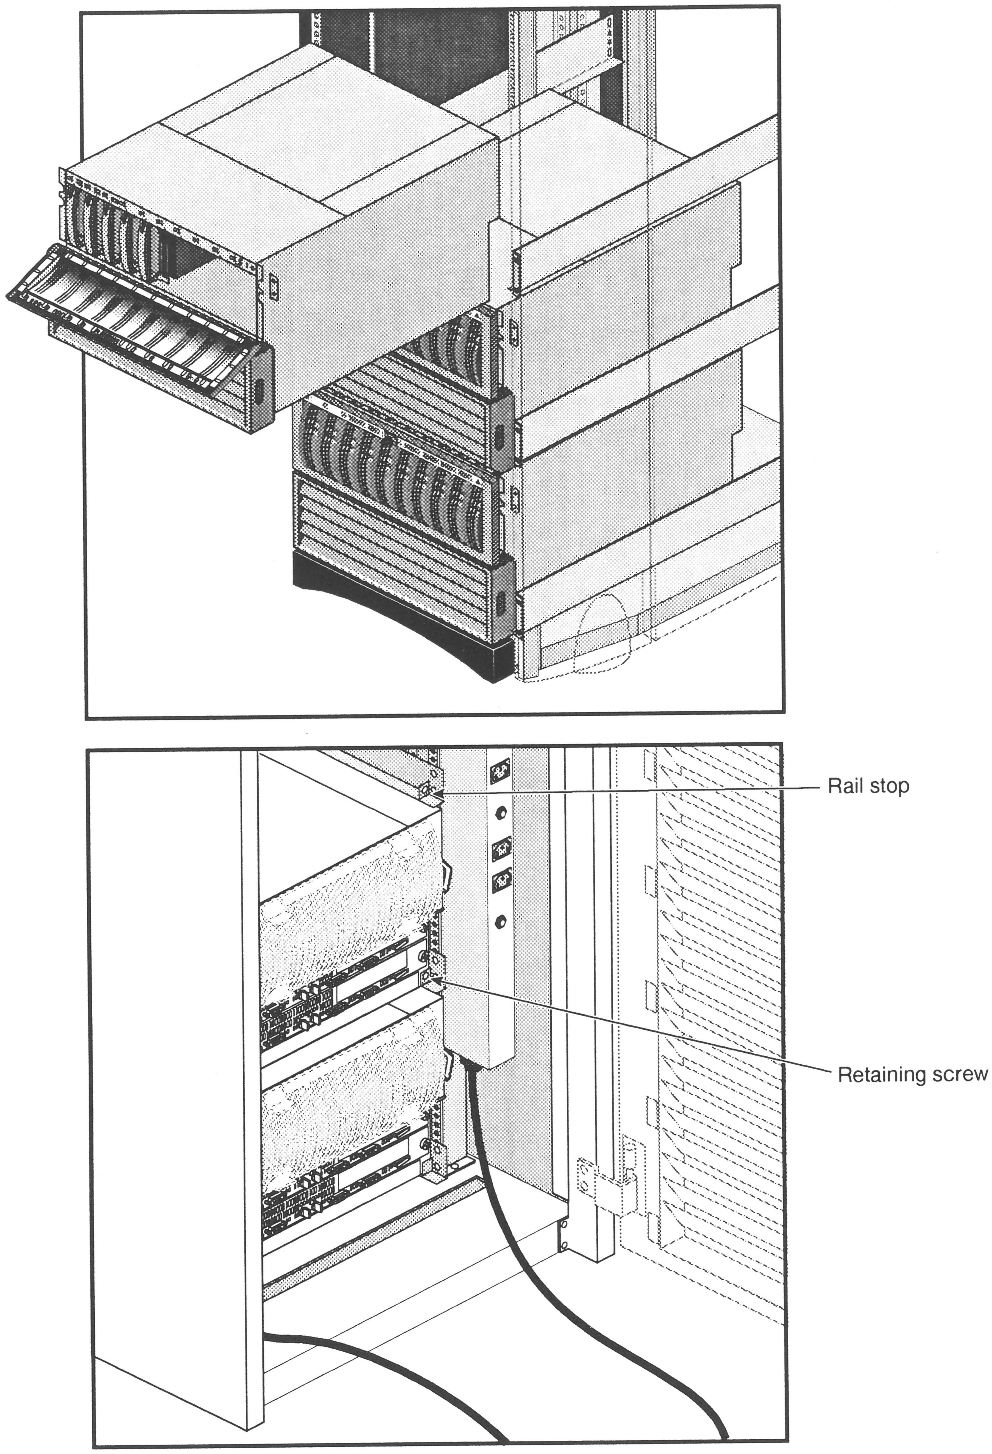

Figure 2-10 Fibre Channel RAID Enclosure Installation Example

2.3.7 Setting the Enclosure Address on the FibreVault

The drives read their FC-AL physical address only at power-on or when the drive is reset. As a result, you must set the enclosure address when you install the FibreVault and power is off. Never change the address while power is on.

Note: Each enclosure on an FC-AL loop *must* have a different ID number.

Set the enclosure address for each unit by using a pencil or ballpoint pen to move the down or up switch. Figure 2-11 shows the address switch.

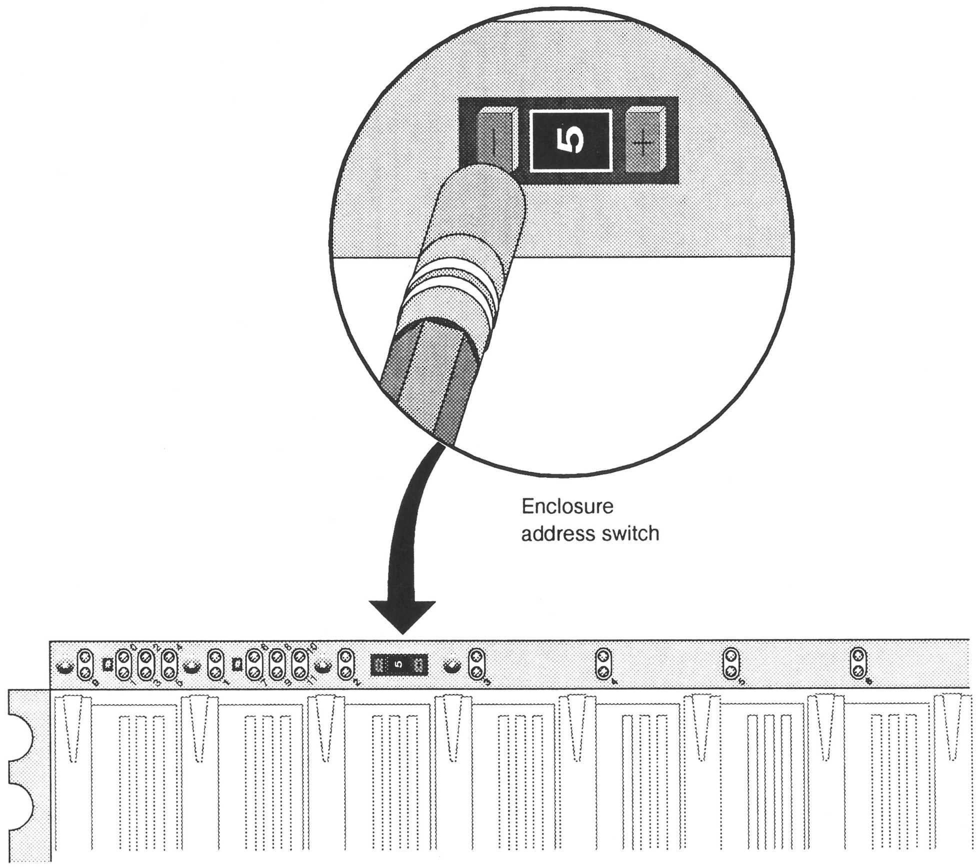

Figure 2-11 Setting the FibreVault Enclosure ID

The disk slots in a FibreVault enclosure are numbered 0-9 (from left to right). Disk modules are numbered based on the ID of the enclosure. When you are finished setting the address, push up the FibreVault's front door until it latches into place. Table 2-2 shows the maximum disk numbering.

Table 2-2 FibreVault Enclosure and Disk Slot Numbering

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

2.3.8 Setting the Enclosure Address on a Fibre Channel RAID Enclosure

Each system processor (SP) must have a unique fibre channel front-end address. Valid address ranges 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 shown in Figure 2-12. Read the scsiha reference (man) page for additional information on this topic.

The Fibre Channel RAID enclosure address is determined and set on the SP board (see Figure 2-12 and Figure 2-13).

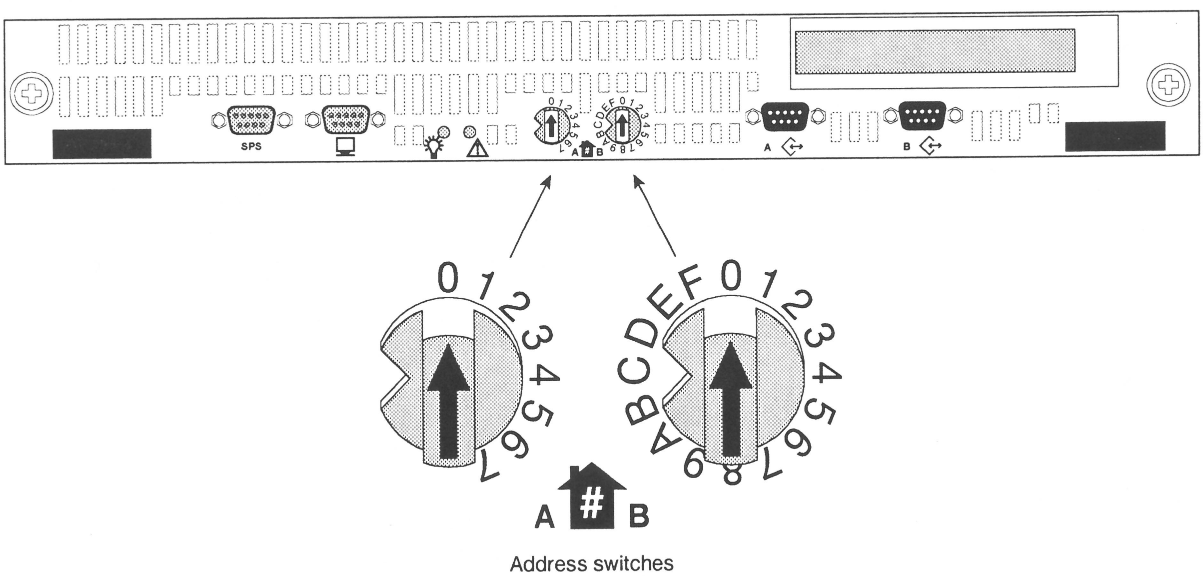

Figure 2-12 Fibre Channel RAID Enclosure Channel Address Switches

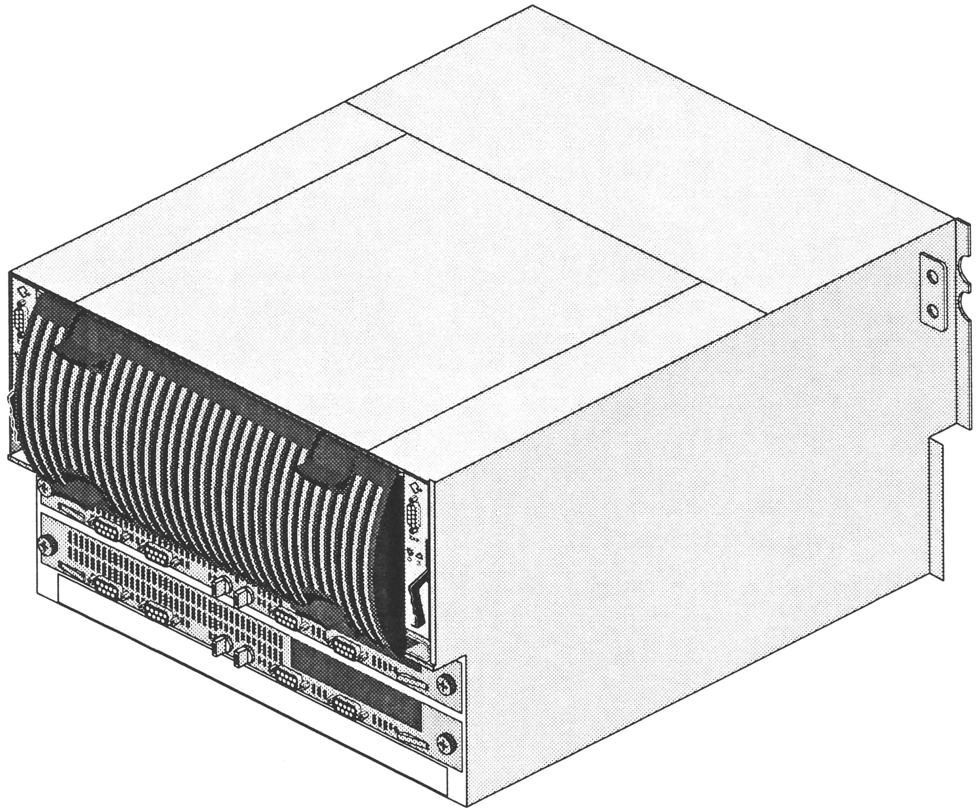

Figure 2-13 Fibre Channel RAID Enclosure Rear View

Caution: Silicon Graphics fibre enclosure addresses are limited to the following ranges:

- 0 through 109 decimal
- 0 through 6D hex

The disk slots in a FibreVault or Fibre Channel RAID enclosure are numbered 0-9. Disk modules are numbered based on the ID of the FibreVault or Fibre Channel RAID enclosure. Table 2-2 shows the maximum disk numbering.

2.3.9 Connecting the Enclosure Power Cord(s)

As a general rule, connect the power cord(s) after the enclosure(s) has been installed in the Fibre Channel Rack.

To attach the enclosure power cord, follow these steps:

1. At the back of the enclosure, remove the fan assembly; see the example in Figure 2-14. Squeeze the fan assembly latches toward each other to release the unit. Go on to step 4 if the fan module is already out.
2. If the assembly does not pull out easily, 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.

3. 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.

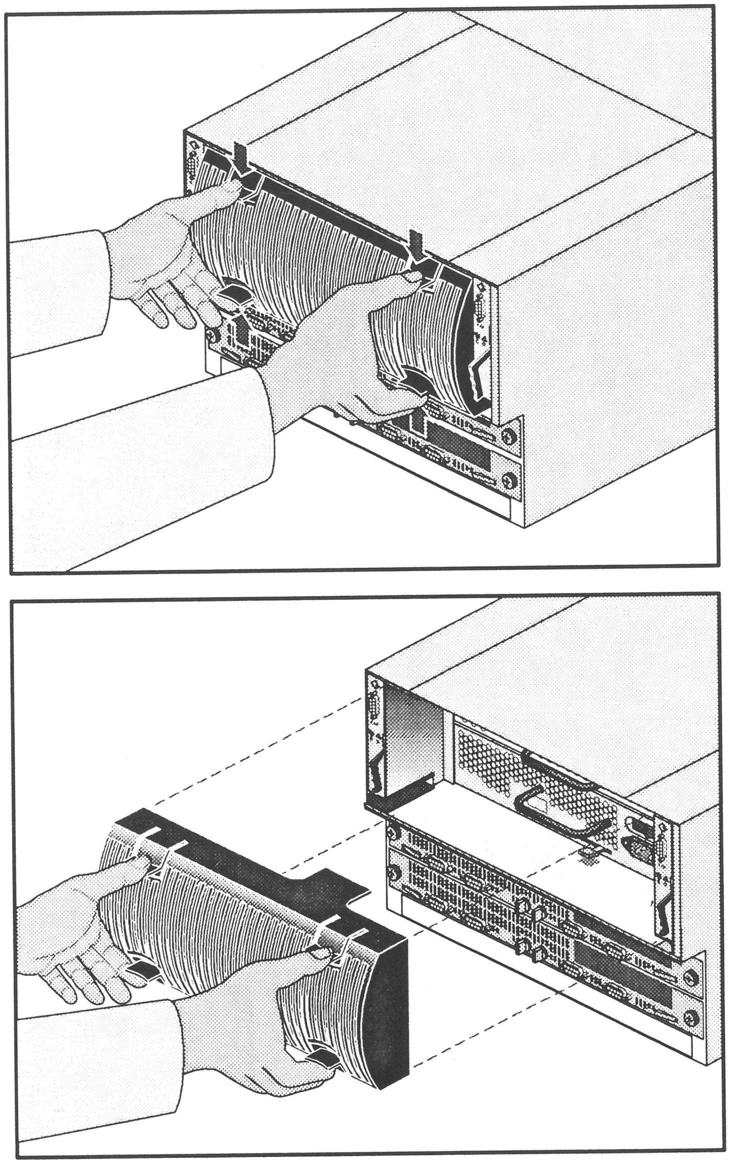

Figure 2-14 Removing the Fan Assembly

4. At the back of the rack, plug the power supply AC line cord supplied with the enclosure into the power supply outlet and route the cord along the power supply to the side of the enclosure chassis, avoiding the space occupied by the fan assembly. See Figure 2-15 and Figure 2-16.

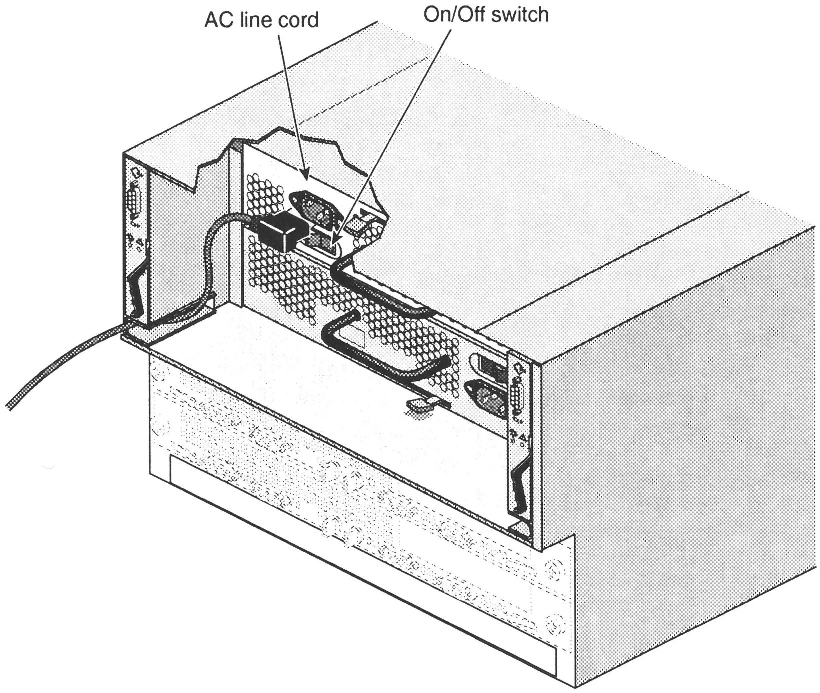

Figure 2-15 Attaching the AC Line Cord: Top Power Supply

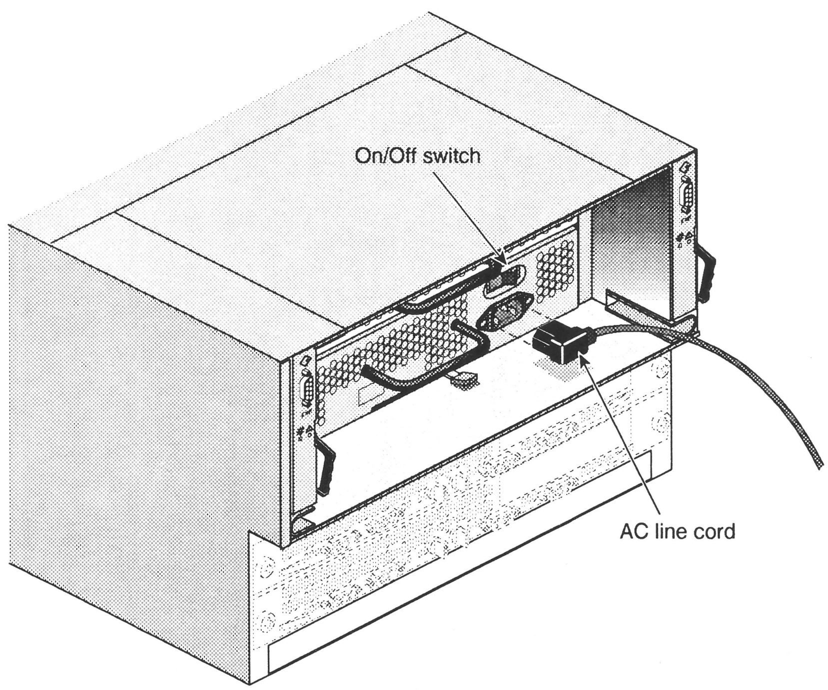

Figure 2-16 Attaching the AC Line Cord: Bottom Power Supply

5. Bend the cord into a U shape and slide it into the channel so that the end of the cable comes out of the channel at the back of the chassis.
6. Plug the end of the cord into one of the rack's outlets, as shown in Figure 2-17.
7. Set the power switch on each power supply to the on position.

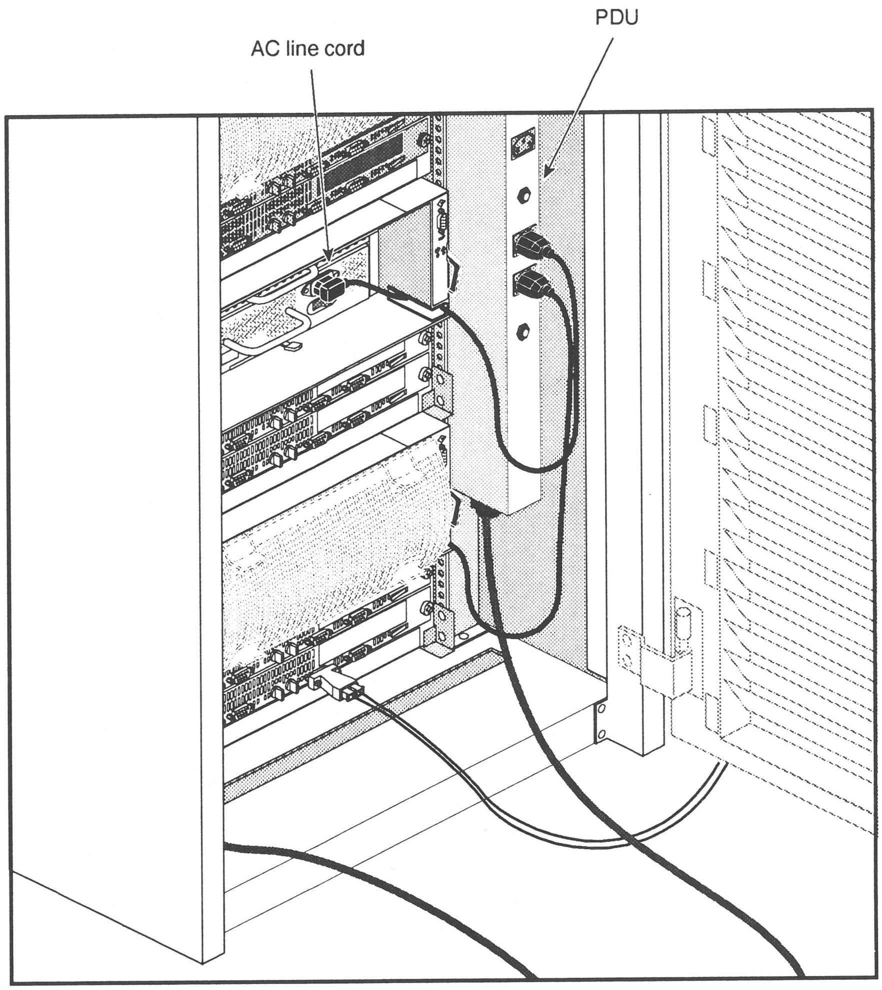

Figure 2-17 Routing the Power Cable to the PDU

8. Reinstall the fan assembly.

   Although the fan assembly attaches to the enclosure with either side up, for consistency among multiple FibreVault or Fibre Channel RAID enclosures, install it with the amber status light in the upper right corner.
9. 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.
10. Repeat the process for all new enclosures you have installed.

2.3.10 Cabling Enclosures in the Fibre Channel Rack

When you connect multiple enclosures on one FC-AL signal, retiming is required to maintain signal integrity. The drive modules in each enclosure must provide this timing.

Signal "retiming" is a functional feature of the fibre channel standard that guarantees signal integrity for the arbitrated loop topology. Each active device within a loop receives a bused signal from the device ahead of it on the FC-AL loop. The signal is retimed and retransmitted so that each device on the loop receives signals within the defined fibre channel minimum signal margins.

To cable the enclosures, follow these steps:

1. Connect one end of the enclosure's 0.3-meter cable to the expansion (**EXP**) port of the enclosure's LCC, as shown in Figure 2-18. If cabling is on one side of the rack only, use the LCC on that side.

   Note: To connect enclosures, use only the 0.3-meter (11.8-inch) copper cables supplied for this purpose. Daisy-chaining to other enclosures in other racks is not supported.

2. Connect the other end of the enclosure cable to the primary (**PRI**) port of the next enclosure.
3. If cabling is on both sides of the rack (dual link control cards), connect the optional second enclosure cable to the other LCC as in steps 1 and 2.

   Note: Make sure you have no cross-cabling between LCC As and LCC Bs.

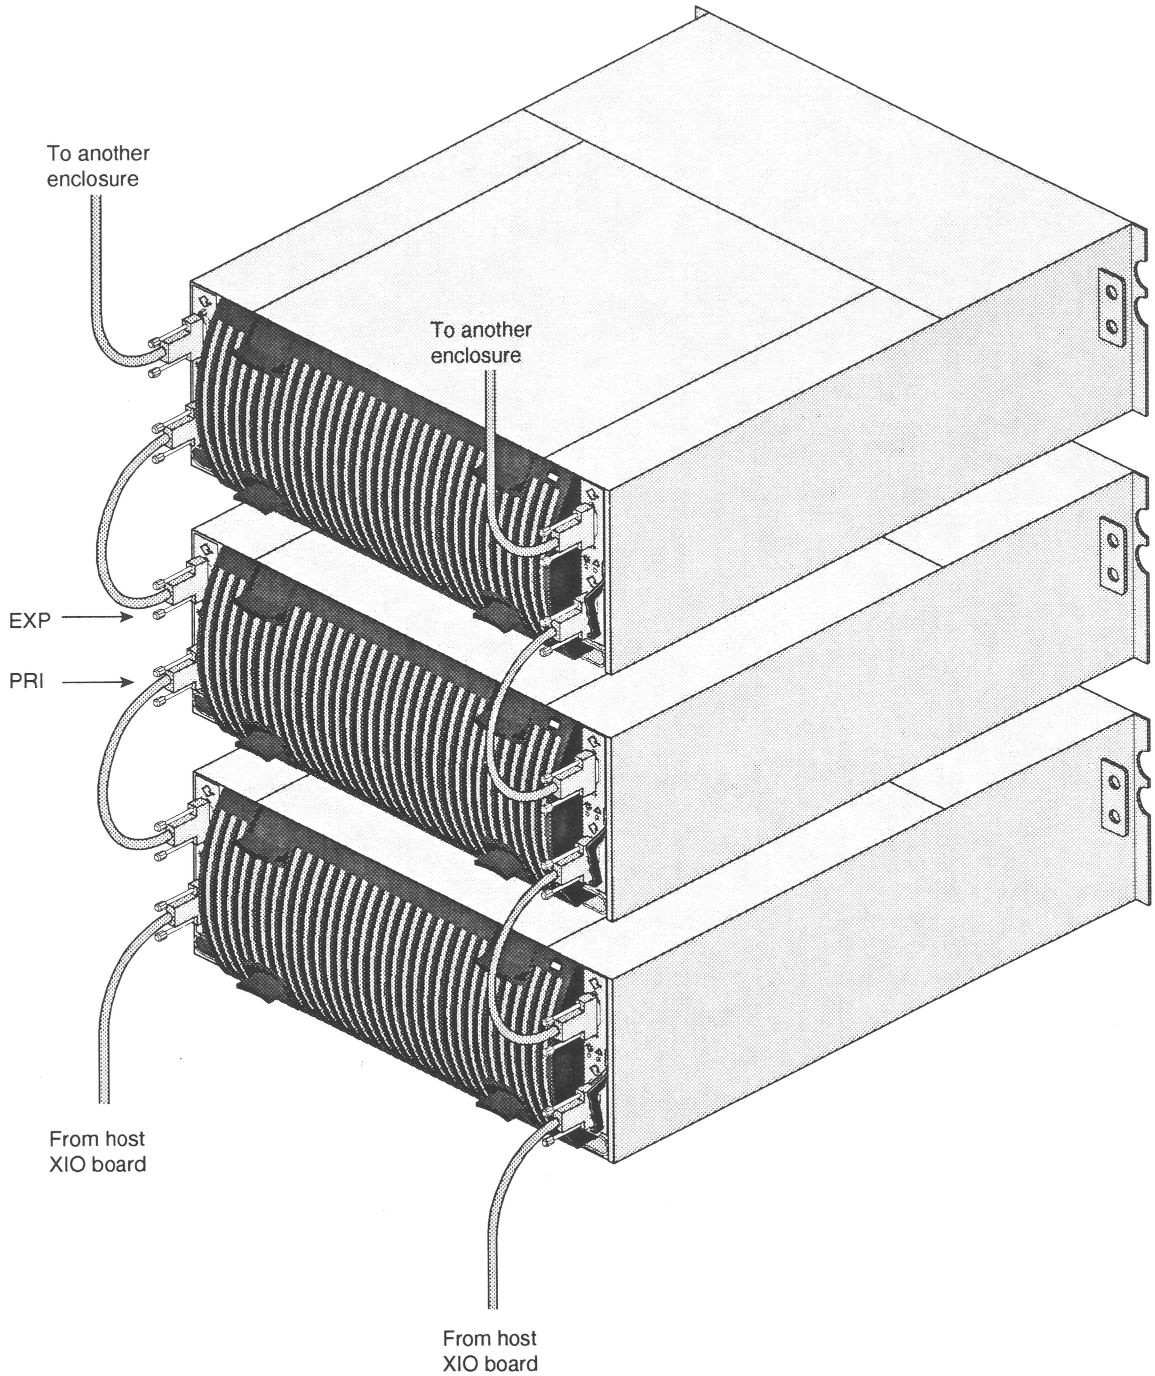

Figure 2-18 Cabling Separate LCC Loops

2.4 Configuring a Fibre Channel Enclosure

Each enclosure on an FC-AL loop must have a unique enclosure address (EA) that identifies the enclosure and determines disk module addresses.

The supported EA range is from 0 through 10. You set the FibreVault or Fibre Channel RAID expansion address with the EA switch, which has one push button for incrementing the address and another for decrementing it. Access these buttons by opening the enclosure's front door. Use a tool, such as a pen or small screwdriver to select the enclosure address. The Fibre Channel RAID enclosure address is set on the SP board as described in "Setting the Enclosure Address on a Fibre Channel RAID Enclosure" on page 2-19 in this chapter.

Each disk module on an FC-AL must have a unique FC-AL physical address. This address is derived from the EA as follows:

- FC-AL physical address = (10 multiplied by EA) + SA
- SA is the slot assignment number of the slot (0-9) containing the disk module

Guidelines for enclosure and drive module combinations on an FC-AL are as follows:

- Configure all fibre channel loops in a linear chain.
- Populate the disk slots starting with the lowest address (0) and adding disks until you reach the highest applicable address (left to right).
- All enclosures attached to an FC-AL should have at least one disk module installed.
- Never connect a FibreVault or Fibre Channel RAID enclosure to an FC-AL if it has no drives installed.
- The non-RAID FibreVault enclosure must always have a disk module installed in either slot 1 or 2. Access to enclosure-specific environmental information depends on at least one of these slots being populated with a disk.

2.5 Configuring the Storage Subsystem for High Availability

You can configure and enclosure as non-high-availability or high-availability. A non-high-availability enclosure has one power supply and one link control card (LCC) with fillers in the slots for the redundant power supply and LCC. A high-availability enclosure has two power supplies and two link control cards. The drive fan pack always provides redundant cooling, no matter which configuration you have.

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 fibre channel link control cards (LCCs) installed in the FibreVault provide an additional physical connection to the second port of each disk in the vault.

Note: Special optional IRIS FailSafe™ software support is required to take advantage of a dual port high-availability configuration. Customer-specific application software may also cover high-availability requirements. High availability is provided in a software and hardware option package; it is not currently 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 immediately replaced without shutting off the disk enclosure.

At the fullest level of fault tolerance, two hosts can be cross-connected to a storage system with redundant power supplies, cooling, dual-ported disks, and so on. As noted previously in this section, you must have optional high-availability or IRIS FailSafe software to fully use these options.

This scheme can provide protection from a failure occurring at a disk, path, power supply, fan, or host.

2.5.1 RAID Data Redundancy

RAID technology provides redundant disk resources in disk-array configurations for high availability of the storage system. A RAID-3 array maintains parity data that lets the disk group survive a disk module failure without losing data. When a disk module in the array does fail, the data is still available and the 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 Section 2.5.2.1) 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 Appendix B and Appendix C.

2.5.2 RAID Parallel Access Array

The Fibre Channel RAID release supports a 4+1 or 8+1 bound RAID-3 group. In a RAID-3 group, the hardware always reads from or writes to all its disk modules. See your system release notes for any changes in the RAID support levels. In a RAID-3 group, the hardware processes disk requests serially. In a RAID-3 group, the parity information is stored on one disk module.

2.5.2.1 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 4+1 or 8+1 RAID array installed. The capacity of a disk module that you bind as a hot spare must be at least as great as the capacity of the largest disk module it might replace.

If any disk in a RAID-3 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 becomes 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: 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.

2.5.3 Online Administrative Disk Swapping

The Silicon Graphics fibre channel implementation allows shutting down and removing individual disk modules using either the command line interface or graphical user interface. This administrative "warm swapping" of disk modules must be executed in a specific sequence to guarantee data integrity. See Appendix B or Appendix C for more information.

Caution: Pulling and replacing (hot swapping) a drive without implementing the proper administrative sequence is not supported.

2.6 Powering On the Fibre Channel Rack

The only power switches on an enclosure are those on the power supplies, which are normally covered by the drive fan pack. As a result, the enclosure is always active as long as power is available.

When AC power is initially applied to an enclosure, the disk drives power up and spin up in a specified sequence. The slot spin up delays are multiples of 12 seconds. The same delays occur when a drive is inserted while the enclosure is powered up. Spindle synchronization is supported within an enclosure, but not between enclosures.

The LCC's customer-replaceable unit (CRU) monitor resets and begins its control loop. The port-bypass circuits are put into the states dictated by their associated drives. The CRU monitor continues to run in this local mode until it receives Enclosure Services Interconnect (ESI) commands that dictate otherwise.

In local mode, the CRU monitor maintains the port-bypass circuits in the same states as the drive command signals. When a drive fault occurs, the corresponding drive status light turns on. ESI commands can take over control of the port-bypass circuits and the drive status lights.

The drives read their FC-AL physical address only at power-up or when the drive is reset. As a result, you must set the enclosure address when you install the enclosure and power is off, and not change the address while power is on.

To power on the Fibre Channel Rack, follow these steps:

1. Make sure that the power switch on each enclosure's supply is set to on.
2. Make sure all the slots in each enclosure contain either disk modules or fillers.
3. Confirm the power cord for each supply in each enclosure is plugged in to an AC power outlet on the power distribution unit (PDU). See Figure 2-19.
4. In the rack, set the PDU's circuit breaker switches to the on position. The enclosures in the cabinet power on.
5. Power on the host system.

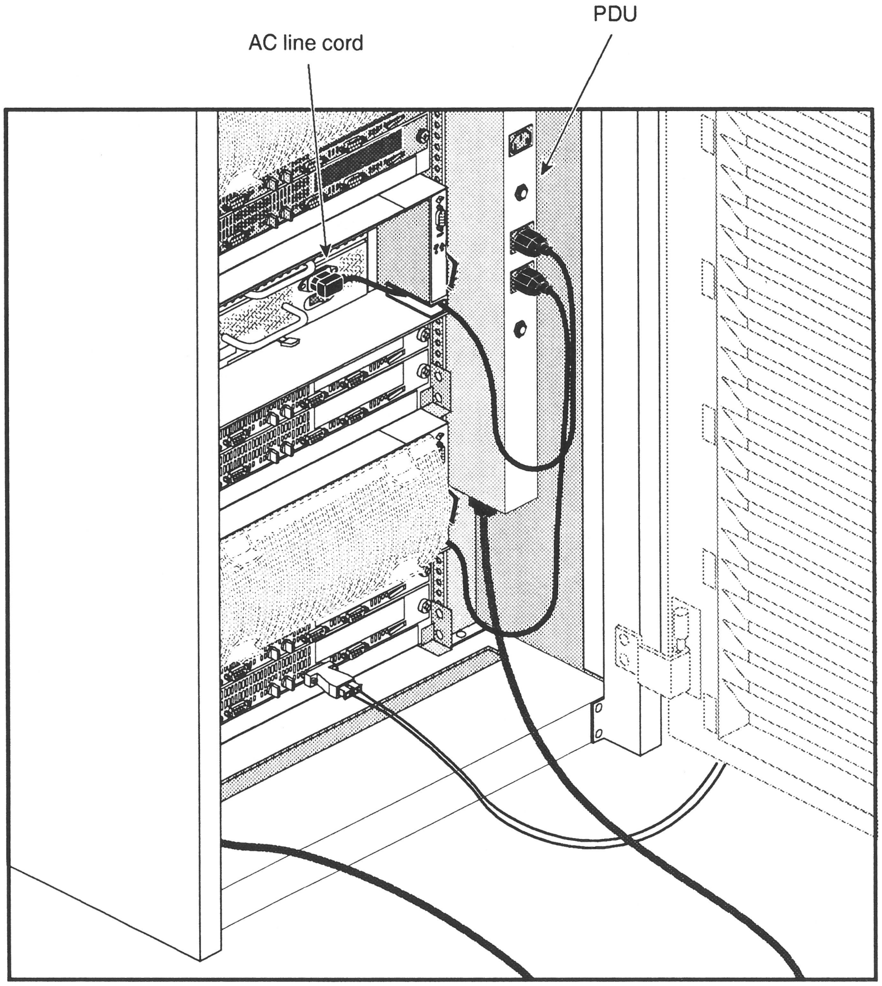

Figure 2-19 Connecting the Power Cords (Back of Fibre Channel Rack)

2.7 Replacing the Power Distribution Unit in a Fibre Channel Rack

This section explains how to swap out a failed power distribution unit from the Fibre Channel Rack. These units are attached to the inner rear wall of the rack. Follow these steps:

1. Work with the system administrator at the site to power off the rack. Unplug all power cables from the power distribution unit(s). See Figure 2-21.
2. Remove or open the rear door. (The procedure should not require removing the rear door.)

   Note: Rehanging the rear door might require a second person because of the difficulty of lining up the pinlocks.

3. As shown in Figure 2-20, remove the rear side panel from the side of the rack where the PDU is located.
   - Grasp the bottom of the side panel and pull firmly outward. The ball locks on the panel should both release.
   - Lift the panel up and away from the chassis.

Note: The panel scratches or mars easily, so set it aside carefully.

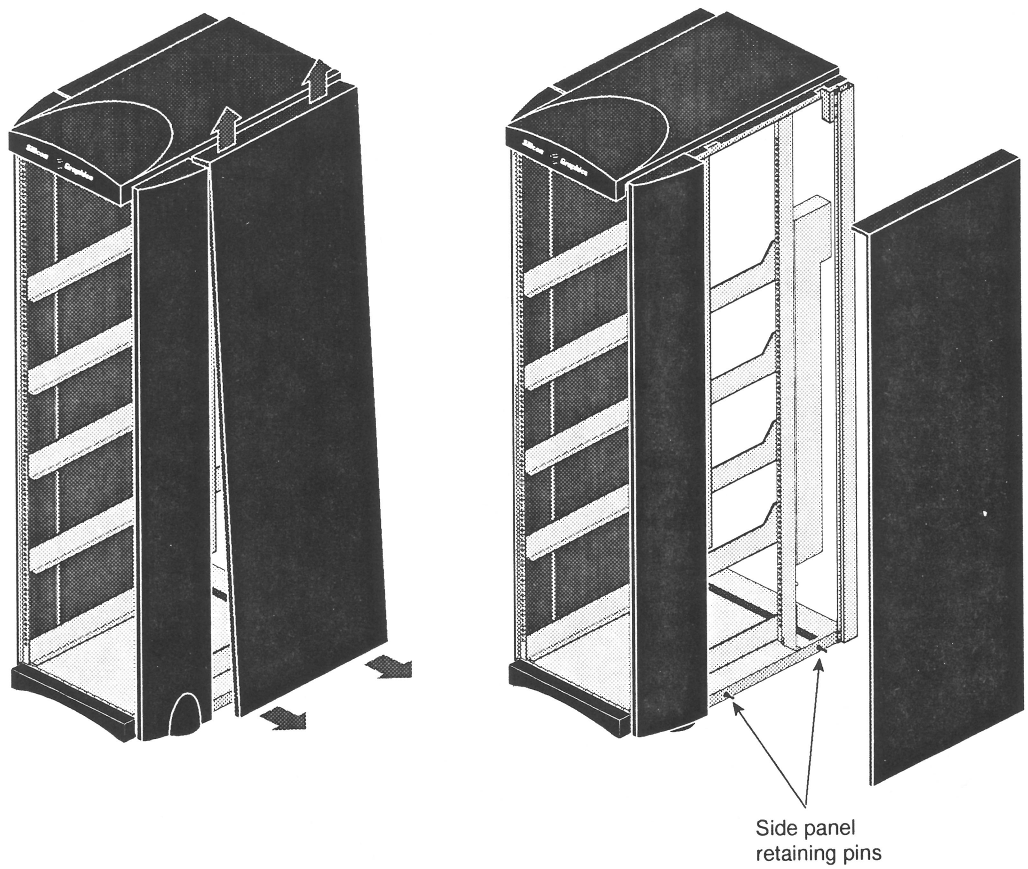

Figure 2-20 Removing the Fibre Channel Rack's Side Panel

4. Undo the screws that attach the power distribution unit to the vertical bracket and remove it, as shown in Figure 2-21.

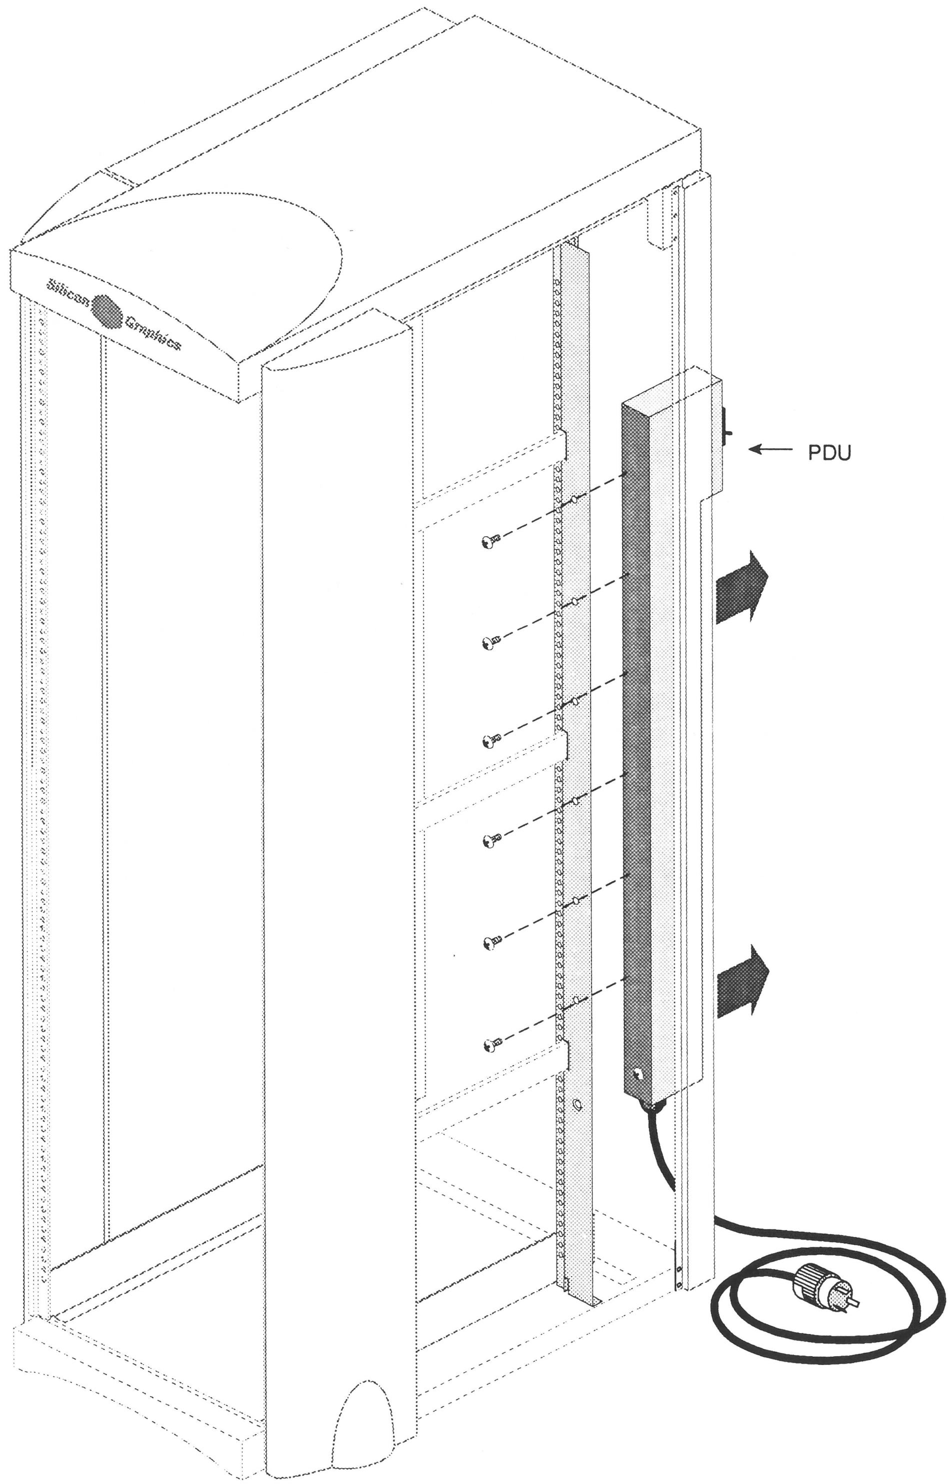

Figure 2-21 Removing the Fibre Channel Rack Power Distribution Unit

5. Attach the replacement power distribution unit.
6. Reattach the side panel, close (or reinstall) the rear door, and plug in the main power cable and all enclosure power cables. Work with the system administrator at the site to power on the system.

Chapter 3

Fibre Channel XIO Board Installation in the Host System

Customers can perform some installation and upgrade operations for the fibre channel storage system. However, the procedures in this chapter are performed *only* by Silicon Graphics trained and qualified installers:

- installing a new fibre channel XIO board
- removing and replacing an existing fibre channel XIO board

A Phillips-head and a medium or thin flat-blade screwdriver are required.

3.1 Fibre Channel XIO Board Components

As with other XIO boards, the fibre XIO boards are installed *only* in the XIO slots of Origin2000 and Onyx2 systems. Each active XIO slot provides up to 800 megabytes per second of bidirectional bandwidth (that is, 400 megabytes in each direction) through a nonblocking crossbar switch on the system's midplane. All XIO slots in a system can be active simultaneously. The main components of the XIO board are listed in Table 3-1.

Table 3-1 Main Fibre Channel XIO Board Components

Component	Explanation
9-pin ports	Two ports: channel 1 is located closest to the hook actuator lever on the board, channel 0 is always farthest away from the actuator.
Compression connector	Provides communication between the board and the system midplane.
Connector hooks	Hold compression connector securely to midplane or frontplane. There is one hook on each side of the compression connector. The hook actuator (next entry in this list) pushes or pulls the hooks into or out of the locked position.
Hook actuator	Device for moving hooks into and out of their locked position on the midplane or frontplane.
Thick side of board	The surface of the board that has the compression connector and the tallest components.
Thin side of board	The side of the board with low-profile components.

Caution: Do not wipe or touch the pads of the board's compression connector (see Figure 3-1). Take the board out of its antistatic bag and remove the compression connector cover only when you are ready to install it. When removing a board, immediately install the connector cover and place the board assembly in an antistatic container. Use all standard electrostatic discharge avoidance guidelines when handling the fibre XIO board(s).

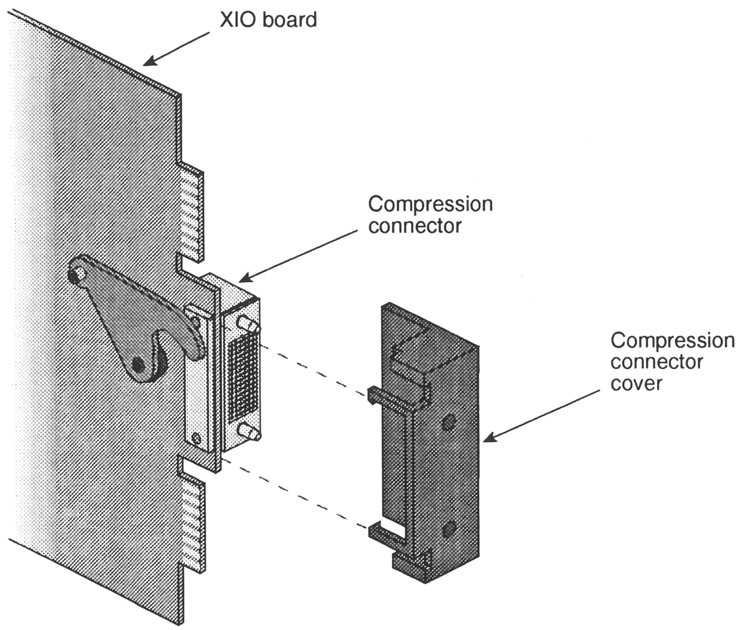

Figure 3-1 FC XIO Compression Connector and Cover

3.2 Installing or Replacing the Fibre Channel XIO Board

This section explains

- preparing the host system for the installation
- installing an FC XIO board in a host system

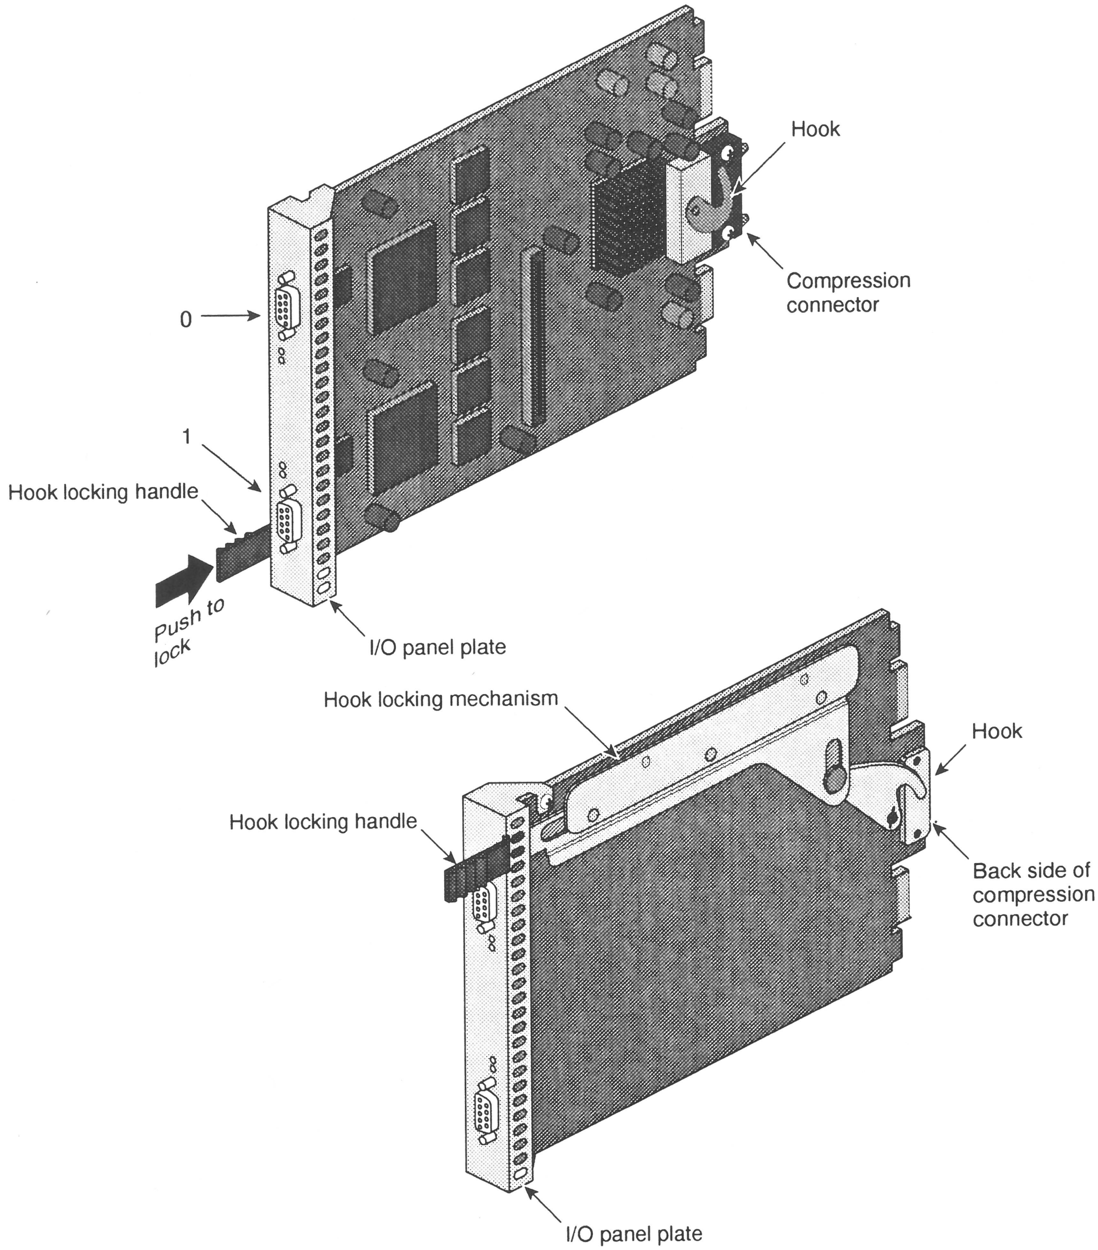

Figure 3-2 Fibre XIO Board

3.2.1 Host System Preparation

Follow the instructions in this section to shut down the host system prior to installation of the fibre XIO board. If you have already shut down the system, go to the next section and begin the board installation.

Use the following steps to prepare the host system for XIO installation:

1. Shutdown the system:

```
% su
Password: the_password
# sync
# /etc/halt
```

2. When the message appears indicating that it is safe to power off the system, turn the key switch on the System Controller to the STANDBY position.
3. At the rear of the system, flip the power switch (circuit breaker) Off (down). Do not disconnect the power plug.
4. Wait 2 full minutes (after turning off the power) to allow the system's stored electrical charge to dissipate.

Caution: Failure to properly shut down the host system may cause damage to equipment and expose the XIO board installer to risk of electric shock.

3.2.2 Chassis Grounding Issues

Grounding issues are very important in Origin family and Onyx2 systems. Each chassis must be well grounded through its power connector. All chassis connected together with XIO copper cables must share the same transformer, must be grounded through the same earthing rod, and must be on the same branch circuit.

If you have any doubts about the quality of the ground connection, it is important that you consult with a qualified electrician. Use of an optical cable between the fibre enclosure(s) and the host XIO 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 together with copper XIO cables can cause severe equipment damage and can 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.

3.2.3 Installing an FC XIO Board in the Host System

The fibre channel XIO board installs in the Silicon Graphics host system in a manner similar to other XIO boards.

The lever (handle) on the XIO board is pulled or pushed to work the compression connector's hook:

- Pushing the handle engages the hook and seats the compression connector to the midplane.
- Pulling the handle releases the hook in preparation for removing the board.

XIO boards are inserted at the rear of the chassis (module), on the right side. Where a slot is not populated, a baffle board is installed to ensure proper airflow. As a general rule, fill available odd-numbered XIO slots before filling even-numbered ones, and fill lower-numbered slots before higher-numbered ones.

To remove a blank panel and insert the new FC XIO board, follow these steps:

1. Remove the new board from its packaging and lay it on an antistatic work surface.
2. If you have not already done so, power off the system. At the back of the chassis, ground yourself and follow other ESD procedures (see the previous section).
3. Locate the panel plate containment bar, loosen each of the bar's screws and pull it towards you. You will need to slide the bar over some rivets. The bar snaps into a holding position so that it stays out of the way.
4. Remove the panel plate and XIO air baffle board from the selected slot and lay it aside.
5. Remove the protective cap from the new FC board's compression connector (if you have not already done so) as illustrated in Figure 3-1. Save this cap to cover the compression connector if you remove the board for any reason.

   Caution: Failure to install the protective cap on the compression connector when the board is removed from the host system can result in irreparable damage to the connector's pads.
6. Orient the board depending on the slot it is to occupy, as diagrammed in Figure 3-3 and Figure 3-4, slide it into the chassis until it is snug against the midplane, and press evenly on the panel plate until the board comes to a stop. If the board's panel plate is not flush with the other panel plates, continue pressing until it is.
7. Lock the board to the midplane by firmly pushing the handle of the engagement lever until it stops. Pushing in on this handle presses the compression connector's hook into its receiver on the midplane.
8. Slide the containment bar back into place so that it holds the panel plates. Tighten its screws.
9. Reverse the previous steps to remove an XIO fibre board.

Go to the next section and connect cables to the appropriate connectors. Do not power the system on until all the fibre channel components are installed and powered on.

Note: IRIX release 6.4 supports a maximum of four fibre XIO controller boards in a host system (not per module). Check the release notes for later OS versions to determine if this maximum number has increased.

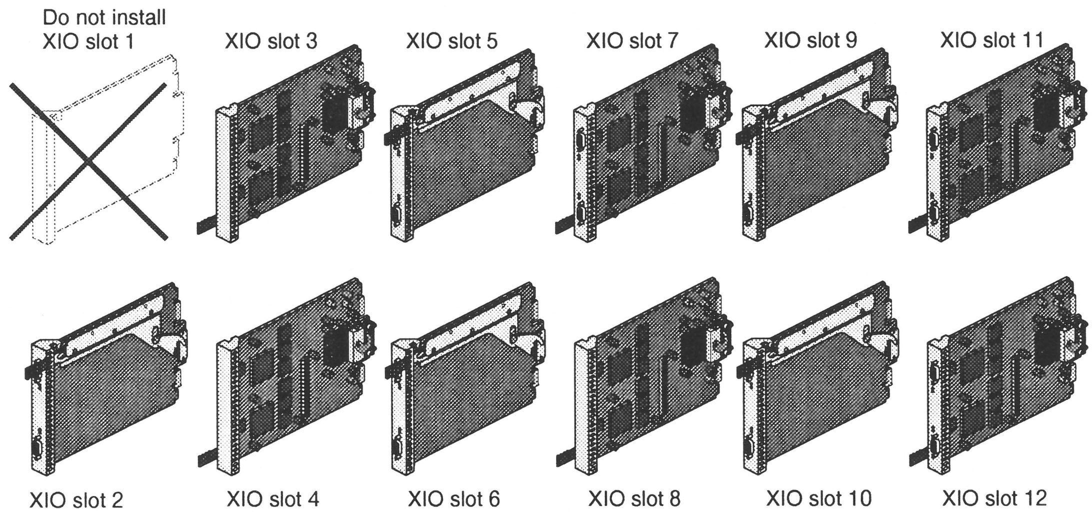

Figure 3-3 Board Orientation for Origin2000 Systems and Onyx2 Rackmount

The Onyx2 deskside system has a different board orientation scheme (see Figure 3-4).

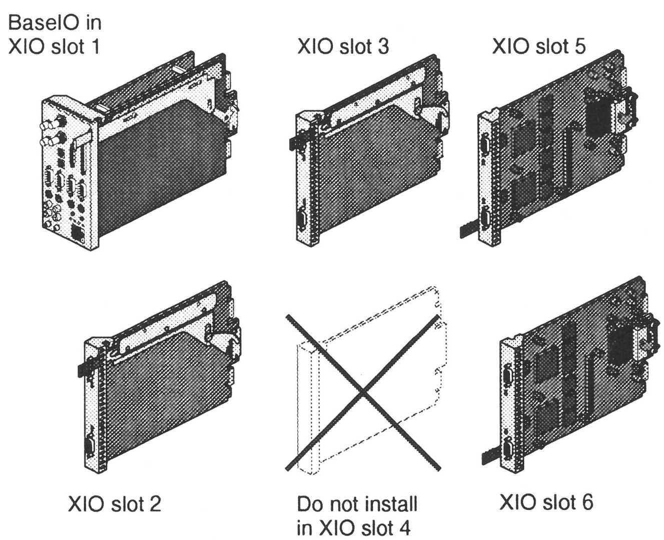

Figure 3-4 Board Orientation for Onyx2 Deskside

Note: In the Onyx2 deskside, slot 4 has no midplane connector.

3.2.4 Fibre Channel Cabling to the XIO Board

This section describes the attachment of the XIO board's external cables.

1. Locate the XIO-to-FC enclosure connection cable(s). The cable(s) may be either copper or fiber optic for this connection. Refer to Table 1-1 for a list of applicable cables.
2. Attach labels on the panel plate and cable(s).
3. If you are installing optical FC cable, place the media interface adapters (MIAs) on the applicable connectors, and remove the protective cap from the cable's connectors as applicable.

   Note: Do not touch the tips of the fiber optic cable connections.
4. If the fiber optic cable tips should become dirty, clean and dry the tip of each fiber within the cable by gently rubbing the tip with a soft, lint-free cloth that has been moistened with reagent grade isopropyl alcohol. If you do not have the proper equipment, skip this step.

   Note: Do not use prepared cleaning compounds, such as tape-head cleaner or denatured (rubbing) alcohol.
5. Attach the external cable(s) to the appropriate ports, as shown in Figure 3-5.

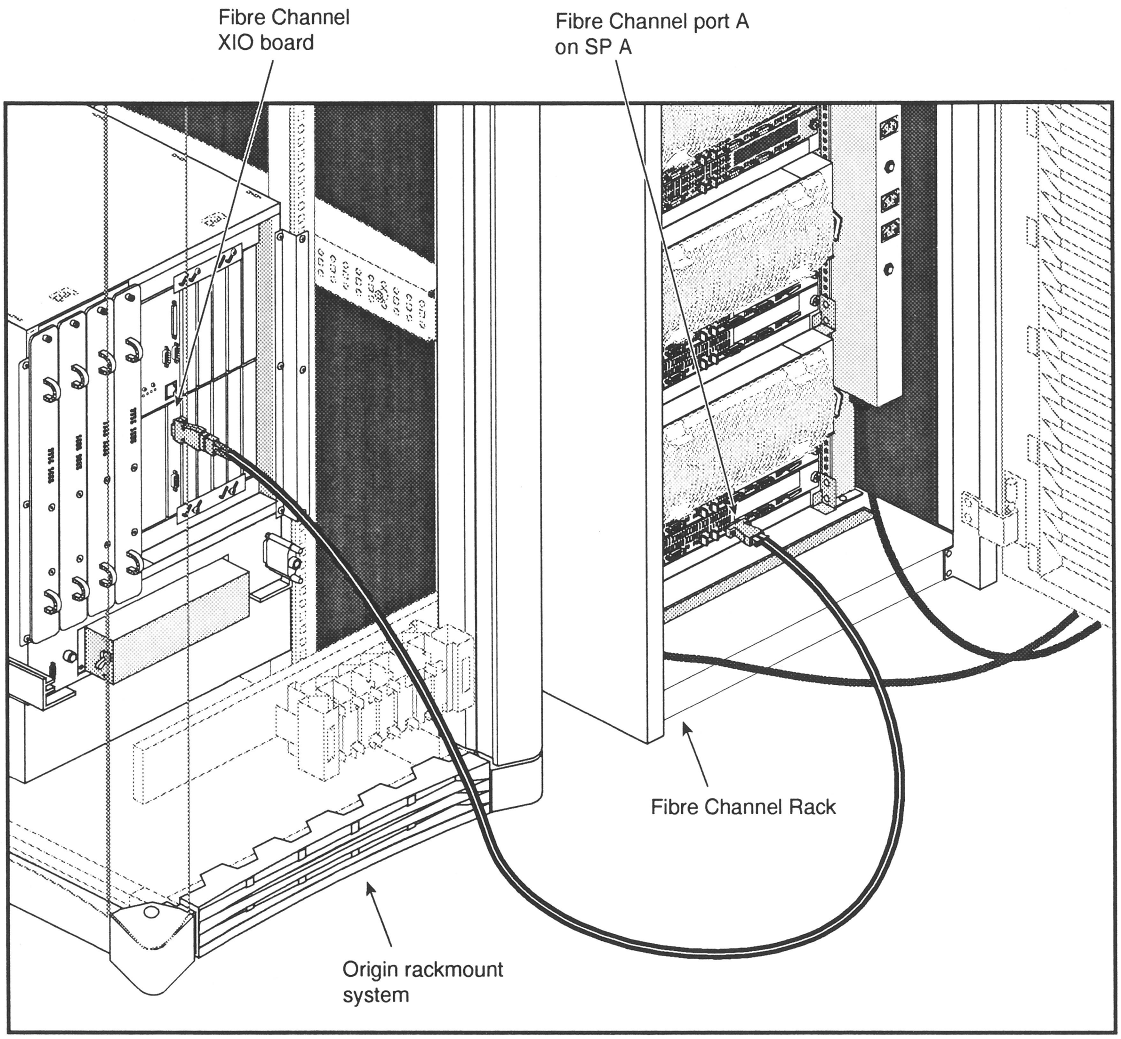

Figure 3-5 FC XIO Board Cabling Example

3.3 Restarting the Fibre Channel Host System

Restart the host only after the fibre channel enclosure(s) is installed, configured, and properly connected. When the board is installed and connected, follow these instructions to start operation:

1. Set the system circuit breaker power switch in the On position.
2. Restart the system by turning the key in the module's System Controller to On.
3. Log on after the system boots.

4. Depending on the IRIX release on the host system, you may have to install a patch or even possibly upgrade the operating system. Consult the release notes.

5. If you have not installed and configured the software, do so now by following the instructions in the release notes. The new fibre channel mass storage system does not function until the software has been properly configured.

   Note: After you finish configuring the software, reboot the system (or run the *autoconfig* command) to build a new operating system (kernel) that includes the new drivers. Then reboot the host again to start running this new operating system.

3.4 Checking the New XIO Installation

Verify that the board has been recognized by the operating system during the bootup. Type

```
% hinv
```

hinv provides a listing similar to the following. Every system will be slightly different; the point of this example is to provide general information on what to look for after the installation and successful bring-up of the fibre channel enclosures and the host.

```
# hinv
FPU: MIPS R10010 Floating Point Chip Revision: 0.0
CPU: MIPS R10000 Processor Chip Revision: 2.6
2 180 MHZ IP27 Processors
Main memory size: 128 Mbytes
Instruction cache size: 32 Kbytes
Data cache size: 32 Kbytes
Secondary unified instruction/data cache size: 1 Mbyte
Integral SCSI controller 0: Version QL1040B
  Disk drive: unit 1 on SCSI controller 0
  Disk drive: unit 2 on SCSI controller 0
Integral SCSI controller 1: Version QL1040B
Integral SCSI controller 2: Version Emerald Fibrechannel
  Disk drive: unit 40 on SCSI controller 2
  Disk drive: unit 41 on SCSI controller 2
  Disk drive: unit 42 on SCSI controller 2
  Disk drive: unit 43 on SCSI controller 2
  Disk drive: unit 44 on SCSI controller 2
  Disk drive: unit 45 on SCSI controller 2
  Disk drive: unit 46 on SCSI controller 2
  Disk drive: unit 47 on SCSI controller 2
  Disk drive: unit 48 on SCSI controller 2
  Disk drive: unit 49 on SCSI controller 2
  Disk drive: unit 50 on SCSI controller 2
  Disk drive: unit 51 on SCSI controller 2
  Disk drive: unit 52 on SCSI controller 2
  Disk drive: unit 53 on SCSI controller 2
  Disk drive: unit 54 on SCSI controller 2
  Disk drive: unit 55 on SCSI controller 2
  Disk drive: unit 56 on SCSI controller 2
  Disk drive: unit 57 on SCSI controller 2
  Disk drive: unit 58 on SCSI controller 2
  Disk drive: unit 59 on SCSI controller 2
```

```
  Disk drive: unit 110 on SCSI controller 2
  Disk drive: unit 111 on SCSI controller 2
  Disk drive: unit 112 on SCSI controller 2
  Disk drive: unit 113 on SCSI controller 2
  Disk drive: unit 114 on SCSI controller 2
  Disk drive: unit 115 on SCSI controller 2
  Disk drive: unit 116 on SCSI controller 2
  Disk drive: unit 117 on SCSI controller 2
  Disk drive: unit 118 on SCSI controller 2
  Disk drive: unit 119 on SCSI controller 2
Integral SCSI controller 3: Version Emerald Fibrechannel
  Disk drive: unit 0, lun 7 on SCSI controller 3
  Disk drive: unit 1 on SCSI controller 3
  Disk drive: unit 2 on SCSI controller 3
  Disk drive: unit 3 on SCSI controller 3
  Disk drive: unit 4 on SCSI controller 3
  Disk drive: unit 5 on SCSI controller 3
IOC3 serial port: tty1
IOC3 serial port: tty2
Integral Fast Ethernet: ef0, version 1
IOC3 external interrupts: 1
```

The first two SCSI controllers in the *hinv* listing are the native controllers that come with every Origin2000 or Onyx2 system. The next two are fibre controllers. The first controls two enclosures with a total of twenty non-RAID fibre drives. The second controls a five-disk RAID LUN that also has a hot spare available.

To confirm or identify the host XIO slot the XIO fibre controller is installed in, enter the following:

```
# find /hw/module -name fibre_channel
```

This provides information similar to the following:

```
/hw/module/1/slot/io6/fibre_channel
```

This should correctly identify the chassis and XIO slot into which you installed the board. Once you have this information you can use the **-l** option for more detail about the option; for example:

```
# ls -l /hw/module/1/slot/io6/fibre_channel
```

The following is a possible output, showing an FC XIO option board entry.

```
lrw-------    1 root     sys          26 May  7 20:40 .master ->
/hw/module/1/slot/n1/node
drwxr-xr-x    2 root     sys           0 May  7 20:40 pci
```

Use the hardware graph listing to determine if the system sees the fibre drives after the board is cabled. For more information, see the *Origin2000 and Onyx2 Deskside and Rackmount Installation Instructions*.

Chapter 4

Deskside FibreVault Disk Tower

This chapter provides general information about the deskside fibre tower. Note that the tower covered in this chapter is a non-RAID (also called JBOD) configuration.

The non-RAID FibreVaults come in deskside (tower) and rackmount versions.

Figure 4-1 shows the deskside FibreVault 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 lock up to free the system to roll.

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

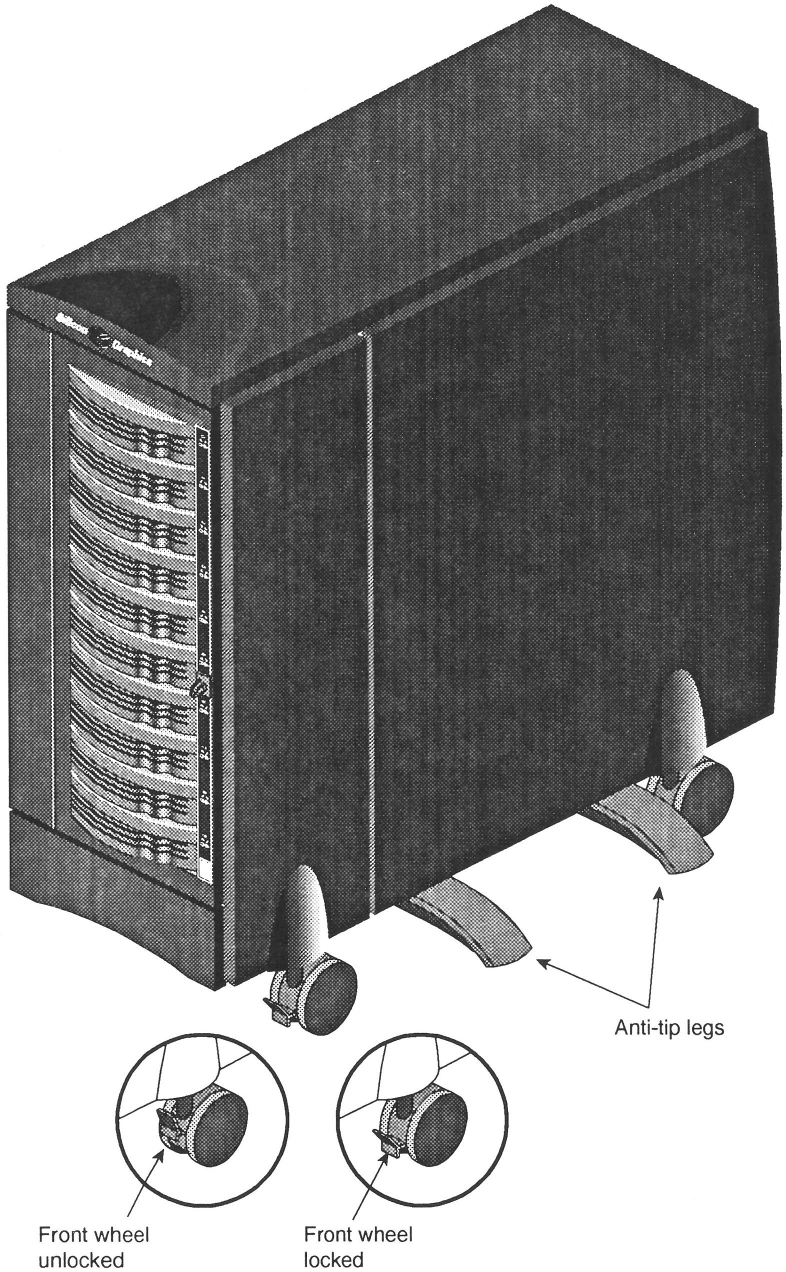

Figure 4-1 FibreVault Deskside (Tower) Unit

4.1 Unpacking the Deskside FibreVault System

The FibreVault deskside tower comes in a pallet-mounted carton. Move the entire package to the assembly area if possible. Use the following guidelines to unpack the unit.

1. Pull the top cover off the pallet and remove the plastic wrap.
2. Unbolt the bracket retainers that hold the unit to the pallet.
3. Roll the deskside tower off the pallet to the designated installation site. Note that you will have to lift the edge of the tower up a bit to clear the anti-tip legs when taking the unit off the pallet.

   Caution: The tower weighs 85 lbs (36 kg) empty and up to 132 lbs (60 kg) fully loaded. Never attempt to pick up the deskside FibreVault without the help of a second person.

4.2 Deskside FibreVault Components

The deskside FibreVault (tower) system has most of the same components the rackmount FibreVault has. It is essentially a FibreVault oriented on its side with skins, wheels, and a power distribution unit attached.

The tower enclosure assembly consists of a sheet-metal housing, front panel, and one or two power distribution units.

Figure 4-2 shows the front panel control and status LEDs.

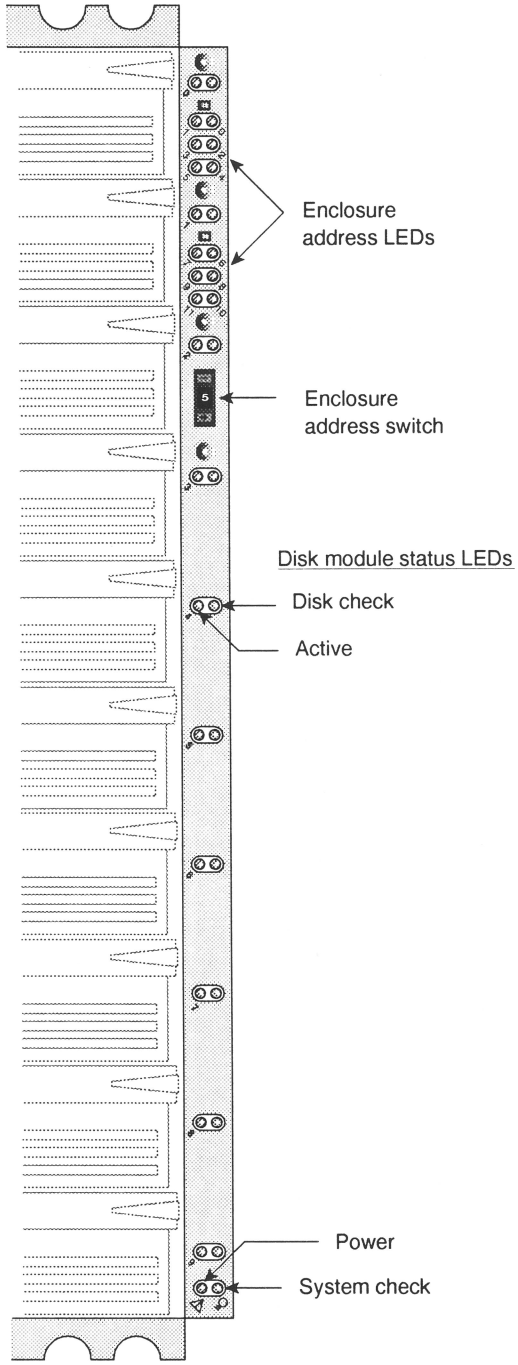

Figure 4-2 Deskside FibreVault Front Panel Features

The power distribution unit (PDU) consists of an AC inlet, a circuit breaker, and three AC outlets.

AC power enters the deskside FibreVault through the PDU when the power switch is in the On (1) position. The AC outlets on the PDU(s) supply power to the enclosure's power supplies.

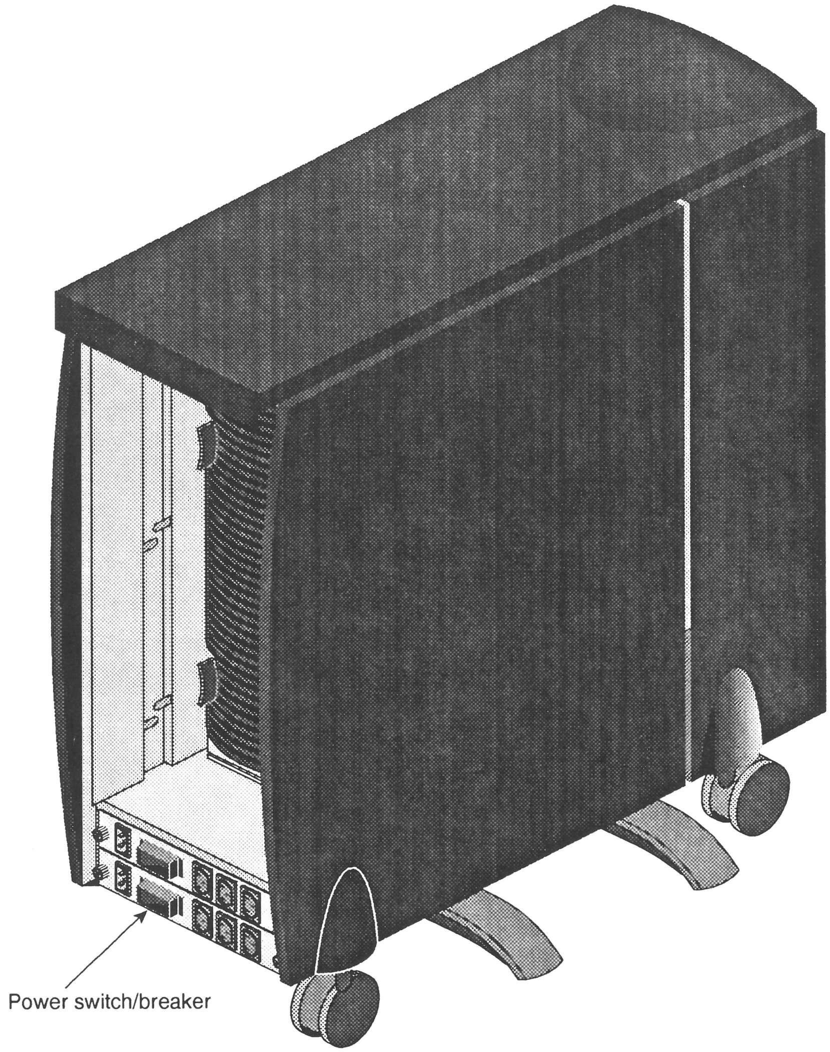

Figure 4-3 Deskside FibreVault Power Distribution Unit (PDU)

4.3 Deskside FibreVault Component Replacement

The deskside FibreVaults use power distribution units that mount on the back of the tower unit below the enclosure. To remove and replace this FRU, use the following information:

1. Take the unit offline, shut off the main circuit breaker on the PDU, and unplug all the power cords.
2. Undo the two screws holding the unit in place and pull it out of the chassis, see Figure 4-4.
3. Slide the replacement PDU into position and secure it with the two screws.
4. Reconnect all applicable power cords and turn the PDU's circuit breaker on.

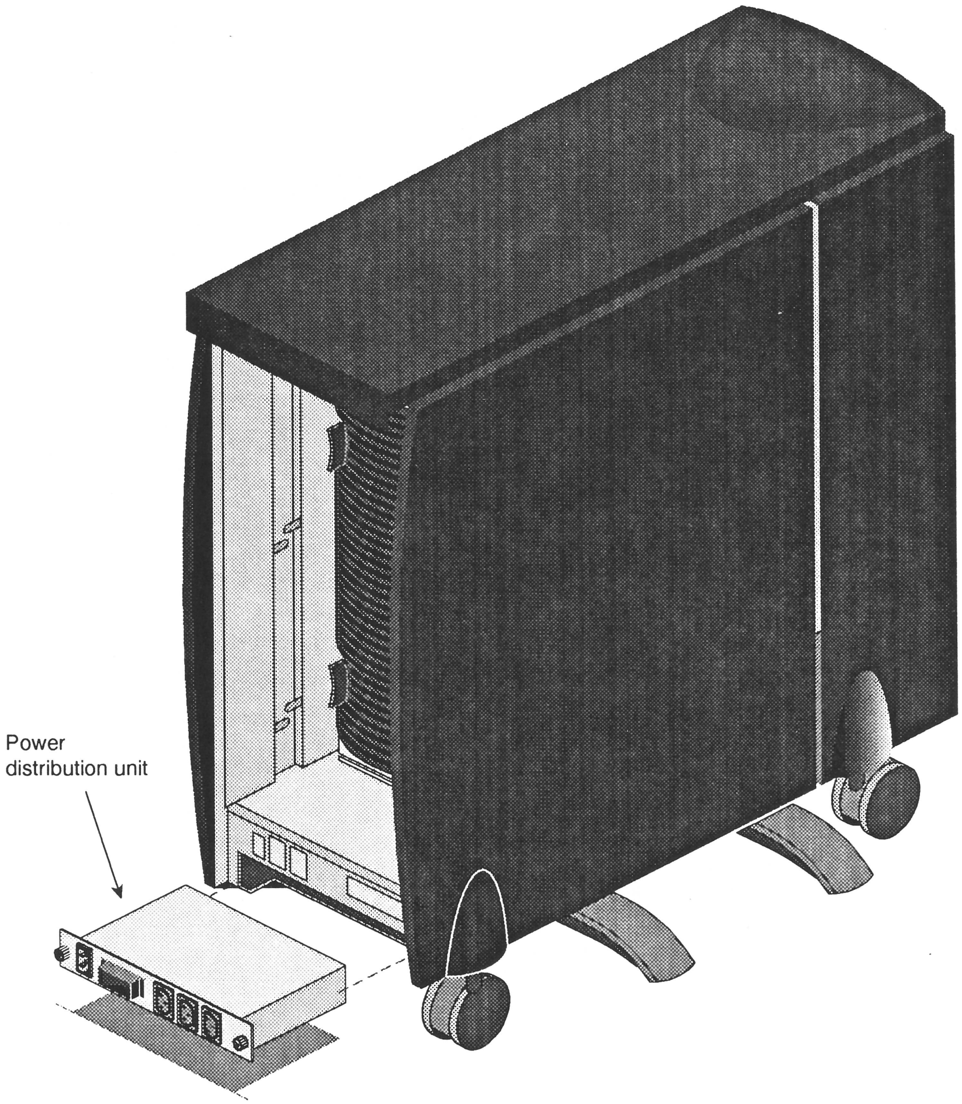

Figure 4-4 Removing the PDU From a Deskside FibreVault

Other FibreVault enclosure FRU components in the tower unit are removed and replaced in the same manner (and with the same rules and restrictions) as those in the rackmounted FibreVault. Chapter 5, "Fibre Channel Component Replacement," provides instructions for removing and replacing the FibreVault parts that are considered field-replaceable units (FRUs).

Chapter 5

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 redundant 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 either a component (drive, power supply, and so on), or a 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 (under carefully controlled procedures) remove a disk module, redundant power supply, link control card (LCC), or filler, but you should replace it as soon as possible with a replacement component or a filler.

Note: Do not remove a faulty component until you have a replacement component available.

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.

Caution: Failure to replace the fan within the two-minute period will cause an uncontrolled shutdown of the disk array and could result in possible data loss.

Table 1-1 at the beginning of Chapter 1 gives marketing codes and part numbers for fibre channel FRU components.

The FC-AL disk modules and Ultra SCSI Origin2000 or Onyx2 chassis disk modules are not interchangeable and cannot be substituted for each other. Fibre channel RAID and non-RAID disks also are not interchangeable and cannot be substituted for each other.

This chapter discusses

- "Adding or Replacing a Disk Module,"
- "Setting an Enclosure Address,"
- "Replacing a Fan Assembly,"
- "Replacing or Adding a Power Supply in a FibreVault and Fibre Channel RAID Enclosure,"
- "Replacing or Adding a Link Control Card (LCC)."

5.1 Adding or Replacing a Disk Module

When adding or replacing a disk module, use only the proper 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:

- P/N 9470138 for (8.8 GB) RAID drives (520-byte sectors)
- P/N 9470140 for (9.1 GB) non-RAID drives (512-byte sectors)

The two types of drives are not compatible and may not be mixed in the same enclosure.

Caution: The disk modules have no ESD shielding; do not stack them on top of each other prior to installation. Use standard ESD precautions as listed in the instructions and always leave disks in an antistatic container until ready for installation.

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

1. Have ready the proper replacement disk module.
2. Use one of the following methods to take the disk off line and prepare for installation:
   - Tell the host that you are removing a (non-RAID) disk module and take it off line using the `fccli REMOVE` option. Appendix B has details for using the line commands.
   - Use the Graphical User Interface (described in Appendix C) to take a RAID disk module offline. Appendix B has RAID CLI information.
   - Shut down the host and fibre channel enclosure systems.
3. 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. Remove the key. Push the door's latch, and pull the door down until it is fully open.

   If you pull the door down with enough force that it becomes unhinged, snap it back on the hinges.
4. 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 enclosure is shut down).
5. 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.

   Note: Do not remove a disk module or filler unless you have its replacement ready to install.

   Grasp the disk module's handle so that your thumb is on its latch. Push the latch and slowly pull the module from the slot, supporting it with your other hand. See Figure 5-1. The disk module's active and check LEDs turn off.

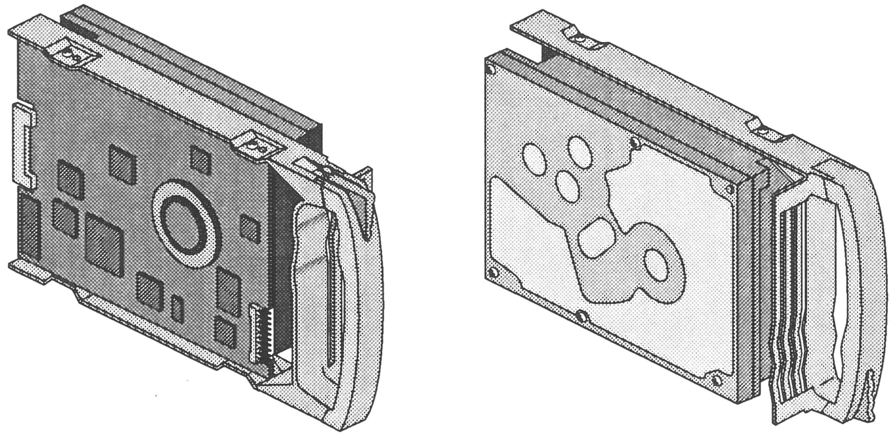

Figure 5-1 Disk Module

6. Grasp the handle of the new disk module and align the module with the guides in the empty slot.
7. 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).
8. Remove and store the ESD wristband.
9. When you have replaced or added the disk module, use one of the following procedures to bring the disk online.
   - Power the host and FC enclosure back on if they were powered off during the installation procedure.
   - Take the disk targeted for replacement offline with the Graphical User Interface (see Appendix C).
   - Use the `fccli INSERT` option in the command-line interface to inform the host (that owns the loop) of its presence; see Appendix B. The fault LED remains lit until the drive is online.
10. Push the fibre enclosure's front door up until it latches into place.

5.1.1 Setting an Enclosure Address

The drives read their FC-AL physical address only at power-on, insertion into a powered enclosure, or when the drive is reset. Each enclosure address should be set when first installed in the Fibre Channel Rack. Reset enclosure addresses only 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. FibreVault enclosures can be numbered greater than 10, but you should use only the numbers 0-10. Figure 5-2 shows the address switch.

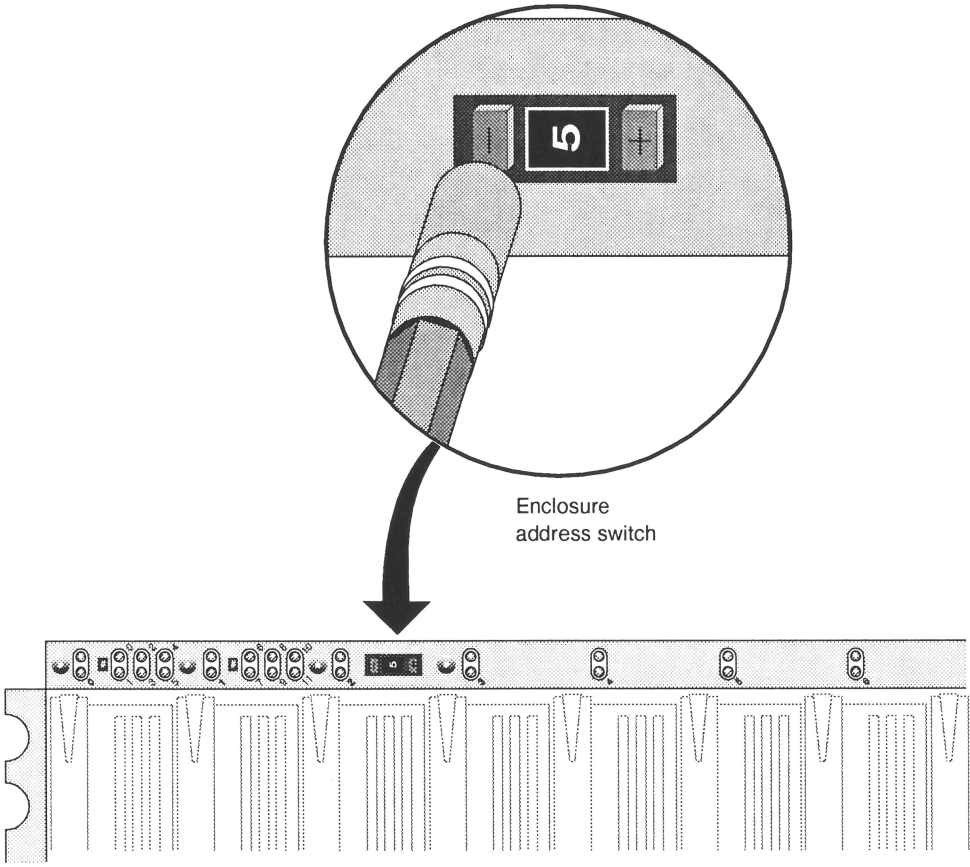

Figure 5-2 Setting the FibreVault ID

Valid address ranges 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, see Figure 5-3.

The Fibre Channel RAID enclosure address is determined and set on the system processor (SP) board, see Figure 5-4. Each SP must have a unique fibre channel front-end address.

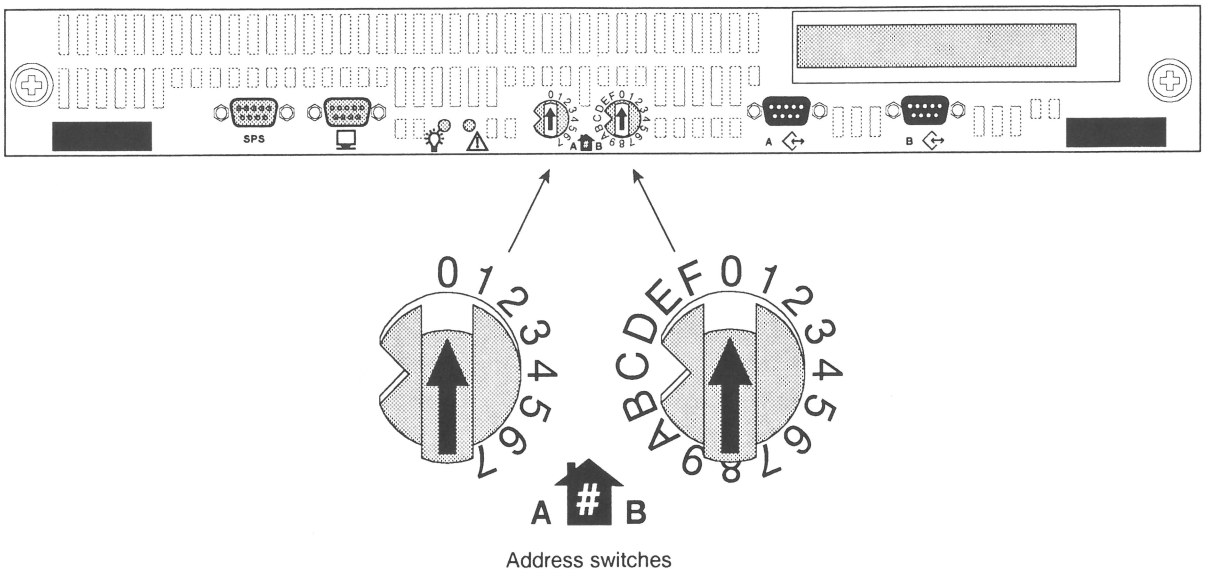

Figure 5-3 Fibre Channel RAID Enclosure Channel Address Switches

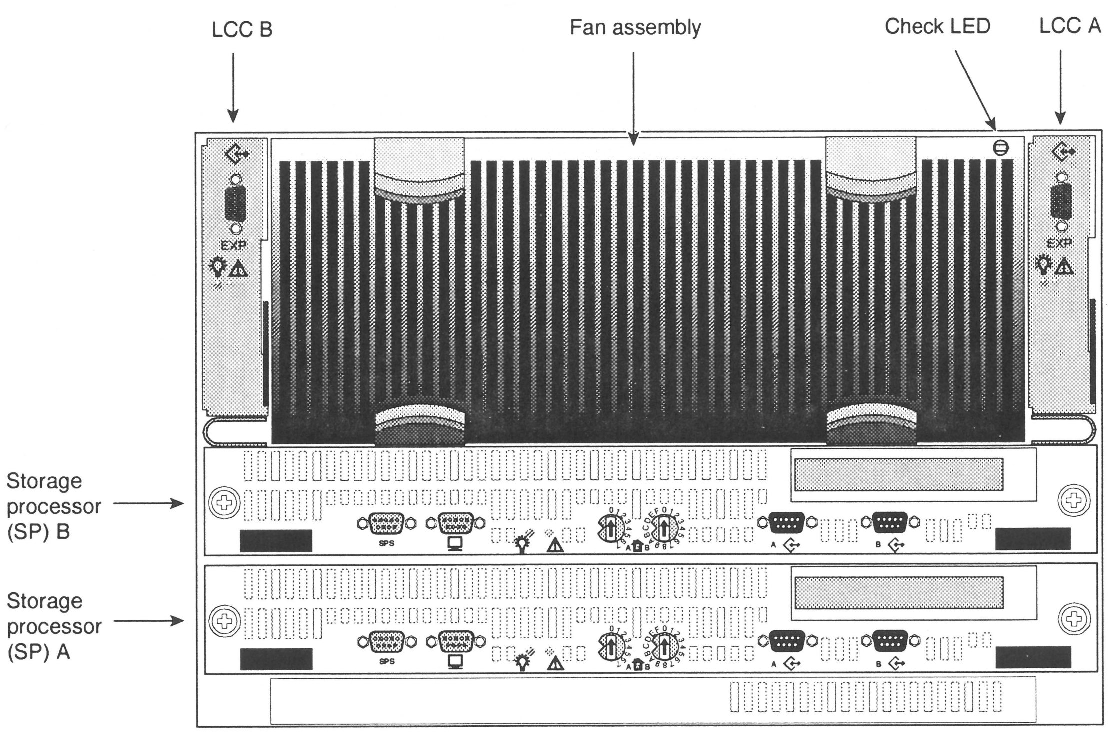

Figure 5-4 Fibre Channel RAID Enclosure Rear View

The disk slots in a FibreVault or Fibre Channel RAID enclosure are numbered 0-9. Disk modules are numbered based on the ID of the FibreVault or Fibre Channel RAID enclosure. Table 5-1 shows the maximum disk numbering.

Table 5-1 Origin 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

When you are finished setting the address, push the enclosure's front door closed until it latches into place. If desired, lock the enclosure's front door.

5.2 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. Follow these steps:

Caution: Do not remove a fan assembly unless you have its replacement ready to install.

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 5-5.

   Caution: Do not leave the fan assembly out of the fibre enclosure for more than two minutes if the enclosure 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 the power supply(ies) flash.

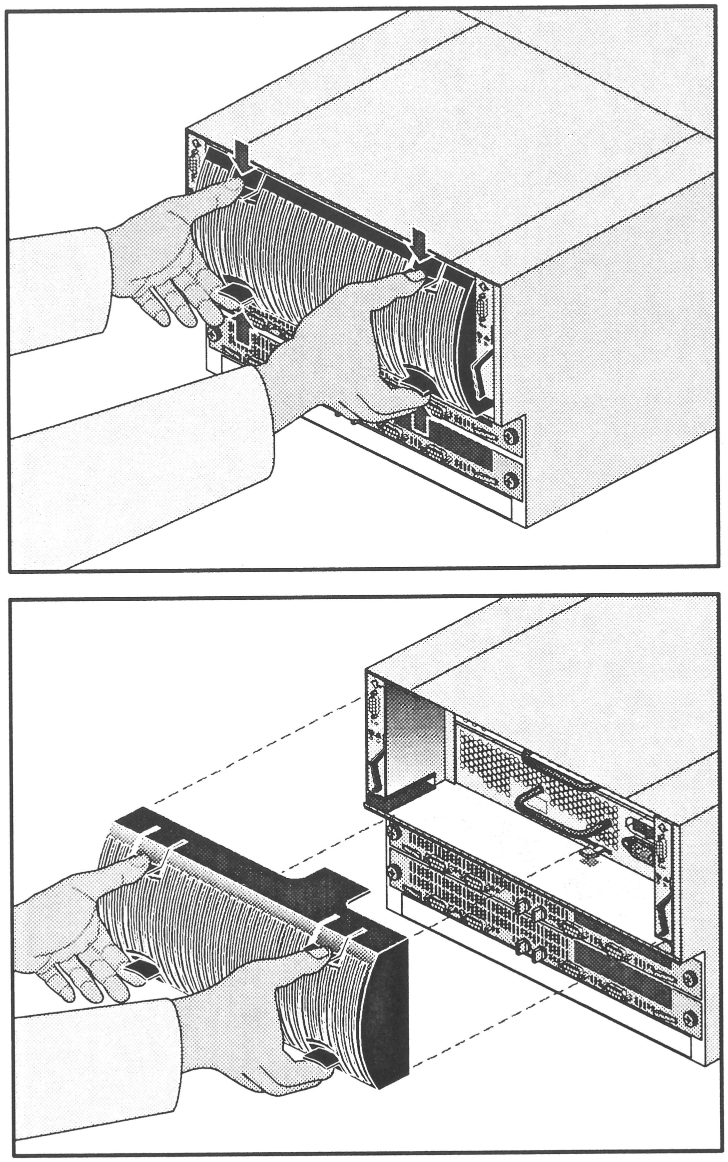

Figure 5-5 Removing the Fan Assembly

5. Orient the new fan assembly. Although the fan assembly attaches to the enclosure with either side up, for consistency among multiple FibreVault or Fibre Channel RAID enclosures, install it with the status light in the upper right corner as shown in Figure 5-6.

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.

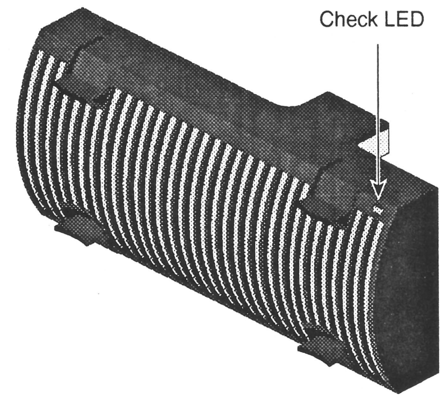

Figure 5-6 Fan Assembly Status LED Orientation

5.3 Replacing or Adding a Power Supply in a FibreVault and Fibre Channel RAID Enclosure

The power supply modules are located behind the fan unit in the Origin FibreVault or Fibre Channel RAID enclosure, as shown in Figure 5-5. This procedure is fully applicable only 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 explains

- preparing the new power supply
- removing a power supply
- inserting a power supply

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

5.3.1 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 a space available to put the fan unit and the power supply you are replacing.
2. Make sure the lever on the new power supply is in the unlocked position, as shown in Figure 5-7.

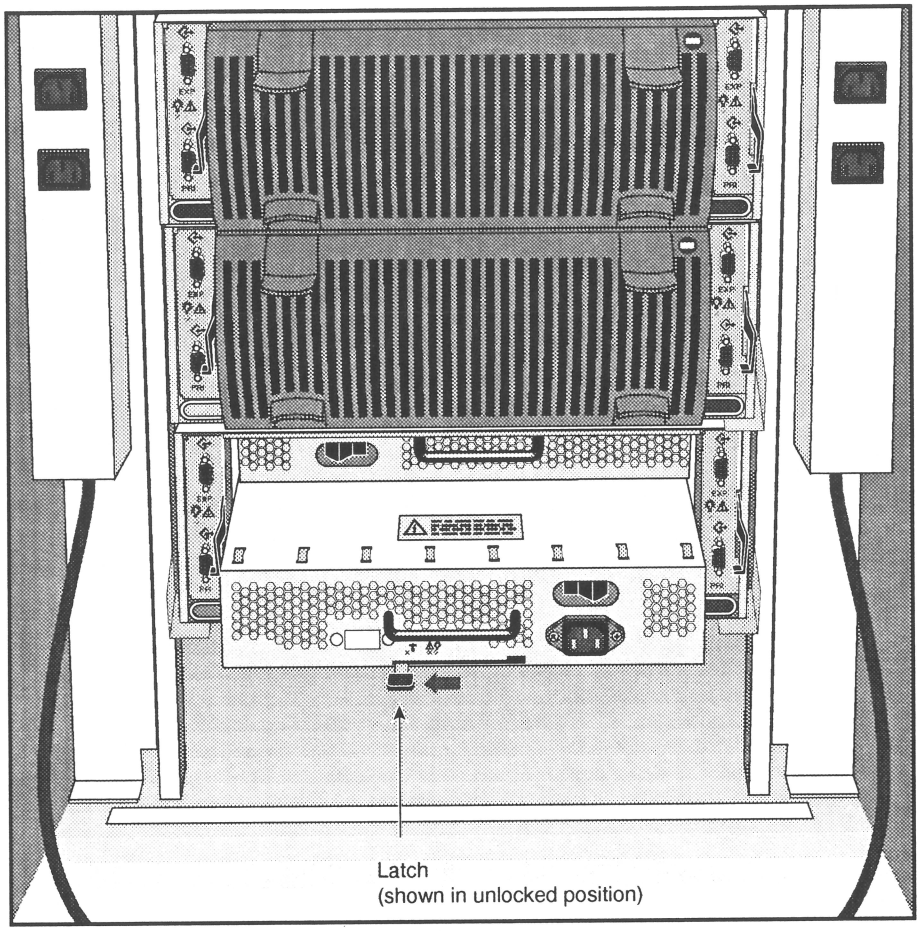

Figure 5-7 Unlatching a New Power Supply

5.3.2 Removing and Replacing a FibreVault Power Supply or Filler

To remove a power supply from a FibreVault enclosure, follow these steps:

Caution: 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.

1. Identify the Origin FibreVault 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 Origin FibreVault 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 5-5. Do not leave the fan assembly out of the Origin FibreVault or Fibre Channel RAID 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 (Figure 5-8).

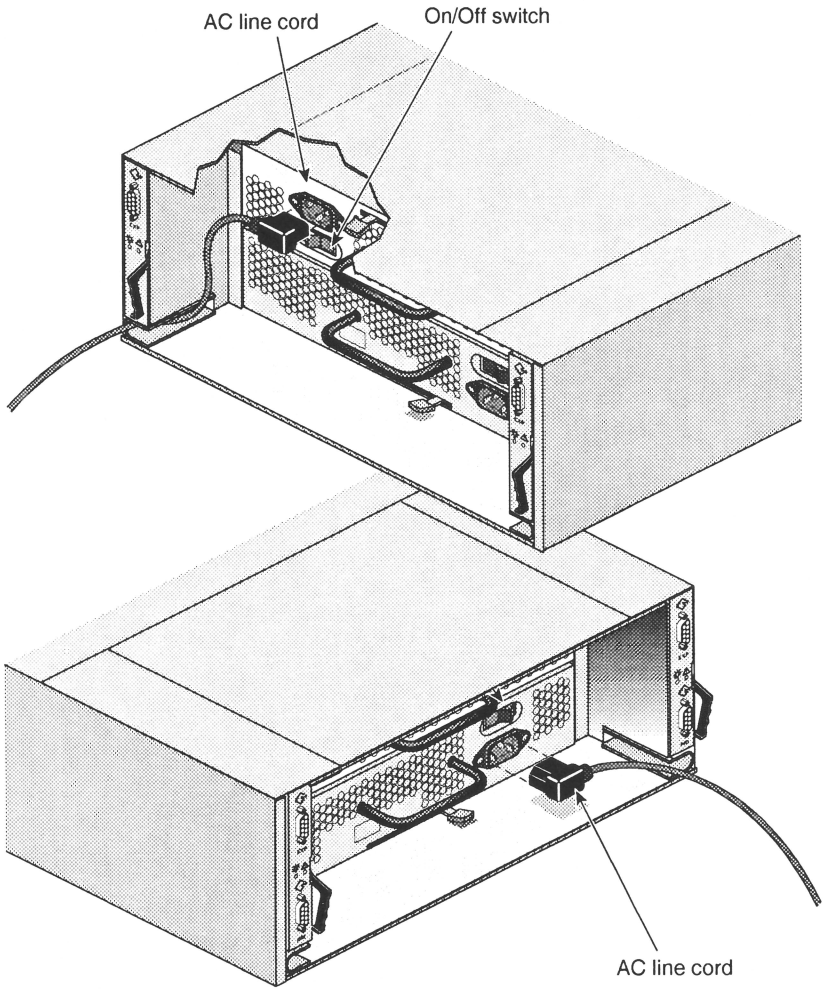

Figure 5-8 Power Supply Power Cord

6. Unlatch the faulty power supply (see Figure 5-9):
   - 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.

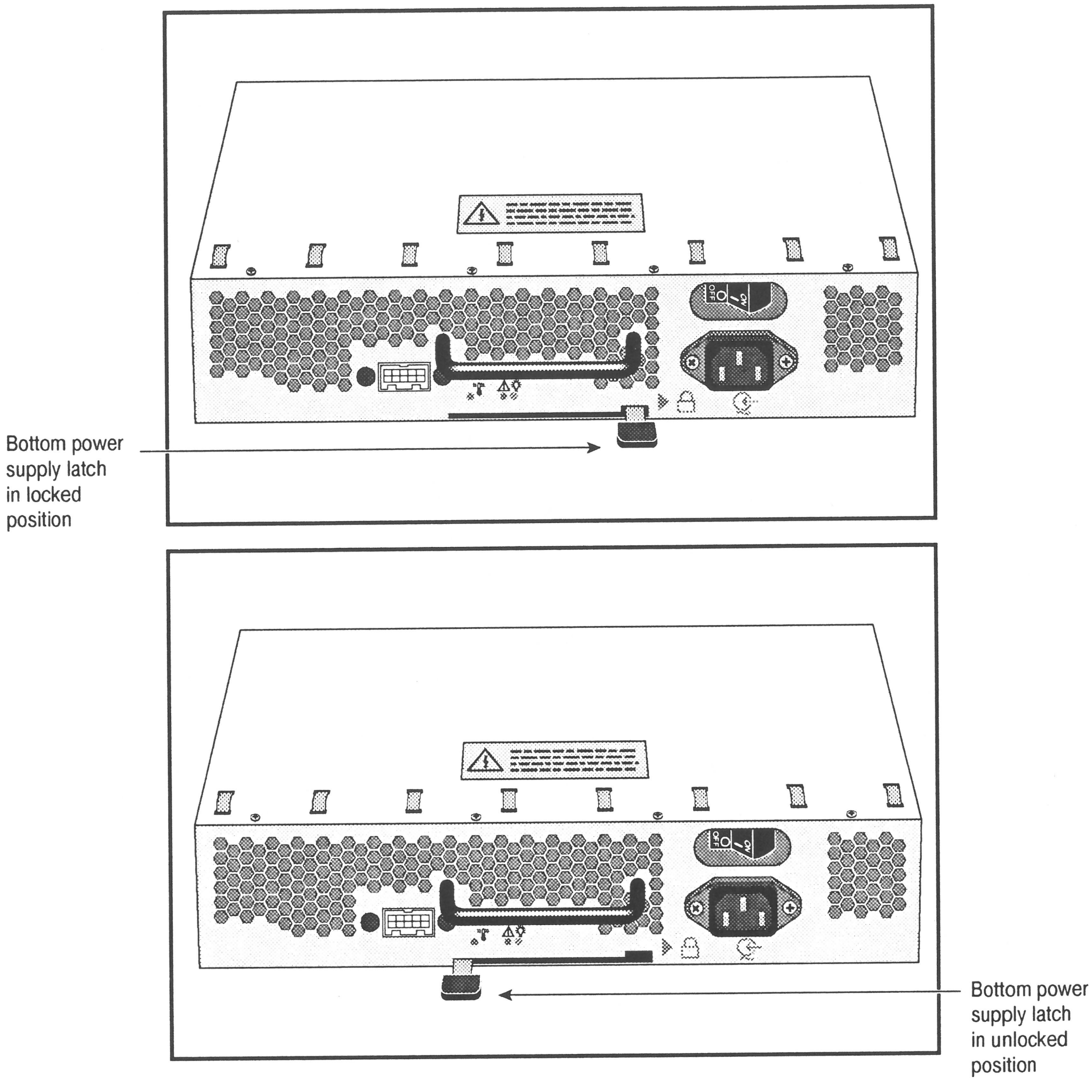

Figure 5-9 Unlatching a Power Supply

7. Pulling the handle, slide the filler or power supply out of the bay, supporting it with your other hand, as shown in Figure 5-10.

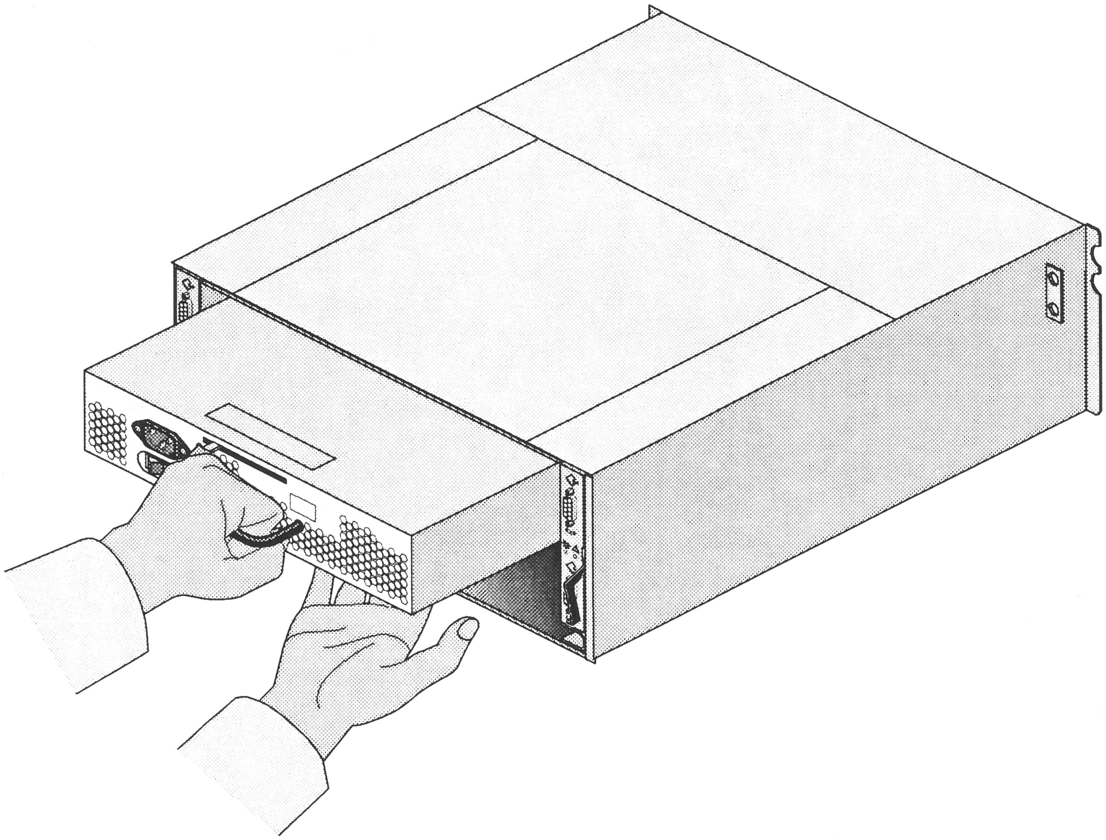

Figure 5-10 Removing a FibreVault Power Supply

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

    The power supply's active LED turns on.
12. Reattach the fan assembly by squeezing the latches toward each other and pushing the assembly into the chassis.

    You can install the fan assembly in either orientation. However, for consistency among multiple FibreVault enclosures, install it with the status light in the upper right corner as shown in Figure 5-6.

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

5.3.3 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 Section 5.2). 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) also shuts down. 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 (Figure 5-11).
3. Remove the power supply:
   - Unlatch the top power supply by moving the handle up and to the far right.
   - Unlatch the bottom power supply by moving the handle down and to the left.
4. Pull the supply straight out using the handle, see Figure 5-12 for an example.

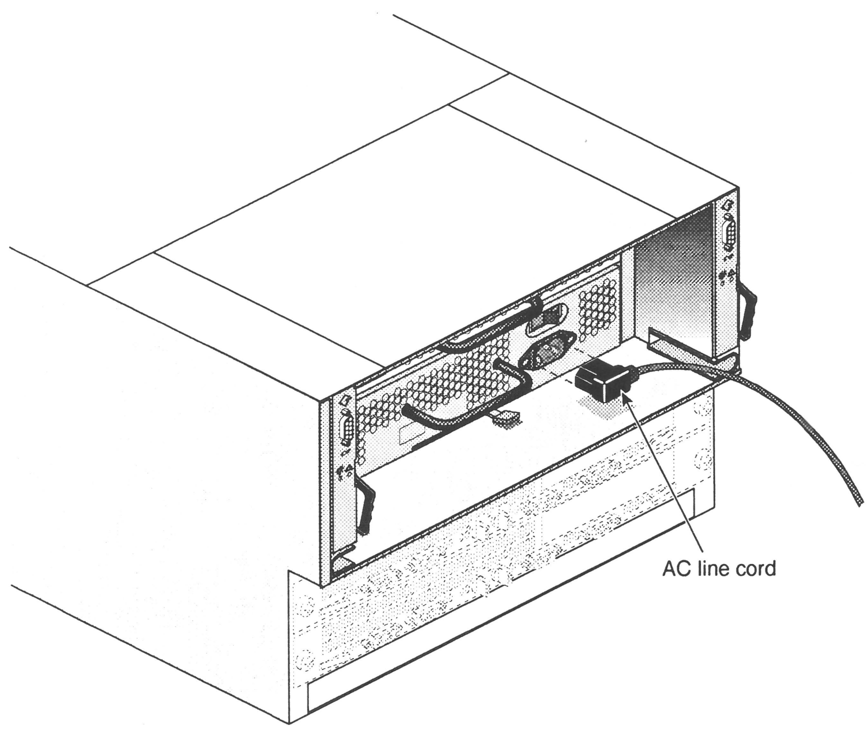

Figure 5-11 Turn Off Power and Unplug Cord from the Supply

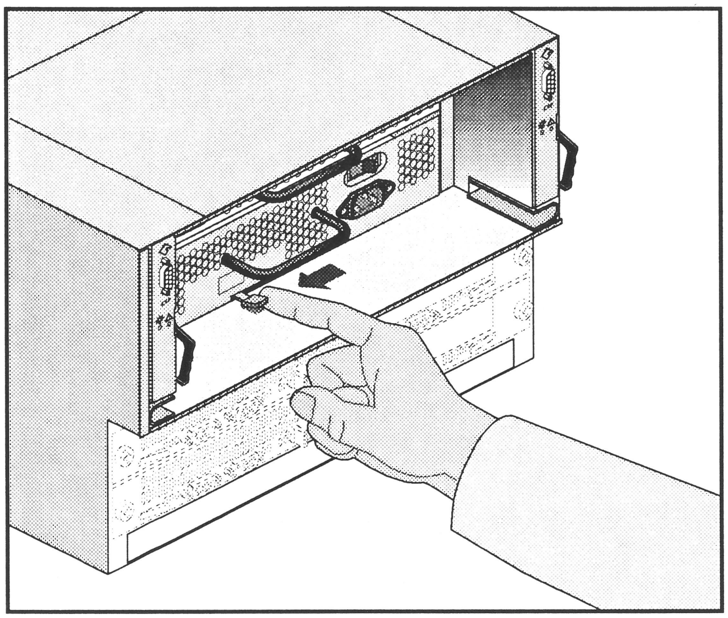

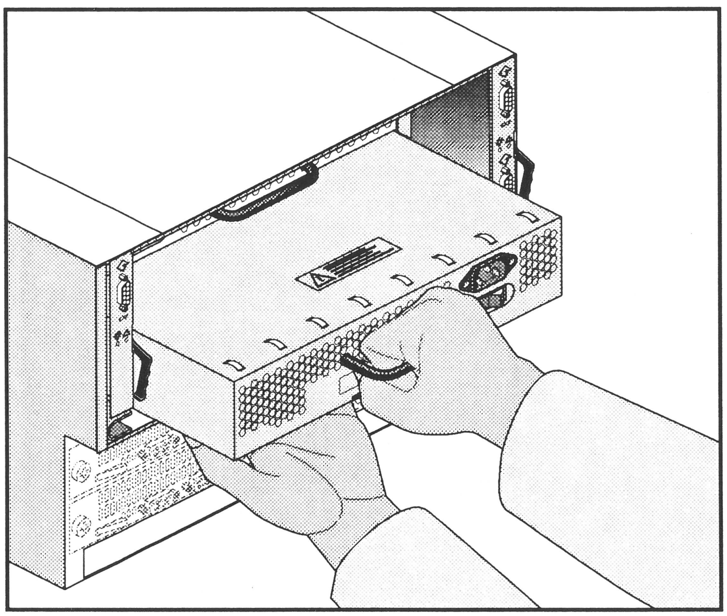

Figure 5-12 Removing an Existing Fibre Channel RAID Power Supply or Filler Panel

5. 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 from that of the top.

6. Push the supply into place until you see the latching lever move toward the latched position (see the example in Figure 5-13), and with your thumb or fingers, push the latch all the way into the locked position.

7. 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.

8. Press the supply's rocker switch to the On position.
9. Reinstall the fan module (see Section 5.2).

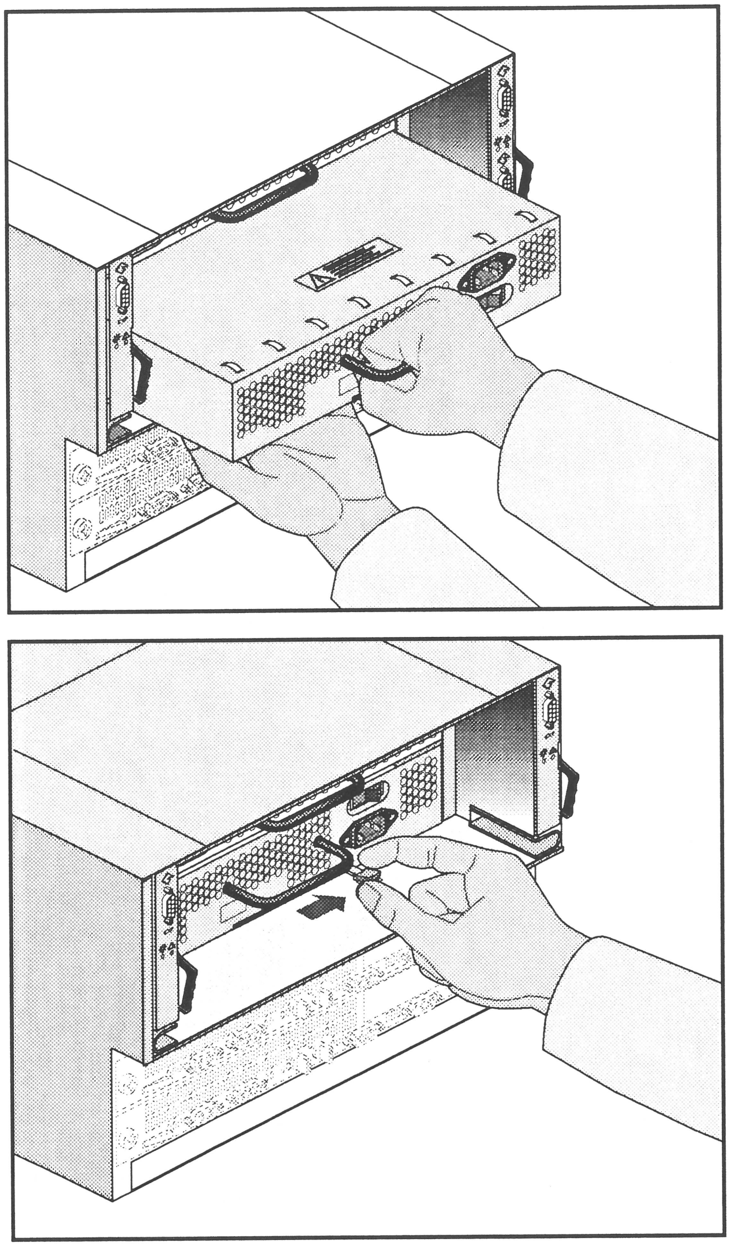

Figure 5-13 Installing the Bottom Fibre Channel RAID Power Supply

5.4 Replacing or Adding a Link Control Card (LCC)

This section explains removing and replacing an LCC or filler panel in a FibreVault or Fibre Channel RAID enclosure.

- removing an LCC or filler panel
- inserting an LCC

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 you are ready to install it.

2. Locate the FibreVault or Fibre Channel RAID enclosure containing the LCC with its check LED illuminated.
3. If the enclosure is mounted in a rack, open the rear door.
4. Attach the clip of an 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.
5. Unplug the cable(s) (as applicable) connected to the faulty LCC, then pull up on the latch of the LCC or filler panel.

   The LCCs in both the FibreVault and Fibre Channel RAID unlatch, remove, and install in the same manner.

6. Grasp the LCC and slowly pull it out of the chassis, as shown in Figure 5-14. The LCC's active and check LEDs turn off if the enclosure is powered on.
7. Align the new LCC with the guides in the empty bay.
8. Slide the new LCC into the slot, as shown in Figure 5-15.
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 that you unplugged in step 5.
11. Remove and store the ESD wristband, then close the enclosure front door.

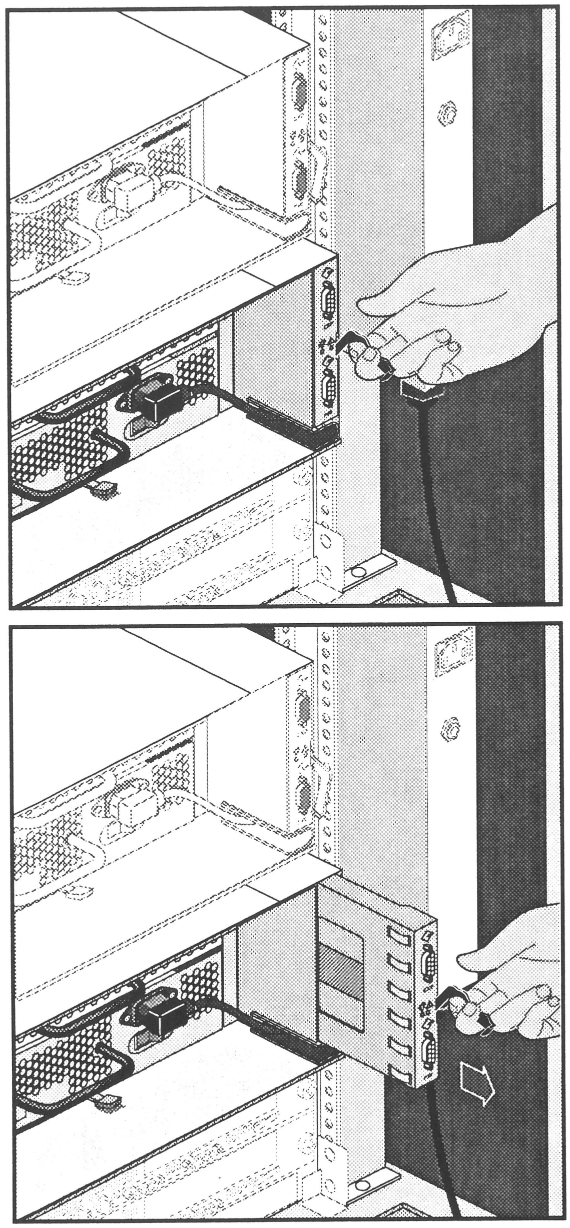

Figure 5-14 LCC Removal Example

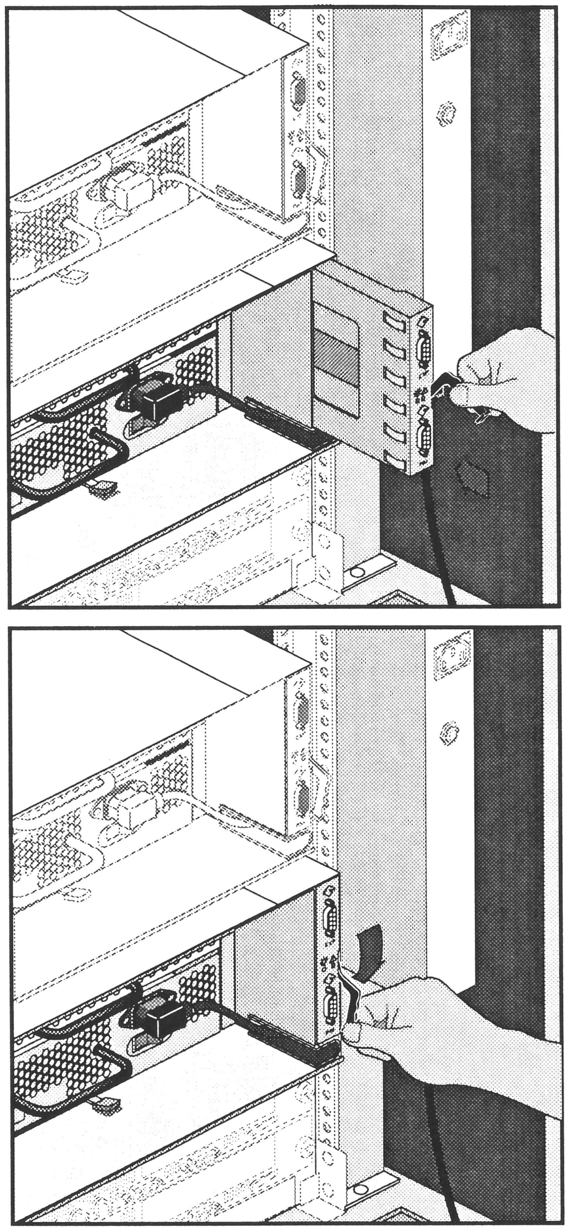

Figure 5-15 LCC Installation Example

Appendix A

A. Technical Specifications

This appendix contains

- site requirements for the Fibre Channel Rack and FibreVault Deskside tower
- technical specifications for the FibreVault
- technical specifications for the FibreVault deskside (tower)
- technical specifications for the Fibre Channel RAID
- technical specifications for the Fibre Channel Rack
- cabling interface descriptions

A.1 Site Requirements for the Fibre Channel Rack and FibreVault Deskside

Before installing the fibre channel rack, make sure the site meets the power and thermal requirements for the fibre rack and its enclosures. See Table A-4 for specific Fibre Channel Rack information. See Section A.6 for cabling interface 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-1 and Table A-2 for all questions regarding the deskside FibreVault.

A.1.1 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 enclosure.

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 are less depending on the number and manufacturer of disk drives. These values represent either the values for the line cord of an enclosure with a single power supply, or the total values shared by the line cords of two power supplies in the same enclosure, with the division between the line cords and supplies at the current sharing ratio.

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.

A.1.2 Enclosure Thermal Requirements

The ambient temperature specification is measured at the enclosure's front door inlet. (See Table A-2 for specifications.) 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).

A.2 Technical Specifications for the FibreVault Enclosure

Table A-1 lists specifications for the FibreVault enclosure (DAE). Note that these specifications are virtually identical for the deskside (standalone) FibreVault enclosure.

Table A-1 FibreVault Enclosure (DAE) Technical Specifications

Classification	Specification	Value
AC and current	AC line voltage	100 to 240 VAC single phase (47 to 63 Hz)
	AC line current	4.5 amps maximum at 100 VAC
	Input power	441 watts maximum
	Inrush current	30.0 amps maximum for 0.5 line cycles
	AC circuit breaker	12 amp (magnetic)
	AC inlet type	IEC320-C14 appliance coupler
	Drive module power	25 watts per bay maximum
Operating limits	Ambient temperatures	10°C to 40°C (50°F to 104°F) 40°C (104°F) at 2,438 m (8,000 ft) 37°C (98.6°F) at 3,050 m (10,000 ft)
	Temperature gradient	10°C per hour (50°F per hour)
	Relative humidity	20% to 80% noncondensing
	Heat dissipation	1588x10^3 J/hr (1500 BTU/hr) maximum
Nonoperating limits	Ambient temp	–40°C to 65°C (–40°F to 149°F)
	Temp gradient	24°C/hr (43.2°F/hr)
	Relative humidity	10% to 90% noncondensing
	Elevation	7625 m (25,000 ft)

Table A-1 (continued) FibreVault Enclosure (DAE) Technical Specifications

Classification	Specification	Value
Dimensions	Height	15.4 cm (6.06 in.) 3.5U (NEMA units)
	Width	44.5 cm (17.5 in.)
	Depth	63.27 cm (24.9 in) front door to back of fan pack
Weights	Minimum weight	32.0 kg (71 lbs) without redundant components
	Maximum weight	36.0 kg (80 lbs) in high-availability configuration

A.3 Technical Specifications for the FibreVault Deskside (Tower)

Table A-2 lists specifications for the standalone deskside tower version of the FibreVault. All environmental operating parameters are the same as those for the rackmounted FibreVault (see Table A-1).

Height, weight, and dimensions are the primary differences between the two products. Most other specifications are the same.

Table A-2 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 amp (magnetic)
	AC inlet type	IEC320-C14 appliance coupler (3 outlets)
Dimensions	Tower height	67.95 cm (26.75 in.)
	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	36.0 kg (80 lbs) empty

A.4 Technical Specifications for the Fibre Channel RAID Enclosure

Table A-3 lists specifications for the Fibre Channel RAID Enclosure, also known as the Disk-array Processor Enclosure (DPE).

Table A-3 Fibre Channel RAID Enclosure (DPE) 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	790 watts maximum
	Inrush current	50.0 amps maximum for 0.5 line cycle
	AC circuit breaker	12 amp (magnetic)
	AC inlet	IEC320-C14 appliance coupler
	Drive module power	25 watts per bay maximum
Operating limits	Ambient temperatures	10°C to 40°C (50°F to 104°F) 40°C (104°F) at 2,438 m (8,000 ft) 37°C (98.6°F) at 3,050 m (10,000 ft)
	Temperature gradient	10°C per hour (50°F per hour) maximum
	Relative humidity	20% to 80% noncondensing
	Heat dissipation	3165 x 10^3 J/hr (3000 BTU/hr) maximum
Nonoperating 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	7625 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	43 kg (95 lbs)
	Maximum weight	49.6 kg (110 lbs)

A.5 Technical Specifications for the Fibre Channel Rack

Table A-4 lists specifications for the fibre channel rack.

Table A-4 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 @ 200 VAC	2.3 amps per FibreVault or RAID expansion 4.1 amps per Fibre Channel RAID enclosure
	Maximum current per rack	24.6 amps maximum
	Power consumption at 200 VAC	225 watts per FibreVault or RAID expansion 390 watts per Fibre Channel RAID enclosure
Power connector	Type	North America: L6-30P International: IEC Plug 309 2-P
Operating limits	Ambient temperature	10°C to 40°C (50°F to 104°F) 40°C (104°F) at 2,438 m (8,000 ft) 37°C (98.6°F) at 3,050 m (10,000 ft)
	Relative humidity	20 to 80% noncondensing
	Heat dissipation	3168 x 10^3 J/hr (3000 BTU/hr) max per chassis assembly
	Maximum heat dissipation for rack	19008 x 10^3 J/hr (18,000 BTU/hr)
	Shock	3 g @ 11 ms
	Vibration	0.25 g peak @ 5 Hz to 500 Hz
Nonoperating limits (shipping and storing)	Ambient temperature	-40°C to 65°C (-40°F to 149°F)
	Temperature gradient	24°C/hr (43.2°F/hr) maximum
	Relative humidity	10 to 90% noncondensing
	Elevation	7625 m (25,010 ft)
Physical	Cabinet dimensions	Height: 192 cm (75.3 in.)
		Width: 65 cm (25.5 in.)
		Depth: 87 cm (34.25 in.)
	Max weight	Approx. 500 kg (1100 lb)
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 Gbit per second transfer rate)

A.6 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 male 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 an MIA is available.

The optical interface (known as the media interface adapter or MIA) uses a shortwave laser (CD-ROM laser) with a wavelength of 780 nm. Silicon Graphics supports optical cabling distances of 25 to 300 meters (between the host and the enclosure). This interface is used with 62.5 µm optical cable and the distance is about 60% of a 50 µm cable. Silicon Graphics fibre products support a maximum distance of 300 meters.

Table A-5 summarizes characteristics for the copper twinax and fiber optic cables.

Table A-5 Cable Characteristics

Characteristic	Copper Twinax	Fibre Optic
Variant	100-TW-EL-S full-duplex (quad conductor)	62.5 µm Type: 780 nm
Range	25 meters maximum	300 meters maximum
Transmitter	PECL	100-M5-SL/N-I
Medium	150 ohm	0.002-0.5
Range	25 m	300 m
Connector	Male DB9; see below	SC duplex connector
Maximum length	Enclosure to server: 25 m Enclosure to enclosure: 0.3 m	Enclosure to server: 300 m Enclosure to enclosure: not used

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.

A copper cable connecting a fibre enclosure to the host system is either 10 or 25 meters. Cables longer than 10 meters must be equalized.

Fibre enclosures are connected only by copper cables, which must be no longer than 10 m (0.3 meters is the Silicon Graphics standard). Fibre enclosure interconnections must 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 connectors for both ends of the twinax cable are male DB9 connectors that plug in at the host and enclosure. 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, there is a potential for grounding problems 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.

Appendix B

Fibre Channel Command-Line Administration

This chapter explains how to use the fibre channel command-line interface (CLI) software to administer, monitor, and adjust FibreVault (non-RAID) and Fibre Channel RAID disk arrays. The command-line interfaces are used when no system graphics are present or if the user is more comfortable with the command line than with a graphical user interface. Either CLI is accessible using a serial-port terminal, remote login, or laptop connected to the host system.

The non-RAID and RAID command-line interfaces may be used with Origin2000 servers or Onyx2 systems. The first part of this Appendix covers the command-line administration of non-RAID FibreVault arrays. The second part covers Fibre Channel RAID command-line administration options.

B.1 Using the (Non-RAID) Command-Line Interface

The non-RAID FibreVault CLI is also generally known as the fibre channel CLI or *fccli*. The fccli works directly with the fibre channel "agent" or *fcagent*.

B.1.1 Fibre Channel Agent (fcagent)

The *fcagent* must be running before the CLI commands are accepted and executed by the fibre channel host system. To start the *fcagent*:

1. Make sure that the *fcagent* flag is "on" by entering *chkconfig fcagent on*.
2. If the agent is not running, you can launch it by entering */etc/init.d/fcagent start*.
3. You can stop the agent by entering */etc/init.d/fcagent stop*.

B.1.2 CLI Command Options

The CLI command options list looks similar to the following:

```
Command Options:

# fccli
Usage : fccli GETVERSION [-h hostname]
      : fccli GETSTATUS [-h hostname] [-c channel ID [ -e eid | -t tid1 [tid2
... [tidn]] ]
       : fccli REMOVE [-h hostname] -c channel ID -t tid1 [tid2 ... [tidn]]
       : fccli INSERT [-h hostname] -c channel ID -t tid1 [tid2 ... [tidn]]
       : fccli LEDON [-h hostname] -c channel ID -t tid1 [tid2 ... [tidn]]
       : fccli LEDOFF [-h hostname] -c channel ID -t tid1 [tid2 ... [tidn]]
       : fccli BYPASS [-h hostname] -c channel ID -t tid1 [tid2 ... [tidn]]
     : fccli UNBYPASS [-h hostname] -c channel ID -t tid1 [tid2 ... [tidn]]
```

B.1.3 Determining the *fcagent* Version Number

To obtain the version number of the *fcagent* that is running on a particular host system, you have two options:

- Get the version information directly from the host system's terminal interface.
- Get the version information remotely.

To get the revision information directly from the host system running the *fcagent*, enter `fccli GETVERSION`

If you are remotely logged into the host, use the command *fccli GETVERSION -h <hostname>*.

A sample output from an *fccli GETVERSION* inquiry would look similar to

```
FC JBOD Agent version 1.1, Wed Mar 12 08:04:26 PST 1997.
```

Note: In this example, the fibre channel agent is identified as a JBOD, or non-RAID, agent version.

B.1.4 Channel and Enclosure Status

To get a general channel and enclosure status for all the fibre channel devices on a host system, use the inquiry *fccli GETSTATUS* at the host, or *fccli GETSTATUS -h hostname* if you are remotely accessing the system.

This command (*fccli GETSTATUS*) provides you with output that looks similar to the following sample output:

```
CHANNEL| DRIVES    POWER     FANS      PEER LCCS
-------+---------------------------------------
01     | OK        FAIL      OK        OK
02     | OK         OK       OK        OK
03     | OK         OK       OK        OK
04     | OK         OK       OK        OK
05     | OK         OK       OK        OK
06     | OK         OK       OK        OK
```

Note that in channel 01 in the previous sample output, a power error is shown. Because the drives, fans, and link control cards (LCCs) show OK in the GETSTATUS output example, you can conclude that there is probably trouble with one of the two power supplies in the enclosure.

Note: In this example channel 01 controls three enclosures. Each fibre channel site will have different configurations installed and a single channel may control up to 11 (non-RAID) enclosures within the same rack.

To determine more about the power failure in the enclosure controlled by fibre channel 01, use the command *fccli GETSTATUS -c 01*. This command causes the *fcagent* to output a more detailed screen display showing the status of the enclosure(s) controlled by channel 01. The example output shows that a power supply has failed:

```
ENCLOSURE | DRIVES    POWER     FANS      PEER LCC
----------+---------------------------------------
      4   |  OK        FAIL      OK         OK
      5   |  OK        FAIL      OK         OK
     11   |  OK        FAIL      OK         OK
```

Always physically inspect the power supply before ordering a new part. The *GETSTATUS* command reports an unplugged power supply or a power supply with its breaker set at 0 (off), as a "failed" power supply. A visual confirmation of the supply is recommended.

If you try to obtain status on a channel that does not exist on the host, the *fcagent* reports an error:

```
# fccli GETSTATUS -c 4
Channel 4 does not exist on 'localhost'.
```

For a further enhanced output of information or to specify one particular enclosure supported by a particular channel, use the **-c** and **-e** arguments to *GETSTATUS*. Entering *fccli GETSTATUS -c 01 -e 01* generates an output that provides specific information for the channel and enclosure you specify. The following example shows a sample result of the request for information on channel 01 and enclosure 01:

```
Enclosure 1, status FAIL
  Vendor ID:      CLARiiON
  Product ID:     DISK ENC
  LCC SN#:
  LCC Mfg. Date:
  LCC code rev:   003
  FRUs:           10 disk slot(s)
                  1 fan slot(s)
                  2 power supply slot(s)
                  1 peer LCC slot(s)
        +---------------------------------------+------+
        | [1] FAIL                              |      |
        +---------------------------------------+      |
        |   |   |   |   |   |   |   |   |   |   | P    |
        | O | O | O | O | O | O | O | O | O | O | E  L |
        | K | K | K | K | K | K | K | K | K | K | E  L |
        | 10| 11| 12| 13| 14| 15| 16| 17| 18| 19| R  C |
        +---------------------------------------+      |
        |              FANS  OK                 |  OK  |
        |      PS0 FAIL          PS1  OK        |      |
        +---------------------------------------+------+
```

If you select a nonexistent (or offline) enclosure with a *GETSTATUS* inquiry, the *fcagent* reports an error similar to the following example:

```
fccli GETSTATUS -c 1 -e 0
Enclosure 0 does not exist on channel 1 on host 'localhost'.
```

B.1.5 Taking Non-RAID Drives Offline

When you need to remove a non-RAID fibre channel drive from an operating enclosure you can use the *fccli REMOVE* command. You must use the **-c** (channel designator argument) and **-t** (target drive argument) with the command in order to properly spin down the target disk.

In the following example, the 10th disk on channel 12 is spun down and taken out of the fibre channel operational loop.

```
fccli REMOVE -c 12 -t 10

        ~~~~~ Start Output from fcagent deamon ~~~
        ********************************************
        CALLOUT - /usr/lib/fcagent/preremove.CO
        HOSTNAME      = perf39
        CALLOUT TYPE = PRE-REMOVE
        TIMESTAMP     = Fri Feb 14 11:21:49 199x
        CHANNEL ID    = 12
        ENCLOSURE ID = -1
        FRU TYPE ID   = DISK
        FRU ID        = 10
        FROM STATUS   = Unknown
        TO STATUS     = Unknown
        ********************************************
        ~~~~~ End Output from fcagent deamon ~~~
```

```
~~~~~ Start Output from fcagent deamon ~~~
       *******************************************
       CALLOUT - /usr/lib/fcagent/postremove.CO
       HOSTNAME     = perf39
       CALLOUT TYPE = POST-REMOVE
       TIMESTAMP    = Fri Feb 14 11:21:49 199x
       CHANNEL ID   = 12
       ENCLOSURE ID = -1
       FRU TYPE ID  = DISK
       FRU ID       = 10
       FROM STATUS  = Unknown
       TO STATUS    = Unknown
       *******************************************
       ~~~~~ End Output from fcagent deamon ~~~
```

The targeted drive can now be physically removed from the enclosure.

If you use the *fccli REMOVE* command to try to remove a disk on a channel that does not exist, you receive an error message similar to the following:

```
fccli REMOVE -c 12 -t 11
Drive 11 does not exist on channel 12 on host 'localhost'.
```

In this case, channel 12 controls only one enclosure with a total of ten disks installed. The targeted drive (drive 11) does not exist and cannot be taken off line.

B.1.6 Bringing Up a Replacement Drive

After you physically install the new drive, you need to alert the host system to spin up the drive and bring it online. This spin-up process uses the `INSERT` function of the *fcagent* and requires that the proper channel (**-c**) and disk targeted (**-t**) be provided.

The following example arguments to the *fcagent* would spin up the replacement drive and bring it back online.

```
fccli INSERT -c 12 -t 10

 SPINNING UP not implemented yet
 DRIVE INSERTION COMPLETED
        ~~~~~ Start Output from fcagent deamon ~~~
         *******************************************
         CALLOUT - /usr/lib/fcagent/postinsert.CO
         HOSTNAME     = perf39
         CALLOUT TYPE = POST-INSERT
         TIMESTAMP    = Fri Feb 14 11:34:45 1997
         CHANNEL ID   = 12
         ENCLOSURE ID = -1
         FRU TYPE ID  = DISK
         FRU ID       = 10
         FROM STATUS  = Unknown
         TO STATUS    = Unknown
         *******************************************
        ~~~~~ End Output from fcagent deamon ~~~
```

This example uses drive 10 on channel 12. If you give the *fcagent* an invalid channel, enclosure, or drive number when bringing a new or replacement drive online, it returns an error message.

B.1.7 Turning Disk Check LEDs On or Off

Each disk in a FibreVault has two LEDs directly above it. The LED closest to the fibre channel disk drive is a green "activity" LED. The green LED lights during normal drive operations.

The amber LED is farther from the drive and lights when there is a problem with the drive, or the drive is offline, or missing. This amber LED can also be turned on using the **LEDON** function under the *fcagent* with the proper arguments.

The user or service person may wish to turn on the LED of a particular drive in order to confirm the identity of a channel, enclosure, or a particular drive within the system. This may be especially helpful in systems with many cables, enclosures, and drives. Proper cable and enclosure labeling should greatly reduce this type of identification problem.

If you needed to turn on the check LEDs for disk 3 and 5 in enclosure 1 on channel 12 you would enter the following:

```
fccli LEDON -c 12 -t 13 -t 15
```

```
+---------------------------------------+------+
| [1] FAIL                              |      |
+---------------------------------------+      |
|    |    |    | *  |    | *  |    |    |    | P      |   <-- '*' = Amber LED on
| O  | O  | O  | O  | O  | O  | O  | O  | O  | O  | E  L |
| K  | K  | K  | K  | K  | K  | K  | K  | K  | K  | E  L |
| 10| 11| 12| 13| 14| 15| 16| 17| 18| 19| R  C |
+---------------------------------------+      |
|                FANS  OK                 |  OK  |
|       PS0 FAIL           PS1  OK        |      |
+---------------------------------------+------+
```

The amber LEDs above disk drives 3 and 5 in enclosure 1 light up and stay lit until you turn them off with the LEDOFF request to the *fcagent*. Note in the sample output for **LEDON** that there is an asterisk above position 13 and 15. This indicates that the amber LED has been turned on for enclosure 1 drive 3, and enclosure one drive 5.

To turn off the amber check LED on the disk modules from the previous example, enter the following:

```
fccli LEDOFF -c 12 -t 13 -t 15
```

Note: The **LEDOFF** function is not intended for turning off a check LED on a drive that is experiencing an actual failure.

B.1.8 Bypassing a Non-RAID Drive

You can have the host bypass one or more disks in a specific enclosure on a specific channel by using the *BYPASS* command and the **-c** and **-e** arguments.

When you use *BYPASS* to take a drive offline, you are cutting off all I/O access. Be sure this does not affect local or remote users before implementing the order.

In the following example, five of the ten drives in enclosure four on channel one are taken offline.

```
# fccli BYPASS -c 01 -t 45 46 47 48 49
# fccli GETSTATUS -c 01 -e 4
 Enclosure 4, status FAIL
  Vendor ID:      CLARiiON
  Product ID:     DISK ENC
  LCC SN#:
  LCC Mfg. Date:
  LCC code rev:   003
  FRUs:           10 disk slot(s)
                  1 fan slot(s)
                  2 power supply slot(s)
                  1 peer LCC slot(s)

  +---------------------------------------+------+
  | [ 4] FAIL                             |      |
  +---------------------------------------+      |
  |   |   |   |   |   |   |   |   |   |   | P    |
  | O | O | O | O | O | B | B | B | B | B | E  L |
  | K | K | K | K | K | P | P | P | P | P | E  L |
  | 40| 41| 42| 43| 44| 45| 46| 47| 48| 49| R  C |
  +---------------------------------------+      |
  |               FANS  OK                |  OK  |
  |     PS0 FAIL             PS1  OK      |      |
  +---------------------------------------+------+
```

The *fcagent* always reports the bypassing of a drive(s) to */var/adm/SYSLOG*. If you checked the *SYSLOG* immediately after the last command was issued, you would see something similar to:

```
perf39 15# view /var/adm/SYSLOG
May  7 17:49:10 5D:perf39 fcagent[427]: FRU reconfiguration: DISK #5 in
enclosure 4 on channel 01: OK --> BYPASSED
May  7 17:49:10 5D:perf39 fcagent[427]: FRU reconfiguration: DISK #6 in
enclosure 4 on channel 01: OK --> BYPASSED
May  7 17:49:11 5D:perf39 fcagent[427]: FRU reconfiguration: DISK #7 in
enclosure 4 on channel 01: OK --> BYPASSED
May  7 17:49:11 5D:perf39 fcagent[427]: FRU reconfiguration: DISK #8 in
enclosure 4 on channel 01: OK --> BYPASSED
May  7 17:49:11 5D:perf39 fcagent[427]: FRU reconfiguration: DISK #9 in
enclosure 4 on channel 01: OK --> BYPASSED
```

B.1.9 Unbypassing a Non-RAID Drive

To bring bypassed drives back online, have the host "unbypass" the disks in an enclosure that have been bypassed. Use the *UNBYPASS* command and the appropriate **-c** and **-e** arguments.

When you use *UNBYPASS* to bring drives back online, the host has almost immediate access. Be sure this does not affect local or remote users before implementing the order.

In the following example, five of the ten drives in enclosure 4 on channel 1 were offline. Use the following command to unbypass the drives and bring them back online:

```
# fccli UNBYPASS -c 01 -t 45 46 47 48 49
# fccli GETSTATUS -c 01 -e 4
 Enclosure 4, status FAIL
  Vendor ID:      CLARiiON
  Product ID:     DISK ENC
  LCC SN#:
  LCC Mfg. Date:
  LCC code rev:   003
  FRUs:           10 disk slot(s)
                  1 fan slot(s)
                  2 power supply slot(s)
                  1 peer LCC slot(s)


+----------------------------------------+------+
| [ 4] FAIL                              |      |
+----------------------------------------+      |
|    |    |    |    |    |    |    |    |    |    | P    |
| O  | O  | O  | O  | O  | B  | B  | B  | B  | B  | E  L |
| K  | K  | K  | K  | K  | P  | P  | P  | P  | P  | E  L |
| 40| 41| 42| 43| 44| 45| 46| 47| 48| 49| R  C |
+----------------------------------------+      |
|              FANS  OK                  |  OK  |
|      PS0 FAIL          PS1  OK         |      |
+----------------------------------------+------+
```

As with the *BYPASS* command, the *fcagent* reports the unbypassing of a drive(s) to */var/adm/SYSLOG*. If you checked the *SYSLOG* immediately after the last command was issued, you would see something similar to:

```
perf39 15# view /var/adm/SYSLOG
May  7 17:59:10 5D:perf39 fcagent[427]: FRU reconfiguration: DISK #5 in
enclosure 4 on channel 01: OK --> UNBYPASSED
May  7 17:59:10 5D:perf39 fcagent[427]: FRU reconfiguration: DISK #6 in
enclosure 4 on channel 01: OK --> UNBYPASSED
May  7 17:59:11 5D:perf39 fcagent[427]: FRU reconfiguration: DISK #7 in
enclosure 4 on channel 01: OK --> UNBYPASSED
May  7 17:59:11 5D:perf39 fcagent[427]: FRU reconfiguration: DISK #8 in
enclosure 4 on channel 01: OK --> UNBYPASSED
May  7 17:59:11 5D:perf39 fcagent[427]: FRU reconfiguration: DISK #9 in
enclosure 4 on channel 01: OK --> UNBYPASSED
```

B.2 Administration Tips for Non-RAID Drives

There are some non-*fcagent* inquiries you can make to the operating system that can help in the administration or troubleshooting of the non-RAID FibreVault disks and enclosures.

B.2.1 Obtaining a Fibre Channel Hardware Graph Output

To obtain hardware graph information for a local host system that is controlling a fibre channel disk array you can enter a grep inquiry such as the following:

```
grep ^controller_number /etc/ioconfig.conf
```

For example, if you wanted to get information on channel 12, you would enter:

```
grep ^12 /etc/ioconfig.conf
```

and see an output similar to:

```
12 /hw/module/1/slot/io4/fibre_channel/pci/0/scsi_ctlr/0
```

This tells you that module 1, slot io4, has a fibre channel enclosure connected to pci 0 (the first channel on the fibre XIO board).

```
12 /hw/module/1/slot/io4/fibre_channel/pci/0/scsi_ctlr/0
```

Detailed layout

```
target/10/lun/0/disk/partition/7
```

Detailed descriptors

```
--> target  -->  targ_id --> lun --> lun_# --> disk --> partition

--> part_# --> char or block
```

B.2.2 Listing Fibre Channel Controller and Drive Numbers

You can use the *hinv* command at the operating system level to determine numbers and types of fiber devices on a particular host system.

Using *hinv* output to get the fibre channel controller numbers listed would look similar to:

```
hinv | grep Fibre

Integral SCSI controller 12: Version Emerald Fibrechannel
Integral SCSI controller 13: Version Emerald Fibrechannel
Integral SCSI controller 10: Version Emerald Fibrechannel
Integral SCSI controller 11: Version Emerald Fibrechannel
Integral SCSI controller 22: Version Emerald Fibrechannel
Integral SCSI controller 23: Version Emerald Fibrechannel
```

If you wanted to use *hinv* output to get the listing of device(s) connected to the fibre channel controller(s), you would enter and see something like:

hinv -c disk

```
Integral SCSI controller 12: Version Emerald Fibrechannel
Disk drive: unit 19 on SCSI controller 12
Disk drive: unit 17 on SCSI controller 12
Disk drive: unit 15 on SCSI controller 12
Disk drive: unit 13 on SCSI controller 12
Disk drive: unit 11 on SCSI controller 12
Disk drive: unit 18 on SCSI controller 12
Disk drive: unit 16 on SCSI controller 12
Disk drive: unit 14 on SCSI controller 12
Disk drive: unit 12 on SCSI controller 12
Disk drive: unit 10 on SCSI controller 12
```

B.3 Fibre Channel RAID Command-Line Interface

This section covers information related to the command-line utility used for administering RAID arrays. These commands do not work with non-RAID fibre drives.

The storage system manager command-line interface (*ssmcli*) utility acts as an interface to the RAID storage processor (SP) and disk-array storage system. With *ssmcli*, you can configure, control, and retrieve status from the RAID array and SP(s).

You can run *ssmcli* either from the local host or from a remote workstation.

Any *ssmcli* command consists of the *ssmcli* command with its arguments together with other commands and their arguments. The *ssmcli* command, if used on its own, does not perform any function. Using *ssmcli* with the other commands described in this section, you can configure and manage Fibre Channel RAID disk array storage systems.

The overall structure of the *ssmcli* command is defined in the next section, and then the commands used with the *ssmcli* command are described. See Appendix C for a description of the Fibre Channel RAID graphical user interface.

B.3.1 Fibre Channel RAID Agent (ssmagent)

The *ssmagent* must be running before the RAID CLI commands are accepted and executed by the fibre channel host system.

To start the *ssmagent*:

1. Make sure that the *ssm* flag is "on" by entering *chkconfig ssm on*.
2. If the agent is not running, you can launch it by entering */etc/init.d/ssm start*.
3. You can stop the agent by entering */etc/init.d/ssm stop*.

B.3.2 Configuring the Path to ssm Executables in .cshrc

The following steps configure the host system's *.cshrc* file to execute *ssmcli* commands from any directory on the host. If this path is not set, you must enter the complete path to */usr/ssm/bin/ssmcli* each time you enter an *ssmcli* command.

Use the following information to set the proper path on the host system:

1. As the superuser, use your favorite line editor to open the *.cshrc* file.
2. Enter the following in the *.cshrc* file: `set path = ($path /usr/ssm/bin)`
3. Write and quit the *.cshrc* file.
4. Enter the following command sequence to reinitialize the *.cshrc* file.

   ```
   # source .cshrc
   ```

B.4 ssmcli Syntax

A synopsis of the *ssmcli* interface to the Fibre Channel RAID storage subsystems is

`ssmcli [-p] -[v|q] [-h` *hostname*`] [-d` *device*`]` *CMD optional-args*

The *ssmcli* command sends RAID storage device management and configuration requests to an application programmer's interface API on a local or remote host. The Fibre Channel RAID storage enclosure consists of several classes of hardware units, including disk drives, one or two power supplies, fan modules, an optional standby power supply (SPS) battery backup unit, and one or two RAID storage processors (SPs). The ssmcli Command Options are:

-p	Parses the *CMD* command without making a remote procedure call (RPC) to the API. If the string does not parse correctly, an error message prints to stderr; otherwise there is no output.
-v	Enables verbose error descriptions
-q	Suppresses error messages. This option may be useful when used as part of a script
-h *hostname*	Specifies the network name of destination server. The environment variable RaidAgentHost is the default value for the server if the **-h** switch is not present. The **-h** switch overrides the RaidAgentHost setting. If RaidAgentHost is not set and the **-h** switch is not present, the local host name is used.
-d *device*	Specifies the target RAID device. If you do not specify the **-d** switch, the environment variable RaidAgentDevice is the default value for the device. If RaidAgentDevice is not set and the **-d** switch is not present, an error is generated on all commands that need device information. Unless the environment variable indicates otherwise, every instance of the *ssmcli* command must have the **-d** device argument. The **-d** switch overrides the RaidAgentDevice environment variable. You can get a list of RAID devices present using the *getagent* command with *ssmcli*.

B.4.1 The CMD Options

The *CMD* part of the *ssmcli* command line is one of a set of commands to configure and manage a storage system.

Use the *CMD* commands described in Section B.5 with the *ssmcli* command. The *CMD* commands are:

bind	*getagent*	*getloop*
chglun	*getcache*	*getlun*
clearlog	*getcontrol*	*setcache*
clearstats	*getcrus*	*setstats*
failback	*getdisk*	*unbind*
firmware	*getlog*	

These commands are described later in this appendix.

B.4.2 Environment Variables RaidAgentHost and RaidAgentDevice

The environment variables RaidAgentHost and RaidAgentDevice are default values for the server and the device if **-h** and **-d** switches are not supplied with the *ssmcli* command set. See the **-h** and **-d** switch descriptions on the previous page. If RaidAgentDevice is not set and the **-d** switch is not present, an error will be generated on all commands that need device information. The following example illustrates a sample environment command, *setenv,* setting the RaidAgentDevice variable.

```
setenv RaidAgentDevice sc3d1l0
```

B.4.3 Verbose Mode Diagnostic Information

The verbose argument to the *ssmcli* command returns numerical diagnostic values.

Enabling the verbose mode (**-v**) with the *ssmcli* command displays detailed information about the returned value. The verbose error numbers and their message descriptions are:

Error Number	Error Description
0	No errors
1	Invalid command-line parameters
2	Too few parameters
3	Too many parameters
4	Invalid bind type
5	Invalid LUN number

Error Number	Error Description
6	Invalid rebuild-time
7	Disk-names for RAID 1/0 must have an even number of arguments
8	Invalid number of disks in bind command
9	Valid disk-names are of format a0, b1, ... e3, etc.
10	Invalid stripe size
11	Invalid disk name
12	Invalid cache flags
13	Invalid default owner
14	Incorrect number of chglun parameters
15	Unable to determine name of target host machine
16	Enable-disable flag invalid
17	Invalid usable cache size
18	Invalid page size
19	Invalid watermark value
20	Hi-watermark less than lo-watermark
21	No device name listed on command line
22	Invalid idle threshold
23	Invalid idle delay
24	Invalid write aside size
25	Disks must be on separate buses for bind
26	The agent was unable to configure any devices
27	LUN does not exist
28	LUN already exists
29	Cannot get current working directory for firmware command
30	Illegal position for hot spare
31	Switch not supported
32	Command is not supported
33	Unknown error
50	Agent - Memory allocation error
51	Agent - Invalid device name
52	Agent - SP not found
53	Agent - Invalid SP suffix

Error Number	Error Description
54	Agent - Invalid LUN number
55	Agent - Invalid page size
56	Agent - Invalid command-line switch
57	Agent - Error bind command
58	Agent - Error chglun command
59	Agent - Error unbind command
60	Agent - Error firmware command
61	Agent - Generic error
62	Agent - Unknown error

B.5 The bind Command

The *bind* command groups physical disks into logical RAID levels.

You must be logged in as the system administrator and you must be at the operating system's root directory.

The *bind* command is used with *ssmcli* as follows:

`ssmcli [-p] -[v|q] [-h` *hostname*`] [-d` *device*`] bind` *raid-type lun-number disk-names* `[-c` *cache-flags*`] [-n` *NMinus1Go*`] [-r` *rebuild-time*`] [-s` *stripe-size*`] [-v` *max-verify-time*`]` `[-z` *stripe-count*`]`

B.5.1 bind Command Arguments

The arguments in the following subsections are used with the *bind* command.

B.5.1.1 raid-type Argument

The *raid-type* argument to the *bind* command specifies the RAID level for the array. Available options are:

- r3 = RAID level 3
- hs = hot spare disk

Note: Only the **r3** and **hs** arguments are supported at first release (IRIX 6.4) of the Silicon Graphics fibre RAID.

B.5.1.2 lun-number Argument

The **lun-number** argument to the *bind* command assigns a logical unit number (LUN) to the unit.

The default LUN number is the next hexadecimal number available, starting with 0. Only available numbers appear in the list. The default number is 0 for the first LUN that you bind, regardless of the number of RAID storage processors (SPs) or hosts attached to the storage system. The default number for the second LUN you bind is 1; for the third LUN, it is 2; for the fourth LUN, it is 3, and so on.

Depending on the operating system running on the host, you can have a variable number of LUNs. If for any reason you want to specify a nondefault number you can do so. After you bind a LUN with a nondefault number, the default number for the next LUN is the lowest number you skipped. For hot spares, you should assign LUN numbers starting with the highest number available and continuing downwards.

B.5.1.3 disk-names Argument

The **disk-names** argument to the *bind* command indicates which physical disks (on that bus) to bind, in the format e_d, where:

- e = enclosure number
- d = disk number

For example, 2_3 represents enclosure 2 disk 3. Valid enclosure values are 0-10 and valid disk numbers are 0-9.

B.5.1.4 -c cache-flags Argument

The optional **cache-flags** argument to *bind* enables either write or read caching, both read and write caching, or neither type of caching for the LUN. If you enable any type of caching for the LUN, caching occurs only when the storage processor system caching is enabled. Whether or not you should use caching for a specific LUN depends on the LUN's RAID type. Recommendations for using caching for the different RAID types are listed later in this appendix. Select the type of caching used as follows:

- none = no caching
- read = read caching
- write = write caching
- rw = read and write caching

The default is none.

B.5.1.5 -n NMinus1Go Argument

The optional **NMinus1Go** argument to the *bind* command is applicable to RAID-3 only. To enable **NMinus1Go** functionality use one of the following:

B.5.1.6 -r rebuild-time Argument

The optional **rebuild-time** argument to the *bind* command is the amount of time allocated to reconstruct data on either a hot spare or a new disk module that replaces a failed disk module in a LUN. It applies to all RAID LUNs except RAID 0. Legal values are any number greater than or equal to zero, and the default is 4 hours. The time you specify determines the amount of resources that the SP devotes to rebuilding instead of to normal I/O activity. The default time of 4 hours is adequate for most situations. A rebuild time of 2 hours rebuilds the disk more quickly, but degrades response time slightly. A rebuild time of 0 hours rebuilds the disk module as quickly as possible, but degrades response time significantly. If your site requires fast response time and you want to minimize degradation to normal I/O activity, you can extend the rebuild over a longer period of time, such as 24 hours.

Note: You cannot change the rebuild time for a LUN while the LUN is rebuilding. The actual rebuild time can differ from the time you specify.

B.5.1.7 -s stripe-size Argument

To have the best flexibility when binding LUNs within an array, bind LUNs in the following order:

- RAID-3
- hot spares

B.5.1.8 -v max-verify-time Argument

The optional **max_verify_time** argument to the *bind* command specifies the maximum amount of time (0-254 hours) to spend rebuilding a LUN if an SP fails and is replaced during I/O operations. A longer value rebuilds the LUN more slowly and has less impact on overall system performance. The default value is 4 hours.

B.5.1.9 -z stripe-count Argument

The optional **stripe-count** argument to the *bind* command sets the number of stripes in a LUN. For example, when binding a RAID-5 with a stripe count of 2, partition the LUN into 2 stripes, thus preventing access to the remainder of the available space. This option is useful for short bind operations. Legal values are any number greater than or equal to zero. The default value is 0, which binds the maximum number of stripes available.

B.5.2 LUN Binding Recommendations

Before you start binding disk modules, read the restrictions and recommendations in this section.

You can bind only unbound disk modules. All disk modules in a LUN must have the same capacity to use the full capacity of all disk modules. If a LUN consists of different size disk modules, the smallest disk module will determine the total disk space available.

In a Silicon Graphics implementation of a RAID-3 LUN, you must bind exactly five or nine disk modules. You can bind one less module per LUN than you will use by selecting an empty slot icon. However, the LUN operates in a degraded mode until a module is installed in the empty slot, and the storage system integrates it into the LUN. You can select the modules in any order. You cannot bind a RAID-3 LUN unless you have allocated storage system memory for the LUN.

When using a hot spare, the capacity of a disk module bound as a hot spare must be at least as great as the capacity of the largest disk module that it might replace.

You can enable either write or read caching, both read and write caching, or neither type of caching for any type of LUN except for a RAID-3 LUN or a hot spare. If you enable any type of caching for the LUN, caching occurs only when storage system caching is enabled. Whether or not you should use caching for a specific LUN depends on the LUN's RAID type. Recommendations for using caching for the different RAID types are given in the next subsection.

B.5.3 Caching Recommendations for RAID-3 Types

You cannot enable either write or read caching for a RAID-3 LUN or a hot spare. When a LUN is assigned to an SP, its state becomes assigned, and the state of its disk modules become enabled when you use the SP that owns the LUN as the communications path to the chassis.

If you need to change the Description (RAID type) or Element Size parameters, you must unbind the LUN and then rebind it with the correct parameters. If you need to change the maximum rebuild time, read cache state, or write cache state, use the *chglun* command.

Example B-1 Caching a RAID-3 LUN

Bind LUN 0 as disks 0, 1, 2, 3, 4 from enclosure 1 into a RAID-3, with a 4-hour maximum rebuild time, a 1 block stripe size per physical disk, and read cache enabled.

```
bind r3 0 0_0 0_1 0_2 0_3 0_4
```

There is no output. Errors are printed to *stderr.*

B.6 The chglun Command

The *chglun* command changes parameters on a currently bound LUN. You must use at least two switches with *chglun* and one of the switches must be the LUN number -l lun.

You must be logged in as the system administrator and you must be at the operating system's root directory. The *chglun* command is used with *ssmcli* as follows:

`ssmcli [-p] -[v|q] [-h` *hostname*`] [-d` *device*`] chglun -l` *lun* `-a` *auto* `-c` *cache-flags* `-d` *default-owner* `-i` *idle-thresh* `-m` *mirror* `-n` *NMinus1Go* `-pc` *maximum-number-of-I/O's* `-pd` *minimum-read-request-at-which-prefetching-is-disabled* `-pf` *prefetch-type* `-pm` *prefetch-size/multiplier* `-r` *rebuild-time* `-t` *idle-delay-time* `-v` *verify-time* `-w` *write-aside*

Errors are printed to *stderr*.

B.6.1 chglun Command Arguments

The arguments in the following subsections are used with the *chglun* command.

B.6.1.1 -l lun Argument

The **-l** argument specifies which logical unit to change. This flag is mandatory.

B.6.1.2 -a auto Argument

The **auto** argument to *chglun* automatically reassigns ownership of a LUN to an active SP if an SP failure occurs. The values are:

- 1 = enable autoassignment for the specified LUN
- 0 = disable autoassignment for the specified LUN

B.6.1.3 -c cache-flags Argument

With the **cache-flags** argument you can enable either write or read caching, both read and write caching, or neither type of caching for the LUN. If you enable any type of caching for the LUN, caching occurs only if storage system caching is enabled.

Whether or not you should use caching for a specific LUN depends on the LUN's RAID type; see recommendations for using caching for the different RAID types. Determine the type of caching used as follows:

- none = no caching
- read = read caching
- write = write caching
- rw = read and write caching

The default is none.

B.6.1.4 -d default-owner Argument

The **default-owner** argument (**-d**) can change the default owner of the LUN.

- 1 = Change ownership of the LUN to SP A
- 0 = Change ownership of the LUN to SP B

B.6.1.5 -i idle-thresh Argument

The **idle-thresh** argument (**-i**) is the maximum number of I/Os that can be outstanding to a LUN and have the LUN still be considered idle. This is used to determine cache flush start time.

B.6.1.6 -m mirror Argument

The **-m mirror** argument enables or disables mirroring on the write cache.

- 0 = Disabled.
- 1 = Enabled.

B.6.1.7 -n NMinus1Go Argument

The **-n** (or NMinus1Go) argument is for RAID-3 only. It enables NMinus1Go (also called minimal latency reads) functionality.

B.6.1.8 -pc maximum-number-of-I/O's Argument

The **-pc** argument to *chglun* defines the maximum number of I/Os which the LUN can have outstanding and still perform prefetching.

B.6.1.9 -pd minimum-read-request-at-which-prefetching-is-disabled Argument

The **-pd** argument gives the minimum read request at which prefetching is disabled Specify 0XFFFF to leave the current value unchanged.

B.6.1.10 -pf prefetch-type Argument

The **-pf** argument gives the prefetch type and determines whether to prefetch data of a variable or constant length or to disable prefetching. The default is 2, "variable length."

- 0 = Disabled.
- 1 = Constant length.
- 2 = Variable length.

B.6.1.11 -pm prefetch-size/multiplier Argument

The **-pm** prefetch size/multiplier argument determines the amount of data prefetched for one host read request. For constant length prefetching, the prefetch size is the number of blocks to prefetch. For variable length prefetching, the prefetch multiplier is the amount of data to prefetch relative to the amount of data requested. For example, if the prefetch size is 8, the amount of data to prefetch is 8 times the amount of data requested. Specify 0XFFFF to leave the setting unchanged. The range is 0 - 32 and the default value is 4.

B.6.1.12 -r rebuild-time Argument

The **-r** rebuild-time argument is the maximum time (in hours) to rebuild a replacement disk module. The value can be any number greater than or equal to zero, and the default is 4 hours. The time that you specify determines the amount of resource that the SP devotes to rebuilding instead of to normal I/O activity. The default time of 4 hours is adequate for most situations. A rebuild time of 2 hours rebuilds the disk more quickly, but degrades response time slightly. A rebuild time of 0 hours rebuilds the disk module as quickly as possible, but degrades response time significantly. If your site requires fast response time and you want to minimize degradation to normal I/O activity, you can extend the rebuild over a longer period of time, such as 24 hours. rebuild-time applies to all RAID LUNs except RAID 0.

Note: You cannot change the rebuild time for a LUN while the LUN is rebuilding.

The actual rebuild time may differ significantly from the time you specify.

B.6.1.13 -t idle-delay-time Argument

The **idle-delay-time** is the amount of time (in 100 ms intervals) that a unit must be below the idle-threshold in order to be considered idle. Once a unit is considered idle, any dirty pages in the cache can begin idle time flushing.

B.6.1.14 -v verify-time Argument

The **-v** verify-time argument is optional. The default value is 4 hours.

The maximum amount of time (0-254 hours) to spend rebuilding a LUN if an SP fails and is replaced during I/O operations. A longer value rebuilds the LUN more slowly and has less impact on overall system performance.

B.6.1.15 -w write-aside Argument

The **-w** write-aside argument indicates the smallest write request size, in blocks, that will bypass the cache and go directly to the disks. RAID-3 LUNs do not use read or write caching.

B.7 The clearlog Command

The *clearlog* command clears the RAID SP error log and destroys the contents of the RAID storage controller's error log. All log counters are reset to zero.

To use *clearlog*, you must be logged in as the system administrator and you must be at the operating system's root directory. The *clearlog* command is used with *ssmcli* as follows:

`ssmcli [-p] -[v|q] [-h` *hostname*`] [-d` *device*`] clearlog`

B.8 The clearstats Command

The *clearstats* command resets statistics logging on the RAID SP. If statistics logging is on, it is turned off and then back on. If statistics logging is off, it is turned on and then off.

You must be logged in as the system administrator and you must be at the operating system's root directory to use the clearstats command. The *clearstats* command is used with *ssmcli* as follows:

`ssmcli [-p] -[v|q] [-h` *hostname*`] [-d` *device*`] clearstats`

B.9 The failback Command

The *failback* command updates the RAID controller firmware by failing back the fiber loop on the specified device.

You must be logged in as the system administrator and you must be at the operating system's root directory. The *failback* command is used with *ssmcli* as follows:

`ssmcli [-p] -[v|q] [-h` *hostname*`] [-d` *device*`] failback`

B.10 The firmware Command

The *firmware* update command updates the RAID controller firmware on the SP.

You must be logged in as the system administrator and you must be at the operating system's root directory. The *firmware* command is used with *ssmcli* as follows:

`ssmcli [-p] -[v|q] [-h` *hostname*`] [-d` *device*`] firmware` *filename* `-o -r -c`

filename	The **filename** argument is the full pathname of the new firmware image on the destination server. This image contains microcode that runs on the SP. It may also contain a microcode image destined for the SP's PROM, which runs power-up diagnostics. If the full pathname is not entered, the current working directory is prepended to the filename.
-c	The **-c** argument acts as a client switch. When -c is present, the firmware files reside on the client.
-o	The **-o** argument is a confirmation override switch. When -o is present, the user is not prompted for confirmation.
-r	The **-r** argument is a reboot switch. When -r is present, the storage system reboots after the firmware is downloaded; for example: `firmware /microcode.rev -o` would download the file microcode.rev, which is located in the root directory of the server. **Note:** Once the microcode has been downloaded, each SP in the cabinet reboots. The reboot takes several minutes.

B.11 The getagent Command

The *getagent* command gets device names and configurations and returns the names of devices controlled by the destination API along with descriptions about each device that the API is configured to communicate with. If you type the getagent command without switches, all of the values are displayed. If you enter the *getagent* command with switches, only the values for those switches will be displayed.

The *getagent* command is used with *ssmcli* as follows:

`ssmcli [-p] -[v|q] [-h` *hostname*`] [-d` *device*`] getagent` *-cabinet -desc -mem -model -name -node -peersig -prom -rev -scsiid -serial -sig -type -ver*

-cabinet	Cabinet type
-desc	ASCII description of the device
-mem	The size in Mbytes (MB) of the SP's memory. Each SP in the storage system must have the same amount of memory to make full use of the memory. Total memory can be 8, 16, 32, or 64 Mbytes.
-model	Model number of the device
-name	Name of the device
-node	Location in file system of device
-peersig	Signature of peer SP
-prom	The revision of the SP's PROM code. Each SP in the storage system will be running the same revision of PROM code. PROM code is updated automatically when you update an SP's LIC.

-rev	Microcode revision number and agent revision number. The revision of the Licensed Internal Code (LIC) that the SP is running. Each SP in the storage system will be running the same revision of LIC. You can update an SP's LIC.
-scsiid	SCSI ID of SP host connection. The SCSI ID number for the SP, which is determined by switch settings on the SP.
-serial	Serial number of the device
-sig	Unique SP signature
-type	Model type - deskside, rackmount
-ver	Version of the *ssmagent*

B.12 The getcache Command

This command provides cache environment information on the RAID SP's caching environment. The getcache command is used with the *ssmcli* as follows:

`ssmcli [-p] -[v|q] [-h` *hostname*`] [-d` *device*`] getcache [`*optional-switches*`]`

The *getcache* optional switches are:

-freebackend	Free memory available in the back end memory bank
-freecontrol	Free memory available in the control memory bank
-freefrontend	Free memory available in the front end memory bank
-pmbackend	Physical memory size for the back end memory bank.
-pmfrontend	Physical memory size for the front end memory bank.
-rm3a	The amount of memory, in Megabytes, reserved for the optimized RAID-3 data buffers on SP A.
-rm3b	The amount of memory, in Megabytes, reserved for the optimized RAID-3 data buffers on SP B.

B.12.1 Sample Output

```
SP Write Cache State Disabled
Cache Page Size: 2
Write Cache Mirrored: YES
Low Watermark: 0
High Watermark: 0
SPA Cache Pages: 0
SPB cache pages: 0
Unassigned cache pages: 0
Read Hit Ratio: 0
Prct Dirty Cache Pages: 0
Prct Cache Pages Owned
SPA Read Cache State Disabled
SPB Read Cache State Disabled
BBU Test Day: Sunday
```

```
BBU Test Time: 1:1
Physical Memory Size of CPU Ram = 3
Physical Memory Size of Front-End = 16
Physical Memory Size of Back-End = 16
SPA Read Cache Size = 0
SPB Read Cache Size = 0
SPA Write Cache Size = 0
SPB Write Cache Size = 0
SPA Optimized RAID 3 Memory Size = 6
SPB Optimized RAID 3 Memory Size = 6
Free Memory Size of CPU RAM = 2
Free Memory Size of Front-End = 10
Free Memory Size of Back-End = 9
```

Note: Some versions of flare do not support RAID-3 memory size value.

An SP's write cache is automatically enabled on power-on if the cache size is a valid number and the storage system has the following components:

B.13 The getcontrol Command

The *getcontrol* command returns information about the SP configuration. The command is used with *ssmcli* as follows:

`ssmcli [-p] -[v|q] [-h` *hostname*`] [-d` *device*`] getcontrol [`*optional-switches*`]`

where the *getcontrol* optional switches are:

-allib	Displays idle and busy ticks for all busses
-arv	The number of times a user request arrived while at least one other request is being processed
-avg	Average number of requests seen by SP
-busy	Percentage of time SP is busy
-cbt	Controller idle and busy ticks
-date	SP's system date
-day	SP's system day
-flush	Displays the number of times the High watermark, Low watermark, Idle unit, turned on flushing, This also displays the number of requests to flush write cache, and the number of write cache blocks flushed.
-he	Number of hard errors SP has seen
-ib1	Displays Internal bus 1 idle and busy ticks
-ib2	Displays Internal bus 2 idle and busy ticks
-ib3	Displays Internal bus 3 idle and busy ticks
-ib4	Displays Internal bus 4 idle and busy ticks
-ib5	Displays Internal bus 5 idle and busy ticks
-idle	Percentage of time SP is idle

-max	Maximum number of requests seen by SP
-read	Number of reads received by SP
-rw	Displays: *Host read requests:* The number of read requests made by the host to the LUN *Host write requests:* The number of write requests made by the host to the LUN *Host blocks read:* The number of blocks the host read from the LUN *Host blocks written:* The number of blocks written from the host to the LUN
-sc	Indicates whether system cache is on or off
-sf	Indicates whether system fault led is on or off
-sl	Indicates whether statistics logging is on or off
-time	SP's system time
-write	Number of writes received by SP

B.14 The getcrus Command

The *getcrus* command returns state information on all or selected CRUs (customer replaceable units) in the subsystem except for disks (see the *getdisk* command). The *getcrus* command displays the state of the following CRUs: fans, fiber cards (lcc), storage processors (SP), battery backup units (SPS), and power supplies (VSC).

getcrus is used with *ssmcli* as follows:

`ssmcli [-p] -[v|q] [-h` *hostname*`] [-d` *device*`]` getcrus [*optional-switches*]

You can use *getcrus* without any switches to get the state information for every CRU except disks. You can also use *getcrus* with the following switches to get the state information for specific CRUs:

-fana	Get the state information on fan A
-fanb	Get the state information on fan B. See -fana
-lcca	Get the state information on the fiber card for SP A
-lccb	Get the state information on the fiber card for SP B
-spa	Get the current operational state of SP A. It can be in any of these states: - Present: The SP is the communication channel you are using to communicate with the chassis. - Not Connected: Agent cannot talk to the SP because a communication channel specifying the SP is not in the agent's configuration file for the selected host. For example, the SP is connected to a different host than the SP in the communications channel for the chassis. - Not Present: SP that is in the communication channel to the selected chassis has failed or been removed. - Removed: SP was not present when the agent was started.
-spb	Get the state information on SP B. See -spa description.
-spsa	Get the state information on battery A

-spsb	Get the state information on battery B
-vsca	Get the state information on power supply A A 30-slot chassis has three VSCs: A, B, and C. A 20-slot chassis has two or three VSCs: A, B, and optionally C. A 10-slot chassis has one or two VSCs: A and optionally B. A 7-slot chassis has two VCS: A and B. If a chassis has the maximum number of VSCs, it can continue operating if one VSC fails, but shuts down immediately if a second VSC fails. If a VSC is Down, you can determine whether it has failed by seeing if its button is amber, looking at the event log for VSC failures, or seeing if the Replace light on the VSC is lit. If the VSC has failed, the system operator or service person should replace the VSC as soon as practical.
-vscb	Get the state information on power supply B

VSC state	Meaning
Present	Operating normally
Not Present	Failed or removed before the agent started running
Not Connected	This is for an SP, the agent cannot talk to the SP because a communication channel specifying the SP is not in the agent's configuration file for the selected host. For example, the SP is connected to a different host than the SP in the communications channel for the chassis. ailed or removed before the agent started running.
Removed	Not present when the agent was started

B.14.1 Sample Output

```
DPE Enclosure 0
Enclosure 0 FAN A State: Present
Enclosure 0 FAN B State: Present
Enclosure 0 Power A State: Present
Enclosure 0 Power B State: Empty
SPA State: Empty
Enclosure 0 SPS A State: Empty
Enclosure 0 SPS B State: Empty
Enclosure 0 LCC A State: Present
Enclosure 0 LCC B State: Present

DPE Enclosure 1
Enclosure 1 FAN A State: Removed
Enclosure 1 Power A State: Removed
Enclosure 1 Power B State: Removed
Enclosure 1 LCC A State: Removed
Enclosure 1 LCC B State: Removed
```

B.15 The getdisk Command

The *getdisk* command returns detailed status information about all or selected disks in the system.

getdisk is used with *ssmcli* as follows:

ssmcli [-p] -[v|q] [-h *hostname*] [-d *device*] getdisk *disk-position* [*optional switches*]

where the *getdisk* switches are:

disk-position	Argument is of the form e_d, where: e = enclosure that the disk is located on (1-10) d = disk number (0-9). If no disk-position is entered, *getdisk* returns a report for all disks in the system.
-bind	Indicates what percent of the disk is bound
-capacity	Indicates the disk capacity in blocks
-hr	Total number of hard read errors for all the disk modules in the LUN that persisted through all the retries. A growing number of hard errors may mean that one or more of the LUN's disk modules is nearing the end of its useful life.
-hs	Hot spare information (NO if not hot spare; bd otherwise)
-hw	Total number of hard write errors for all the disk modules in the LUN that persisted through all the retries. A growing number of hard errors may mean that one or more of the LUN's disk modules is nearing the end of its useful life.
-lun	The hexadecimal number identifying the LUN to which the disk module is bound. You specified this number in the Bind window when you bound the LUN. If not bound, returns "No."
-private	Private space of disk in blocks
-product	Product id of the disk
-rb	Indicates what percent of the disk is rebuilt
-read	Number of reads. Total number of read requests made to the disk module. You might find the LUN read information in getlun more useful because it is for the entire LUN, and not just for one of the disk modules in the LUN.
-rr	Number of read retries. Total number of times read and write requests to the disk module were retried.
-rs	Number of remapped sectors. Total number of disk sectors that were faulty when written to, and thus were remapped to a different part of the disk module.
-state	State of disk
-serial	Serial number of the disk
-sig	Bind signature (assigned during bind)

-sr	Total number of soft read errors for all the disk modules in the LUN that disappeared before all the retries. A growing number of soft errors may indicate that one of the LUN's disk modules is nearing the end of its useful life.
-sw	Total number of soft write errors for all the disk modules in the LUN that disappeared before all the retries. A growing number of soft errors may indicate that one of the LUN's disk modules is nearing the end of its useful life.
-vendor	Vendor id of the disk
-wr	Number of write retries. Total number of times read and write requests to the disk module were retried.
-write	Number of writes. Total number of write requests made to the disk module. You might find the LUN write information in getlun more useful because it is for the entire LUN, and not just for one of the disk modules in the LUN.

The *getdisk 0_0* command returns information for disk 1 in disk array enclosure 1 (DAE 1).

B.15.1 Sample Output

A sample output might look like:

```
Enclosure 0 Disk 0
Vendor Id: SEAGATE
Product Id: ST19171FC
Product Revision: CC46
Lun: 3
State: Ready
Hot Spare: NO
Prct Rebuilt: 0
Prct Bound: 0
Serial Number: LA142784
Capacity: 8429
Private: 184320
Bind Signature: 9337
Hard Read Errors: 0
Hard Write Errors 0
Soft Read Errors 0
Soft Write Errors 0
Read Retries 0
Write Retries 0
Remapped Sectors 0
Number of Reads 0
Number of Writes 0
```

B.15.2 Disk Module Information

A disk module can be in any of these states:

State	Meaning
Binding	Being bound into a LUN
Enabled	Either a hot spare on standby or part of a bound LUN that is assigned to (owned by) the SP you are using as the communication channel to the chassis. If the storage system has another SP, this module's status is Ready when you use the other SP as the communication channel to the chassis.
Equalizing	Data from a hot spare is being copied onto a replacement disk module
Failed	Powered down or inaccessible
Formatting	Being hardware formatted. Generally, modules do not need formatting.
Off	Powered off by the SP, which can happen if a wrong size module is inserted
Powering Up	Power is being applied to the disk module
Ready	Module is part of a broken LUN or a LUN that is bound and unassigned. This can mean that the disk module is part of a LUN that is not owned by the SP that you are using as the communication channel to the chassis. If the disk module is part of a LUN assigned to an SP other than the one you are using as the communication chassis, the module's status is either Enabled or Ready. It is Enabled when you use the other SP as the communication channel to the chassis.
Rebuilding	Module is either a hot spare or replacement disk module that replaced a failed module in a LUN. The data is being rebuilt on a hot spare or a replacement disk module.
Removed	Removed from the chassis; applies only to a disk module that is part of a LUN
Standby	Module is either a hot spare or replacement disk module that replaced a failed module in a LUN. The data is being rebuilt on a hot spare or replacement disk module
Synchronizing	Write cache on the disk module is being synchronized
Unbound	Ready to be bound into a LUN

B.16 The getlog Command

The *getlog* command returns the entire SP log, the newest entries to the log, or the oldest entries to the log.

Note: Depending on the size of the unsolicited log, the *getlog* command may take several minutes to execute.

getlog is used with *ssmcli* as follows:

ssmcli [-p] -[v|q] [-h *hostname*] [-d *device*] getlog [+-]*N*

where the *getlog* switches are:

+N	Display the oldest N entries in the log, with the oldest entry first. If there are fewer than N entries, the entire log is displayed. N must be greater than 0.
-N	Display the newest N entries in the log, with the oldest entry first. If there are fewer than N entries, the entire log is displayed. N must be greater than 0.
NL	If a range of entries is not specified, the entire log is displayed with the oldest entry first.

For example:

```
getlog -5 o
```

generates the following sample output:

```
04/10/97 12:43:20 Enclosure 0 Disk 0 (938) (Non R3 LUN-Can't Assign BW)
[0x03]   0    0
04/10/97 12:43:20 Enclosure 0 Disk 2 (938) (Non R3 LUN-Can't Assign BW)
[0x01]   0    0
04/10/97 12:43:19 SP B (6e5) (FE Fibre loop operational)
[0x00]   0    0
04/10/97 12:43:19 SP B (6e4) (FE Fibre loop down)
[0x03]    0   0
04/10/97 12:43:01 SP B (6e5) (FE Fibre loop operational)
[0x00]   0    0
```

The above example shows the most recent entries in the log, which show a CRU that has failed be replaced, rebuilt, and become enabled.

B.17 The getloop Command

The *getloop* command returns information about the fiber loop.

getloop is used with *ssmcli* as follows:

ssmcli [-p] -[v|q] [-h *hostname*] [-d *device*] getloop

The following is a sample output:

```
Needs Discovery: NO
Data Loop Failover: NO
Command Loop Failover: NO
Can Peer SP Failover: YES
Manual Discovery Enabled: NO
Can This SP Failover: NO
Illegal Cross Loop: NO
Connection #0: Enclosure 0
```

B.18 The getlun Command

The *getlun* command returns information about a LUN and the CRUs which make up that LUN. If statistics logging is disabled, some fields are not printed (see switches below).

getlun is used with *ssmcli* as follows:

`ssmcli [-p] -[v|q] [-h` *hostname*`] [-d` *device*`] getlun` *lun-number* `[`*optional flags*`]`

where the *getlun* arguments are:

lun-number	The number of the LUN to get information about. This switch is mandatory. All others are optional.
-aa	Get whether auto assignment is enabled or disabled
-at	Get whether auto-trespass is enabled or disabled
-bind	Get the percent of the LUN that is bound
-bp	Get the number of blocks prefetched by read cache
-bread	Get the number of blocks read for each CRU
-brw	Get the number of host block read and write requests
-busy	Get the percentage of time that the disk is busy
-bwrite	Get the number of blocks written for each CRU
-capacity	Get the LUN capacity in megabytes
-crus	Get the names and states of disks in this LUN
-default	Get whether or not the SP is default owner of LUN
-disk	Get the disk statistics for the LUN. Disk statistics switches are only valid when the -disk switch is included.
-ena	Get whether the disk is enabled. You must also include the -disk switch.
-ff	Get the number of times that a write had to flush a page to make room in the cache
-idle	Get the percentage of time that the LUN is idle
-idt	Get the idle delay time for caching (in 100 ms intervals)
-it	Get the idle threshold value for caching
-mp	For variable length prefetching, get the absolute maximum number of disk blocks to prefetch in response to a read request
-nminus	Get whether or not NMinus1Go (also called minimal-latency reads) is supported
-owner	Get which SP currently owns the LUN, SPA or SPB
-pd	Get the prefetch disable, which is the size in blocks of the smallest read request for which prefetching is disabled. Prefetching is disabled for all read requests whose size is equal or larger than the Prefetch Disable size.
-pic	Get the maximum number of I/Os that can be outstanding to the unit and have the unit still perform prefetching
-prb	Get the percentage of the LUN that is rebuilt

-prefetch	Get the prefetch type. The type of prefetching that the LUN is using for read caching.
-prf	Get the percentage of read requests that caused a flush
-ps	Get the prefetch size
-psm	Get the prefetch size/multiplier. The amount of data prefetched for one host read request. For constant length prefetching, gets the read request in blocks. For variable length prefetching, gets the multiplier used to calculate the amount of data to prefetch. For example, if the prefetch size is 8, the amount of data to prefetch is 8 times the amount of data requested.
-qavg	Get the average queue depth for each CRU
-qmax	Get the maximum queue depth for each CRU
-rb	Get the maximum time to rebuild failed disk (in hours)
-rc	Get the LUN's read cache state (enabled or disabled)
-rchm	Get the read and write cache hits and misses
-reads	Get the number of reads received for each CRU
-remap	Get the number of remapped sectors
-ret	Get the retention priority. The priority for retaining prefetched data when the read cache becomes full. If the priority is Equal Priority, prefetched data has priority over host requested data. If it is Equal, both types of data have the same priority.
-rh	Get read hit info if statistics logging is enabled
-rhist	Get the Read Histogram, which is an array of 10 locations that contain the number of reads. Element n of the array contains the number of reads that were larger than or equal to 2n-1 and less than 2 n blocks in size, e.g. the first bucket contains 1 block reads, the second contains 2-3 block, the third 4-7 and so on. rhist also displays read histogram overflows which are the number of I/O operations that were larger than 512 blocks.
-rr	Get the number of read retries. Total number of times read requests to all the disk modules in the LUN were retried.
-rwf	Get the percent of write requests that caused a flush
-rwr	Get the host read and write requests
-service	Get the average service time for each CRU
-srcp	Get the read cache configuration
-ps	Get the segment size
-ssm	Get the prefetch segment size/multiplier. The size of the segments that make up a prefetch operation. For constant length prefetching, -ssm gets the segment size (in blocks) of data prefetched in one read operation from the LUN. For variable length prefetching, -ssm gets the multiplier which determines the amount of data, relative to the amount of requested data, prefetched in one read operation. For example, if the segment multiplier is 4, the segment size is 4 times the amount of data requested.
-stripe	Get the stripe size (in blocks). Number of sectors that the storage system can read or write to a single disk module in the LUN. The default size is 128 sectors. This size was specified when the LUN was bound. Stripe element size does not apply to a RAID 1 LUN, individual unit, or hot spare.

-type	Get the RAID type of the LUN as follows: RAID 0 - nonredundant individual access array RAID 1 - mirrored pair RAID 1/0 - mirrored RAID 0 group RAID 3 - parallel access array RAID 5 - individual access array DISK - individual unit HOT SPARE - hot spare
-upb	Get the number of prefetched blocks not used by the read cache
-verify	Get the maximum verify time in hours
-was	Get the minimum block size request that bypasses cache
-wc	Get the write cache state. The write cache is enabled or disabled
-wch	Get the number of times that a write was completely satisfied by cache, avoiding a disk operation
-wh	Get write hit information, if statistics logging is enabled
-whist	Get the write histogram which is an array of 10 locations that contain the number of writes. Element n of the array contains the number of writes that were larger than or equal to 2n-1 and less than 2 n blocks in size, e.g. the first bucket contains 1 block writes, the second contains 2-3 block, the third 4-7 and so on. whist also displays the write histogram overflows which are the number of I/O that were larger than 512 blocks.
-wr	Get the number of write retries. Total number of times write requests to all the disk modules in the LUN were retried.
-writes	Get the number of writes received for each CRU

For example:

```
getlun 0 o
```

This command retrieves information about LUN 0.

This is a sample output:

```
Prefetch size (blocks) = 4
Prefetch multiplier = 4
Segment size (blocks) = 4
Segment multiplier = 4
Maximum prefetch (blocks) = 512
Prefetch Disable Size (blocks) = 129
Prefetch count = 40

Variable length prefetching YES

Prefetched data retained YES
```

```
Max Verify Time Hours
Read Histogram[0] 0
Read Histogram[1] 0
Read Histogram[2] 0
Read Histogram[3] 0
Read Histogram[4] 0
Read Histogram[5] 0
Read Histogram[6] 0
Read Histogram[7] 0
Read Histogram[8] 0
Read Histogram[9] 0
Read Histogram overflows 0

Write Histogram[0] 0
Write Histogram[1] 0
Write Histogram[2] 0
Write Histogram[3] 0
Write Histogram[4] 0
Write Histogram[5] 0
Write Histogram[6] 0
Write Histogram[7] 0
Write Histogram[8] 0
Write Histogram[9] 0
Write Histogram overflows 0

Read Requests 0
Write Requests 0
Blocks read 0
Blocks written 0
Read cache hits 0
Prefetched blocks 0
Unused prefetched blocks 0
Write cache hits 0
Forced flushes 0
Read Hit Ratio 0
Write Hit Ratio 0

RAID Type: Disk
Stripe Crossing: 0
Current owner:
Auto-trespass: DISABLED
Auto-assign: DISABLED
Write cache: ENABLED
Read cache: ENABLED
Idle Threshold: 0
Idle Delay Time: 20
Write Aside Size: 2048
Default Owner: SPB
Rebuild Time: 0
Read Hit Ratio: 0
Write Hit Ratio: 0
Prct Reads Forced Flushed: 0
Prct Writes Forced Flushed: 0
Prct Rebuilt: 0
Prct Bound: 0
Lun Capacity: 8429
```

```
Enclosure 0 Disk 3 Ready
Reads: 0
Writes: 0
Blocks Read: 0
Blocks Written: 0
Queue Max: 0
Queue Ave:0
Avg Service Time: 0
Prct Idle Not Available
Prct Busy Not Available
Remapped Sectors: 0
Read Retries: 0
Write Retries: 0
```

B.19 The setcache Command

The *setcache* command changes system caching on the RAID SP. Using *setcache* with its various switches, you can enable, disable, and configure the cache.

You must be logged in as the system administrator and you must be at the operating system's root directory.

setcache is used with *ssmcli* as follows:

`ssmcli [-p] -[v|q] [-h` *hostname*`] [-d` *device*`] setcache` *enable-disable* `-h` *hi-watermark* `-l` *low-watermark* `-m` *mirror* `-p` *page-size* `-r3a` *raid3-size-spa* `-r3b` *raid3-size-spb* `-rca` *read-cache-spa* `-rcb` *read-cache-spb* `-rsza` *read-size-spa* `-rszb` *read-size-spb* `-wc` *write-cache* `-wsza` *write-size-spa* `-wszb` *write-size-spb*

where the *setcache* arguments are:

enable-disable	When enable-disable is used with other options: 1 = enable system write cache 0 = disable system write cache When enable-disable is the only option: 1 = enables system read and write cache 0 = disables system read and write cache
hi-watermark	The high watermark is the percentage of dirty pages, which when reached, causes the SP to begin flushing the cache. Valid hi-watermark values are 0-100 and must be greater than the low-watermark.
low-watermark	The percentage of cache dirty pages that determines when cache flushing stops. When the write cache low watermark is reached during a flush operation, the SP stops flushing the cache. Valid low-watermark values are 0-100. The low watermark is set to 50 regardless of enable-disable.
mirror	Enable mirroring. The default is enable. 1 = enable mirroring 0 = disable mirroring
page-size	Size (in kilobytes) of pages to partition cache into. Valid sizes are 2, 4, 8, and 16. The default is 2. page size is set to 2 regardless of enable/disable. As a general guideline, the cache page size should be 8 KB for general UNIX file server applications, and 2 KB or 4 KB for database applications.

raid3-size-spa	Set the optimized RAID 3 size for SP A. This size must be the sum of the memory allocated for all RAID 3 LUNs to be bound. For example, if you bind two RAID 3 LUNs, each with the recommended 6 Mbytes of memory, you must set this field to 12 (6 + 6). For any other storage systems, set this field to 0, and do not change it.
raid3-size-spb	Set the optimized RAID 3 size for SP B. This size must be the sum of the memory allocated for all RAID 3 LUNs to be bound. For example, if you bind two RAID 3 LUNs, each with the recommended 6 Mbytes of memory, you must set this field to 12 (6 + 6). For any other storage systems, set this field to 0, and do not change it.
read-cache-size-spa	Set the read cache size for SP A
read-cache-size-spb	Set the read cache size for SP B
read-cache-spa	Set the read cache state for SP A 0 = read cache disabled 1 = read cache enabled
read-cache-spb	Set the read cache state for SP B 0 = read cache disabled 1 = read cache enabled
write-cache	Set the write cache state 0 = write cache disabled 1 = write cache enabled
write-cache-size-spa	Set the write cache size for SP A
write-cache-size-spb	Set the write cache size for SP B

Note: The sum of an SP's Write Cache size, and Optimized RAID-3 Memory size must be less than or equal to the SP's Front End Size. The sum of an SP's Read Cache size, and Optimized RAID-3 Memory size must be less than or equal to the SP 's Back End Size. RAID-3 memory must be the same size in the front and back ends.

This is an example of using setcache:

```
setcache -wc 1 -p 8 -l 50 -h 75 o
```

This example enables the system write cache partitioned into 8K pages, with a 50% low watermark value and a 75% high watermark value.

The *setcache* command generates no output

B.20 The setstats Command

The *setstats* command sets statistics logging. If no optional switches are present, the command returns the current state of Statistics logging. The user must specify whether or not to enable or disable the cache.

The SP maintains a log of statistics for the LUNs, disk modules, and storage system caching that you can turn on and off. When the log is on it affects storage system performance, so you may want to leave it off unless you have a reason to monitor performance.

The log uses a 32-bit counter to maintain the statistics numbers, so when the counter is full, the statistics numbers start over at zero. As a result, you will see a sudden decrease in a statistics number if you view it shortly before the counter is full and shortly after the counter restarts at zero. If you want to keep the log turned on for more than two weeks, we recommend that you reset the log about every two weeks, so you know when the numbers start at zero.

You must be logged in as the system administrator and you must be at the operating system's root directory.

setstats is used with *ssmcli* as follows:

`ssmcli [-p] -[v|q] [-h` *hostname*`] [-d` *device*`] setstats -on -off`

where the *setstats* arguments are:

-on	Set statistics logging on
-off	Set statistics logging off

Note: If statistics logging is off, certain options in other commands are disabled. The options for the other commands specify whether statistics logging must be on for their options to be valid.

This is an example of using setstats:

```
setstats -on
```

This is the output of the example:

```
Statistics logging is enabled.
```

B.21 The unbind Command

The *unbind* command deconfigures LUNs from their current logical configuration. unbind destroys all data on the LUN, therefore when you execute an unbind command, you are prompted to verify that you really want to unbind the LUN. If you do not want the prompt to appear, you can override prompting using the **-o** switch.

Storage systems typically ship from the factory with each disk module bound as an individual unit. In such a situation, none of the disk module buttons have select buttons beside them, and you need to unbind LUNs to create the LUNs you want.

Caution: Never unbind the last LUN because doing so may prevent you from accessing the storage system. Never reboot a server unless at least one LUN is bound. If you unbind the last LUN and lose access to the storage system, refer to the release notice for the agent media for information on what to do.

You must be logged in as the system administrator and you must be at the operating system's root directory.

unbind is used with *ssmcli* as follows:

ssmcli [-p] -[v|q] [-h *hostname*] [-d *device*] unbind *lun-number* [-o]

where the *unbind* arguments are:

lun-number	Logical unit to deconfigure
-o	Do not prompt the user for permission

This is an example of using setstats:

```
unbind 3 -o
```

This command destroys LUN 3 without prompting you for permission and frees its physical disks to be reconfigured.

B.22 General Fibre Channel Performance Tips

Several factors affect performance on an FC-AL:

- frame overhead, which is fairly fixed and includes headers, CRC, and more
- frame size: increasing frame size means that a larger percentage of the packet (overhead + payload) is used for data. Frame sizes of 256, 512, 1024, and 2048 bytes are supported. Maximum frame size is not always better because buffer credit is directly related to frame size: cutting frame size can double buffer credit
- I/O size: use large I/O request sizes (64 KB or more) for good bandwidth performance
- buffer credit: when buffer credit goes to zero, transfers on the loop stop. Therefore both the host and the targets must provide enough buffers and credit
- loop length effect on performance is due to propagation delays
- number of nodes has a great effect on performance, since each node adds about a 6-word (38 ns each) delay
- initiator behavior: the host chip should be well behaved, for example, it should not hold onto the FC-AL loop while executing firmware; the host chip should also be unfair when arbitrating for the loop when it needs to send a command to a device
- target device behavior: disk drives should be well behaved, for example, it should not hold onto the FC-AL loop while executing firmware; targets should also avoid sending out-of-order frames or frames with a non-zero relative offset

Appendix C

Graphical User Interface Administration

The Fibre Channel RAID drives and controllers can be monitored and administered using a graphical user interface (GUI) program. This program is included with the CD that ships with every Silicon Graphics Fibre Channel RAID mass storage system.

A graphical user interface (GUI) is available for all Onyx2 systems. The GUI can also be used on Origin2000 systems with optional visualization console (SI) graphics.

This chapter discuses the basic operations of the Fibre Channel RAID GUI.

Note: There is no GUI support for non-RAID fibre disk arrays.

You can rlogin to a server and run the GUI from a remote graphics workstation. To do this, you must be an authorized user. A user is authorized if the ssmagent on the server contains an entry for the user in */etc/config/ssmagent.config*.

C.1 The ssmagent Interface

If you need to confirm that the Fibre Channel RAID ssmagent is installed on the host system, type the following at the command line:

```
# chkconfig -f ssm on
```

If the GUI is not available on the host, install it using the information in the release notes and the CD included with the Fibre Channel RAID option. Any user can run an ssmagent session from any management station on which it is installed, and can use it to monitor the storage system. Only authorized users can use ssmagent to configure or reconfigure a storage system. A user is authorized if the ssmagent agent on the server contains an entry for the user.

Caution: The agent allows more than one ssmagent session to access the same storage system at the same time. You must make sure that two authorized users do not use ssmagent to configure or reconfigure the same storage-system chassis at the same time.

C.1.1 Launching the Fibre Channel RAID GUI

Before *ssmagent* opens its Storage Management window, it must know which servers have storage systems that you want to manage during the session. You can specify these hosts in one of the following ways:

- Put the fibre host system names in an ssmhost file in your home directory.
- Use the *ssmagent*'s Host Administration dialog box, which ssmagent opens if the previous two options are not available.

To create your hosts ssmagent config file:

1. Copy the content of the */etc/config/ssmagent.config.proto* file to */etc/config/ssmagent.config*.
2. Edit the new file using your favorite line editor and enter the valid device and host names that apply to your local host system.
3. Save the file.

Note that any changes to the file that are made after you start the agent require that you stop and then restart the agent before they take affect.

Launch the Fibre Channel RAID agent by entering

```
/etc/init.d/ssm start
```

To stop the Fibre Channel RAID agent enter

```
/etc/init.d/ssm stop
```

C.1.2 Configuring the Path to ssm Executables in .cshrc

The following steps configure the host system's *.cshrc* file to execute *ssmgui* commands from any directory on the host. If this path is not set, you must be in the */usr/ssm/lib* directory and enter the complete path (*/usr/ssm/bin/ssmgui*) each time you want to launch *ssmgui*.

Use the following information to set the proper path on the host system:

1. As the superuser, use your favorite line editor to open the *.cshrc* file.
2. Enter the following in the *.cshrc* file: `set path = ($path /usr/ssm/bin)`
3. Write and quit the *.cshrc* file.
4. Enter the following command sequence to reinitialize the *.cshrc* file.

   ```
   # source .cshrc
   ```

   Launch the Fibre Channel RAID GUI session by entering the following from any direcetory

   ```
   # ssmgui &
   ```

5. Click on OK after the Fibre Channel RAID SSM startup screen appears. At this point an inquiry window appears and asks if you "want to search for a host file specified by the environment variable." See Figure C-1 for an example of the inquiry window.

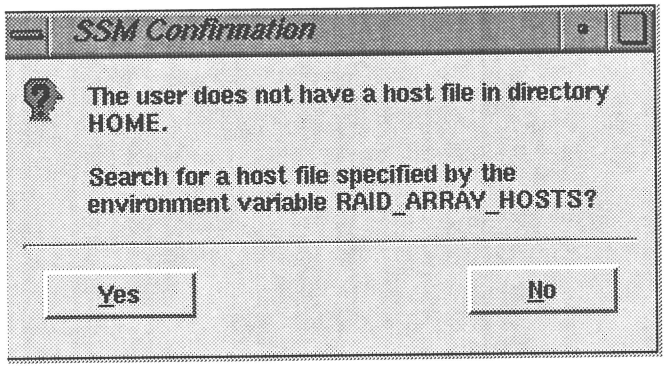

Figure C-1 Hosts Inquiry Box

C.1.3 Host Administration Selection

If you want to set up an alternate *ssmhosts* file to use as the source of the RAID array host variables, go to the same command line you launched the agent from and enter something similar to `setenv RAID_ARRAY_HOSTS /ssmhosts`

If you click No in the inquiry box, the Host Administration window appears, see Figure C-2. You can use this window to select or confirm the fibre channel host system that you wish to administer.

The Host Identifier sub-window (Figure C-2) shows the current host system name(s) that the administration GUI can access and control. The Host Administration dialog box opens if no server host names are in the ssmagent file or the variable *RAID_ARRAY_HOSTS*.

In this situation, the list contains the host names of the systems with currently managed arrays. You can add or remove system host names from this list. Host systems on this list are called managed systems.

To add a system to the list, follow these steps:

1. In the Host Identifier field, enter the system host name.
2. Click the Add button. The host name appears in the list of systems below the field.
3. When the list contains all the system host names you want, click the OK button. Click Cancel if you wish to return to the GUI entry window.

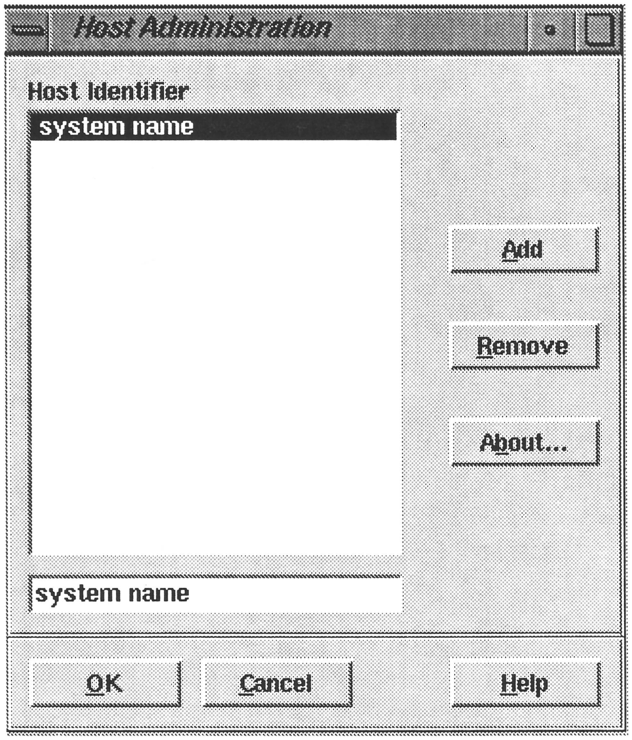

Figure C-2 Host Identifier Window Example

A popup opens asking if you want to save the new host list to the ssmagent file in your home directory. You should save the list if it contains only systems you want to manage in most future ssmagent sessions. You should not save the list if you added systems that you want to manage *only* in the current ssmagent session.

4. In the popup, click the Yes button to save the new system list to the *ssmagent* file; otherwise, click the No button.

Note: If you save the host list, it is always written to the ssmagent file and never to the variable *RAID_ARRAY_HOSTS*. The ssmagent never changes the variable *RAID_ARRAY_HOSTS*.

If you add the wrong name, you can remove it from the list as described in the procedure that follows, and then add the correct name.

The *ssmagent* creates an icon for each array that it finds connected to each system whose host name you added to the list, provided that the icon does not already exist. The *ssmagent* displays these icons in the array selection area of its Storage System Manager window. Since *ssmagent* creates only one icon for each array, only one icon is displayed for an array connected to two different host systems.

Note: If the ssmagent cannot communicate with a host system you added to the list, it displays an error message.

To remove a server from the list, follow these steps:

1. Either enter the server host name to remove in the Host Identifier field or select the host name from the list. The selected host name appears in Host Identifier field.
2. Click the Remove button. A confirmation popup opens.
3. In the popup, click the Yes button to confirm the removal.

C.1.4 The Storage System Manager Window

Table C-1 identifies and describes the administrative icons in the Storage System Manager GUI window.

Table C-1 Storage System Manager Configuration Icon Descriptions

Icon	Functional Description
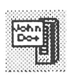	The "John Doe" icon allows you to name a selected array. Enter an array name in the dialog box or click Cancel to exit.
	This icon displays the main monitoring dialog box for a selected array. You can perform several monitoring functions that are available in the Storage System Manager window. Flashing a specific enclosure's LEDs and stopping the flashing are two functions unique to this dialog box.
	The icon displays the Array Configuration dialog box for each selected array in the window. This window and its icon functions are covered in Figure C-6 and Table C-2.
	Displays the Array Information window for each selected array in the window. This window provides operation, polling priority, disk formatting, and host connection information.
	The SPA icon provides event log information for a selected array. If the enclosure has no SPA, the information box tells you there are no events in the log for SPA.
	Clicking the SPB icon provides event log information for a selected array. If the enclosure has no SPB, the information box tells you there are no events in the log for SPB.
	The manual polling icon allows you to poll the selected arrays.
	This icon disables automatic polling for the selected arrays in the window.

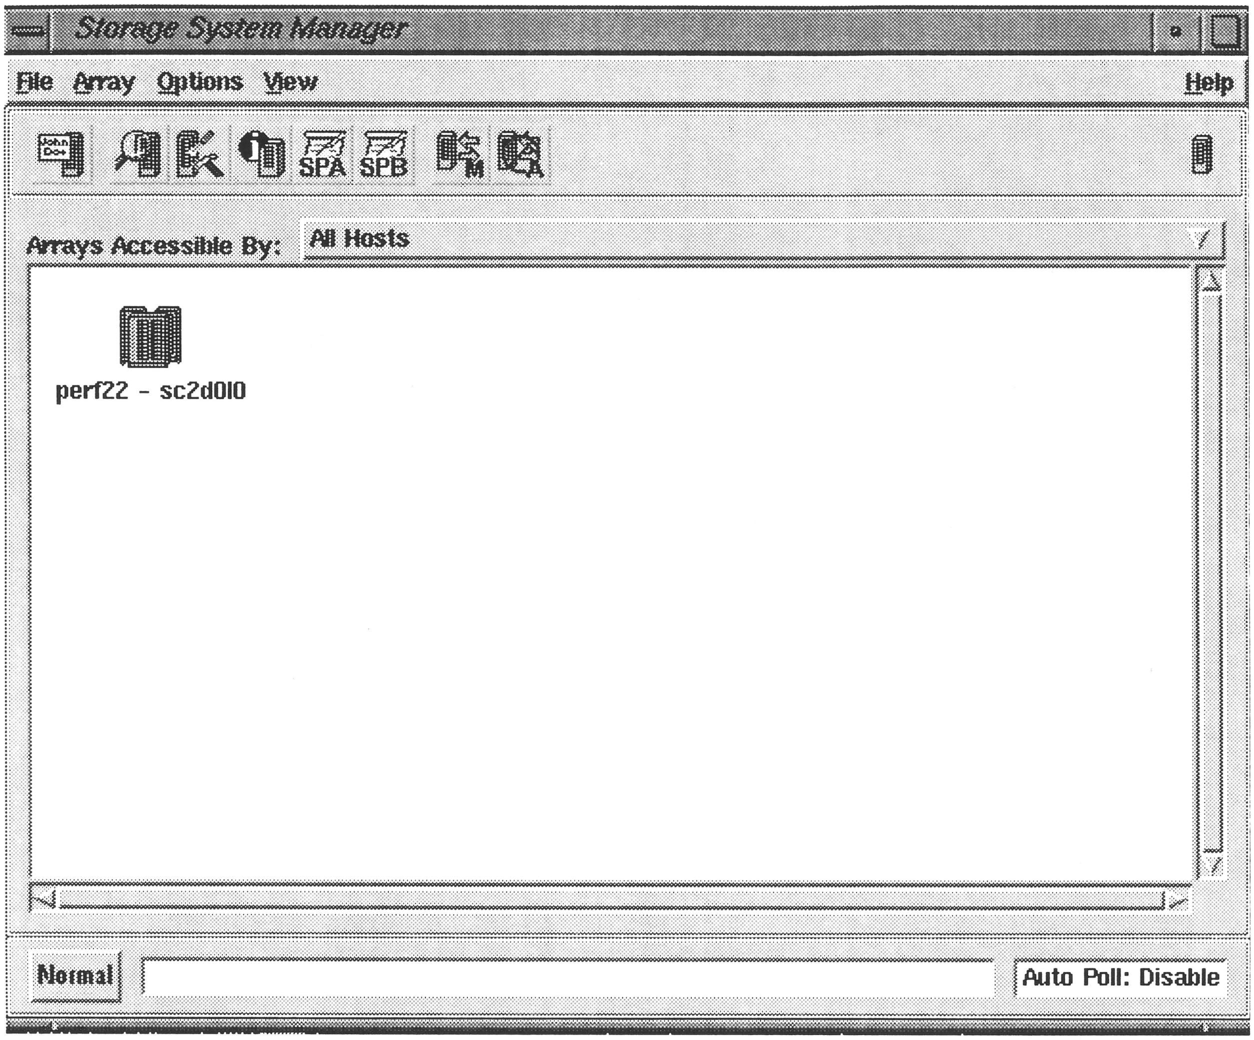

Figure C-3 Main GUI Administration Window

The array selection area is where you select the arrays on which you want to perform operations. It contains an array icon for the arrays specified by the array selection filter. The icon for an array varies depending on whether the array is accessible, inaccessible, or unsupported. If the ssmgui can communicate with an array, it is accessible. The icon for an accessible array consists of a graphic that represents the array type and the array name.

C.1.5 Checking and Saving System Event Messages

Click the Event Log icon to check on SPA or SPB errors seen by the administrative GUI system. This event log is similar in function to the error log on a System Controller. It holds a finite number of messages. Unless the user writes these to a separate file, the oldest messages are deleted after it reaches capacity. Figure C-4 shows an example of the Event Log window for a RAID array controlled by storage processor A (SPA).

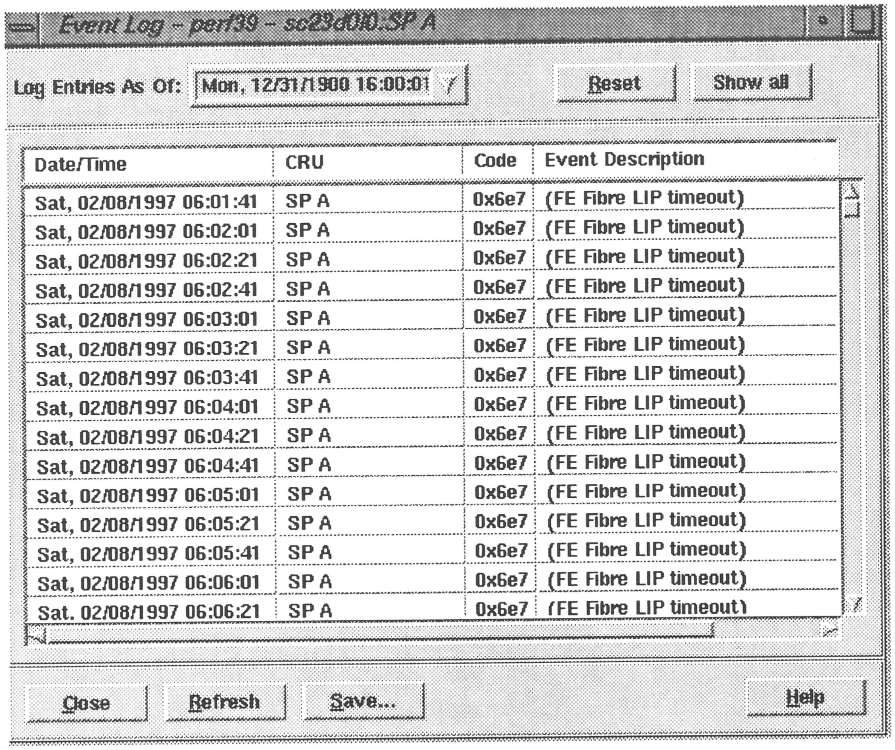

Figure C-4 System Event Log Window Example

To keep a record of all the current fibre related status messages in the *Event Log*, click on the Save button at the bottom of the Event Log window. The Save Log Events window comes up (see Figure C-5).

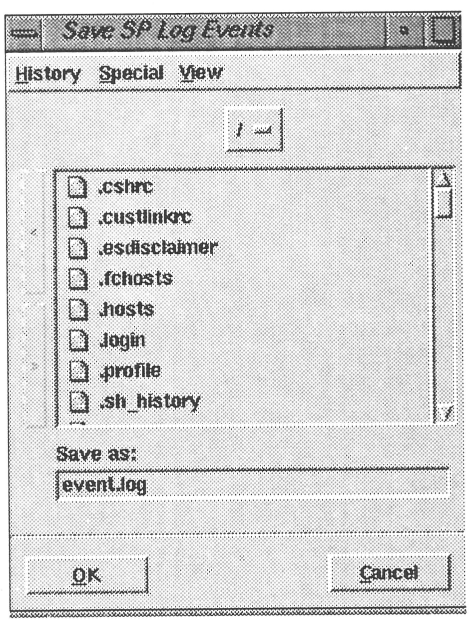

Figure C-5 Saving the System Events Log to Another File

To save all the Event Log information to a new file:

1. Select the directory and type in the designated file name for the information.
2. Click OK, and the contents of the *Event Log* file are copied to the new location and filename.
3. Click the Cancel button in the Save Log Events window.
4. Click the Close button on the Event Log window when you are finished reviewing the event messages.

You should be back at the main GUI administration interface window.

C.2 Fibre Channel RAID Array Configuration Options

When you click on a displayed RAID array icon in the GUI administration interface window (refer to Figure C-3), you access an Array Configuration administration window similar to Figure C-6. This window offers graphical display of LUNs that are owned by SPA, SPB, or that are unowned.

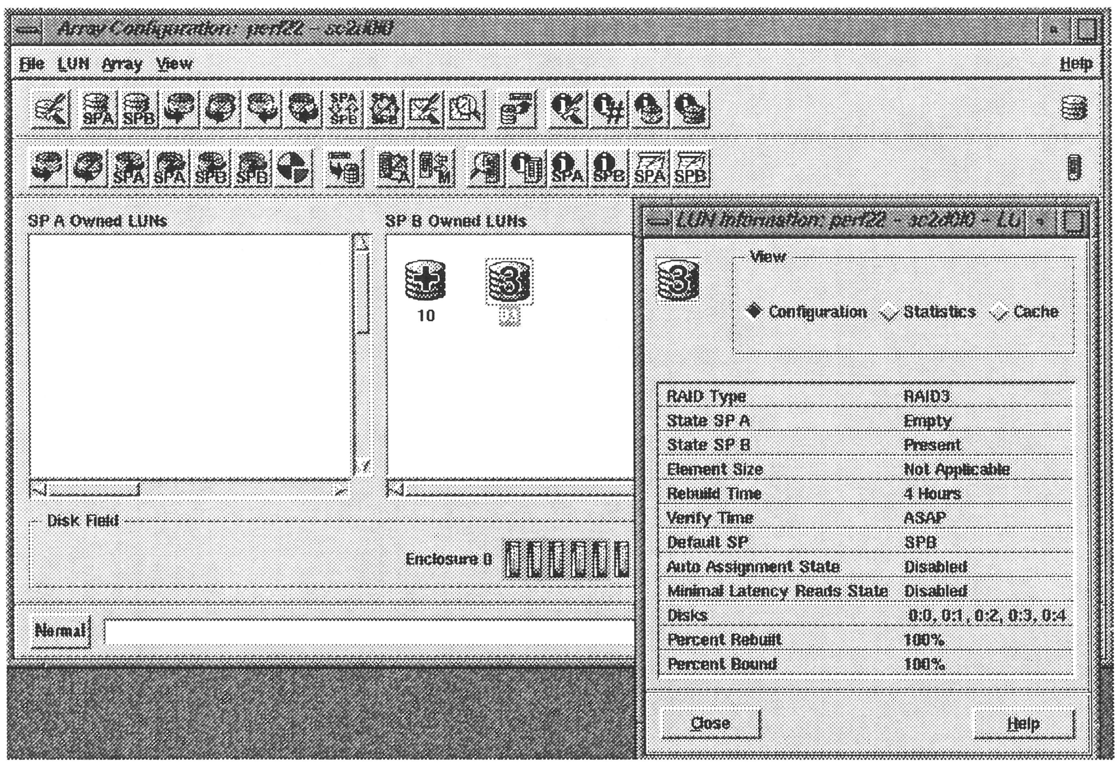

Figure C-6 Fibre Channel RAID LUNs and Arrays Accessible from the Host

This window provides multiple optional administration interface tools and information windows. Your particular Fibre Channel RAID configuration may not have all the components that appear in the Array Configuration window.

Note: The icons for administering particular functions appear regardless of weather your system has that functional hardware installed.

The administrative icons in the Fibre Channel RAID Array Configuration window are divided into two groups on two separate lines:

- The top line of icons is concerned specifically with disk array administration. An icon of vertically stacked disks appears to the far right of the icons.
- The bottom line of icons is mainly concerned with enclosure functionality, and an icon of a fibre enclosure is displayed at the far right of the displayed icons.

The example window (Figure C-6) also shows an icon of the number 3 over vertically stacked disks in the center sub-window. This icon represents a RAID-3 array. The logical unit number (LUN) information window (at the right side of Figure C-6) pops up to provide information and options regarding the LUN when you click the disk stack icon.

Table C-2 provides a functional summary of each of the administration icons in the Array Configuration window.

Table C-2 Array Configuration Icon Descriptions

Icon	Functional Description
	This icon can change the bind parameters for a selected LUN. Options include rebuild times and default storage processor selection.
	The storage processor A (SPA) icon button allows you to set the default SP to SPA ownership for particular LUNs.
	The SPB icon allows you to set the default SP ownership for a particular LUN to SPB.
	This icon enables write caching for particular LUNs.
	This icon allows you to disable write caching for particular LUNs.
	This icon enables read caching for particular LUNs.
	This icon allows you to disable read caching for a selected LUN.

Table C-2 (continued) Array Configuration Icon Descriptions

Icon	Functional Description
	This icon enables auto assignment for the selected LUN.
	This icon disables auto assignment for a selected LUN.
	Clicking this icon allows changing of the rebuild time for a selected LUN. You enter the rebuild time in hours.
	Clicking this icon gives you a query box to enter and change the verify time for a selected LUN in hours.
	This icon allows you unbind a selected LUN. Note that unbinding a LUN causes loss of all data stored on the LUN.
	Clicking this icon displays selected LUN information. This window automatically displays configuration information for each selected LUN. Statistics and cache information are also accessible from this window.
	This icon automatically displays selected LUN statistics. Configuration and cache information are also accessible from this window.
	Displays cache information for a selected LUN as the default. Configuration and statistics information are also accessible from this window.
	Displays a fault description for a selected LUN. The information window reports none of the selected LUNs are faulted when operation is normal.
	This icon permits you to enable the write cache for a selected array.
	This icon allows you disable the write cache for a selected array. The information box indicates if the write cache is already disabled or being disabled for a selected array.
	Enable the read cache for SPA in the selected array.
	Disable the read cache for SPA in the selected array.
	Enable the read cache for SPB in the selected array.

Table C-2 (continued) Array Configuration Icon Descriptions

Icon	Functional Description
	Disable the read cache for SPB in the selected array.
	This icon allows you to partition the array's memory. Each SP is represented for a selected array. If one SP is not present, the box is grayed out and not accessible.
	Clicking this icon displays the bind LUN dialog box. Options include binding and unbinding disks. Bind parameters for a selected SP are adjustable using this box.
	Allows you to disable automatic polling for the selected array.
	Allows you manually poll the array selected.
	Displays the main monitoring dialog box for the selected array. Identifies specific enclosures in the array.
	Displays the Array Information dialog box.
	Clicking this icon displays the array information dialog box for storage processor A.
	Clicking this icon displays the array information dialog box for storage processor B.
	Displays the event log window for storage processor A. If SPA is not present, an information window tells you there are no events in the log.
	Displays the event log window for storage processor B. If SPB is not present, an information window tells you there are no events in the log.

C.3 Configuring an Array

To configure an array initially, you perform the following tasks:

1. Set up array memory to create RAID-3 LUNs.
2. Create LUNs.
3. After initial configuration, you may need to perform one or more of the following tasks:
   - Change array caching parameters
   - Upgrade an array for caching
   - Update Licensed Internal Code and PROM code.

C.3.1 Setting Up Array Memory

To use read or write caching or bind RAID-3 LUNs, you must have the required hardware, and specify the array memory partitions. This section lists the hardware required for caching and RAID-3 LUNs and describes how to partition memory.

C.3.1.1 Hardware Requirements for RAID-3 LUNs

All Fibre Channel RAID enclosure implementations support RAID-3 LUNs. Each SP has at least as much memory as the sum of the following:

- 16 Mbytes for system buffers in a fibre channel array
- At least 15 Mbytes of memory should be allocated for each SP in an enclosure

C.3.1.2 Partitioning Array Memory

You partition memory for each SP using the Memory Partition dialog. The Total box in dialog specifies the amount of SP memory that you can partition. It is equal to the SP's total memory minus the memory the SP requires for itself.

To specify memory partitions open the storage management window and:

1. If you want to specify the memory for just one array, simply double-click the array icon to open the Array Configuration dialog and skip to step 2.
2. In the Array Configuration dialog for an array, either click the Partition Memory button or follow the menu path Array > Partition Memory. A Memory Partition dialog for the array opens.
3. Move the slider, and the pie chart changes to show portions of memory allocated to the SP B Read Cache partition.

Note: If you change the RAID-3 partition size, the flare reboots the array. Rebooting restarts the SPs in the array, which terminates all outstanding I/O to the array. If you enable automatic reboot, do not complete this procedure until you have unmounted any file systems or partitions on the array and quiesced the bus.

If you want a different size for the RAID-3 partition, either enter the size in the RAID-3 field for one SP or drag the slider to the right or left until the desired number of Mbytes (MB) appears in the field.

The RAID-3 partition on both SPs must have same size. As a result, when you move the slider for the RAID-3 partition on one SP, the slider for the RAID-3 partition on the other SP also moves, and both pie charts change to show portion of memory allocated to the RAID-3 partition on both SPs.

- Click the OK button to change the memory partitions to the sizes you specified. A confirmation popup opens.
- Click the Yes button in the popup to confirm the new size(s).

You may have to restart the ssmagent on the host system connected to the array whose memory you have partitioned.

Stop the agent and then restart it as discussed at the beginning of this appendix.

C.4 Creating Logical Unit Numbers (LUNs)

You must bind disk modules into LUNs so that it the server's operating system will recognize them.

Before you create LUNs, you need to know which disk modules you want to use for each LUN. You need to make sure that the array has disk modules in each of the required slots, and that these disk modules are not already bound into LUNs.

In the Array Configuration window click on the icon that brings up the bind LUN dialog box and select a specific enclosure.

Under certain circumstances you may need to move disk modules. Generally, modules should not be moved from one slot to another slot; but if moving one is absolutely necessary, the system operator or service person can move it with the following cautions:

The disk module must be unbound. Moving a module that is part of a LUN to another slot makes all information on the LUN inaccessible. The "Unbinding a LUN" section discusses this topic.

If you remove and install a disk module while the array is powered, you must use the proper procedures. Note that drives 0, 1, and 2 have specific code resident and should not be swapped with other disks in a RAID array.

C.4.1 Binding Disk Modules as LUNs

In a fibre RAID array, since each disk module is on only one of several internal buses, you should be careful about which disk module you choose for a LUN. To have good choices of disk modules available for different RAID types, you should bind LUNs in a fibre array in this order:

1. RAID-3
2. Hot spares

Note: The binding of a disk array into a LUN may take as long as two hours to complete. During this time, the disks are unavailable for any I/O processes until the bind is complete.

Before you start binding disk modules in either type of array, be sure you understand the following restrictions:

- You can bind only unbound disk modules. All disk modules in a LUN must have the same capacity to fully use the modules' storage space.
- You must bind exactly five or nine disk modules in a RAID-3 fibre array.
- The capacity of a disk module bound as a hot spare must be at least as great as the capacity of the largest disk module that it might replace.

Note: If you bind disk modules into a RAID-3 LUN without allocating adequate memory to the RAID-3 partition, the LUN will be unowned. To correct this situation, you need to allocate memory to the RAID-3 partition, unbind the LUN and then rebind the LUN.

Open the Storage Management window and do the following:

1. Select the array with disk modules you want to bind into LUN(s).
2. If you want to bind disk modules for just one array, you can simply double-click the array icon to open the Array Configuration dialog for the array, and then skip to the next step.
3. In the Array Configuration dialog, either click the Bind LUN button on the array toolbar or follow the menu path Array > Bind LUN. The Bind LUNs dialogs for the selected array opens.
4. In the Bind LUNs dialog for the array, select the enclosure containing the disk
5. If the Unbound Disks field contains the name of the desired enclosure, continue to the next step. If it contains anything else, pull down its selection list and select the name of the desired enclosure.
6. In the Unbound Disks area, select disk modules by clicking on them; then move them to the bind disks box by clicking the right arrow. Repeat for all additional disks. The disk modules you selected move from the Unbound Disks area to the Bind Disks area.
7. If you move a wrong disk module into the Bind Disks area, simply select it and click the left arrow button to move it back into the Unbound Disks area.

Note: Do not select any bind parameters for the LUN until you have selected all the disk modules for it. The number of disk modules you select determines which RAID types are available, and the RAID type determines availability of the other bind parameters.

C.4.2 Selecting a RAID Type

If you want the LUN to be a different RAID type than the one in the RAID Type field, click the field list button and select the desired type from the list that opens.

Only the RAID types that are available for the number of selected disk modules appear in the list. Be sure the correct RAID type is selected prior to clicking the Bind button.

Raid Type	Number of Disk Modules	Description
RAID-3	5 or 9 fibre disks	Parallel access group
Hot spare	1 fibre disk	The hot spare must have data capacity equal to or greater than the disks in the RAID array.

C.4.3 Selecting a LUN ID

If you want the LUN to have a different hexadecimal identifier (ID) than the number in the LUN ID field, click the field list button and select the desired number from the list that opens.The default LUN ID is the next hexadecimal number available, starting with 0 and ending with 31. Only available numbers appear in the list. The default number is 0 for the first LUN that you bind, regardless of the number of SPs or servers attached to the array. The default number for the second LUN you bind is 1; for the third LUN, it is 2; for the fourth LUN, it is 3, and so on.

If for any reason you want to specify a nondefault number you can do so. After you bind a LUN with a nondefault number, the default number for the next LUN is the lowest number you skipped. For hot spares, you should assign LUN numbers starting with the highest number available and continue downwards.

C.4.4 Changing the Rebuild Time

If you want to have a different rebuild time than the time in the Rebuild Time field, either enter the number of hours in the field or click the field list button and select the desired number of hours from the list that opens.

The rebuild time is the amount of time that the array allots to reconstruct the data on either a hot spare or a new disk module that replaces a failed disk module in a LUN. The time you specify determines the amount of resource the SP devotes to rebuilding instead of to normal I/O activity. The default time of 4 hours is adequate for most situations. A rebuild time of 2 hours rebuilds the disk more quickly, but degrades response time slightly. A rebuild time of ASAP (as soon as possible) rebuilds the disk module as quickly as possible, but degrades response time significantly.

The actual rebuild time can differ significantly from the time you specify, especially for a LUN containing 8.8 Gbyte disk modules.

C.4.5 RAID Array Verify Time

The verify time is the amount of time that the array allots to checking parity during a rebuild.

If you want to have a different verify time than the time in the Verify Time field, click the field list button beside and select the desired number of hours from the list that opens. If an SP detects parity inconsistencies, it starts a background process to check all the parity sectors in the LUN. The time you specify determines the amount of resource the SP devotes to verifying instead of to normal I/O activity.

The default time of 4 hours is adequate for most situations. A verify time of 2 hours rebuilds the disk more quickly, but degrades response time slightly. A verify time of ASAP (as soon as possible) checks the parity sectors as quickly as possible, but degrades response time significantly. If your site requires fast response time and you want to minimize degradation to normal I/O activity, you can extend the rebuilding process over the maximum time of 4 hours.

C.4.6 The Stripe Element Size

The stripe element size is the number of disk sectors that the array can read or write to a single disk module without requiring access to another disk module (assuming that the transfer starts at the first sector in the stripe).

A RAID-3 LUN has a fixed stripe element size of one sector, a value that you cannot change.

C.4.7 The Auto Assign Option

If the LUN is not a hot spare and you want to change the state of the auto assign option for a LUN, click the Auto Assign button. The presence or absence of a visual indicator in the button tells you whether the option is enabled or disabled for the LUN.

Note: You can enable auto assign for arrays with two SPs only. The ssmgui disables this option for arrays with one SP.

Auto assign controls the ownership of the LUN when an SP fails in a array with two SPs. With auto assign enabled, if the SP that owns a LUN fails and the server tries to access that LUN through the second SP, the second SP assumes ownership of the LUN so the access can occur. The second SP continues to own the LUN until the SP's power is cycled (turned off and on again). When the power is cycled, ownership of each LUN returns to its default SP. If auto assign is disabled in the previous situation, the other SP does not assume ownership of the LUN, so the access to the LUN does not occur.

As a general rule, you should disable auto assign; however, some applications require that it be enabled.

C.4.8 Setting Fibre Channel RAID Minimal Latency Reads

If you want to change the state of the minimal latency reads option for a RAID-3 LUN, click the Minimal Latency Reads button. Minimal latency reads provide a more constant bandwidth between the SP and the server when the bandwidth between the disk modules in a RAID-3 LUN and the SP decreases because of a slow response from a disk module. Minimal latency reads change how the SP responds to a request from the server to read a RAID-3 LUN. This option is called **NMinusOneNGo** in the CLI (see Appendix B).

The presence or absence of a visual indicator in the button tells you whether the option is enabled or disabled for the LUN.

You can enable minimal latency reads for fibre channel arrays only.

When minimal latency reads are disabled, the SP responds to a read request to a RAID-3 LUN with n disk modules as follows:

- It issues a read request to the n-1 disk modules in the RAID-3 LUN that contain data.
- It performs a checksum on the data as it receives the data from the n-1 disk modules.
- If the checksum is correct, it sends the data to the server; if the checksum is incorrect, it reads the parity from the remaining disk module and uses it to reconstruct the data before sending the data to the server.

When minimal latency reads are enabled, the SP responds to a read request to a RAID-3 LUN with n disk modules as follows:

- It issues a read request to all n disk modules in the LUN.
- It performs a checksum on the data as it receives it from the first n-1 drives to respond.
- If the checksum is correct, it sends the data to the server; if the checksum is incorrect, it uses the parity it has received to reconstruct the data before sending the data to the server.

Because minimal latency reads always require access to all the disk modules in the RAID-3 LUN, they nominally decrease the bandwidth between the disk modules and the SPs. For a RAID-3 LUN with five disk modules, the bandwidth decreases by about 25%; for one with nine disk modules, it decreases about 12%.

C.4.9 Setting the Default SP Option

If the LUN is *not* a hot spare and you want the LUN to have a different default SP owner than the one shown in the Default SP box, click the button for the other SP. When you select one SP, the other SP is automatically deselected.

You can change the default SP in arrays with two SPs only.

The default SP is the one that assumes ownership of the LUN after the array's power is cycled. If the array has two SPs, you can choose to bind some LUNs using one SP as the default, and the rest using the other SP as the default. You would do this to balance the load across the SPs and/or to establish the primary route for the LUN in dual-server configuration. The primary route to a LUN is the route through the default SP, and the secondary route is through the other SP.

Use the following steps to finish the bind:

1. When all bind parameters for the LUN are set as you want, click the Bind button. A confirmation popup opens.
2. In the popup, click the Yes button to start the bind operation. A popup opens stating that the bind operation was successfully initiated.

The icon for the LUN appears in the Unowned LUNs area of the Array Configuration dialog. When the bind operation is completed, the icon moves to the area for its default SP.

Binding takes awhile (up to two hours); exactly how long varies with the type of SP and size of the disk modules. When the bind operation for a LUN is complete, an icon for the LUN appears in the LUN selection area of the Array Configuration dialog box for the array.

C.5 Updating Licensed Internal (flare) Code

The licensed internal code (flare code) media may also include an update to the SP programmable read-only memory (PROM) code. When you install flare code, the SP tries to copy it to reserved areas outside operating system control on several disk modules, which are called the database disks. Having multiple copies of code offers higher availability if a disk module fails. The database disk modules for the different array types are disk modules 0, 1, and 2 in the Fibre Channel RAID enclosure.

When you install flare code, at *least* two of the database disk modules must be on line, and ideally, all of them should be on line. A disk module is on line if it is fully powered up and not faulted; that is, if its state in the Disk Information dialog Normal. If you try to power up the array without two of these disk modules in place, the powerup fails.

The following procedure assumes that the system operator or service person has loaded the new revision of flare onto the server to which the array is connected. If the flare media also includes new PROM code, the PROM code is loaded along with the flare.

To update flare code and PROM code:

1. Open the Storage Management window and select the array whose flare code you want to update.

   Note: If you want update the flare on just one array, you can simply double-click the array icon to open the Array Configuration dialog for the array, and then skip the next step.

2. Either click the Configure Array button in the array toolbar or follow the menu path Array > Configure. An Array Configuration dialog opens for each selected array.

3. In the Array Configuration dialog, click the Firmware Download button on the array toolbar. A Firmware Download dialog box opens.

Caution: Rebooting restarts the SPs in the array, which terminates all outstanding I/O to the array. If you enable automatic reboot, do not complete this procedure until you have unmounted any file systems or partitions on the array and quiesced the bus.

The presence or absence of a visual indicator in the button tells you whether automatic reboot is enabled or disabled.

The SPs must be rebooted for the new revision of LIC (flare code) to take effect. If you select automatic reboot, the SPs are rebooted automatically when the LIC download operation is complete. When you reboot an array's SP(s) you may have to restart the ssmagent on the array's host. Enter the full path name for the LIC (flare code) file in the upper right box.

```
/usr/lib/flare/K5/flarecode.bin
```

In the Firmware Download dialog box, click the OK button. A confirmation popup box opens.

In the popup, click the Yes button to start downloading flare code to the selected array.

C.6 Configuring a LUN

You configured a LUN initially when you created it as described in "Binding Disk Modules as LUNs" on page C-14. After you create a LUN, you can change some of the LUN bind parameters, that you specified when you created the LUN. Exceptions are the RAID type, capacity, and LUN number.

If you want to change the LUN RAID type or the number or capacity of the disk modules in the LUN, you must unbind the disk modules that comprise it and then rebind the desired disk modules into another LUN.

This section describes how to

- Unbind a LUN
- Change the LUN RAID type or the number or capacity of its disk modules

C.6.1 Unbinding a LUN

Typically, you unbind a LUN only if you want to change its disk modules to be part of a different LUN or change the capacity of its disk modules (size of the LUN). In either case, you should make sure that the LUN contains the disk modules that you want.

To determine which disk modules make up a LUN:

1. In the Storage Management window, select the array with the LUN whose disk modules you want to identify.
2. Either click the Configure button on the toolbar or follow the menu path Array > Configure. The Array Configuration dialog for the selected array opens.
3. In the LUN selection area of the Array Configuration dialog, select the LUN whose disk modules you want to identify.
4. In the Disk Field, confirm that each disk module comprising the LUN is surrounded by a box with the same color as the box surrounding the selected LUN.

You can also determine which disk modules comprise a LUN using the LUN IDs. If the LUN IDs do not appear in any disk modules in the Disk Field, follow the menu path from the Array Configuration dialog: View > Show LUN IDs.

Caution: Unbinding a LUN destroys any data on it. Before unbinding any LUN, make a backup copy of any data on the unit that you want to retain.

Unbind a LUN using the following steps:

1. In the Storage Management window, select the array with the LUN you want to unbind.
2. Either click the Configure button on the toolbar or follow the menu path Array > Configure
3. The Array Configuration dialog for the selected array opens.
4. In the LUN selection area of the Array Configuration dialog, select the LUNs you want to unbind.
5. Either click the Unbind LUN button on the LUN toolbar or follow the menu path LUN > Unbind.
6. A popup opens warning you that unbinding the causes data loss and asking you to confirm the unbind operation.
7. In the popup, click the Yes button to confirm the operation.
8. A popup opens telling you that the LUNs were unbound. The LUN icon disappears from the LUN selection area of the Storage Management window.

C.6.2 Changing the LUN RAID Type or Number of Disks

If you want to change the LUN RAID type or the number or capacity of its disk modules, you must change the bound configuration of a LUN, which means in effect, that you destroy the LUN and recreate it.

When you want to change the number of disk modules, use the same guidelines as for a new configuration.

Caution: The disk modules that you use for a LUN must have the same capacity to fully use the modules' disk space.

To change a LUN configuration you must:

1. Back up any information on the LUN that you want to save because unbinding a LUN destroys all information on the LUN.
2. Unbind the LUN.
3. Install any new disk modules you need for the new configuration using the installation information in this manual.
4. Move any disk modules necessary for the new configuration. Remember that disks 0, 1, and 2 in the RAID array contain special internal code and should not be moved.
5. Bind the disk modules into the desired LUNs.
6. Make the array available to the host.

If you want to replace all the disk modules in an array, use the procedures in the next section.

C.6.3 Replacing all Disk Modules in a LUN

Do not power off the array enclosure while replacing all the disks in an array.

Replace all the disk modules in a LUN using the following steps:

1. Unbind the LUNs whose disk modules you want to replace.
2. Replace all the disk modules you want to replace except for disk module 0 and 1 in the base Fibre Channel RAID enclosure. In Fibre Channel RAID expansion enclosures, you can remove all drives in the enclosure at this point. If you have a Fibre Channel RAID expansion enclosure, skip to step 8.
3. Reload the Fibre Channel RAID flare code, "Updating Licensed Internal (flare) Code" on page C-18 gives instructions for this procedure.
4. Remove disk module 1 and insert a replacement disk module.
5. Repeat step 3.
6. Remove disk module 0 and insert a replacement disk module.
7. Repeat step 3.
8. Bind the disk modules into the desired LUN.

C.7 Monitoring an Array

You monitor the health of managed arrays using the Storage System Manager window. A description of the primary administrative icons is provided in the section "The Storage System Manager Window" on page C-5. You can also monitor arrays by checking the event messages that the ssmgui sends to an SP log as described in "Checking and Saving System Event Messages" on page C-6.

This section describes more details about using the Storage System Manager window status icons and pull-down menus. It also explains how to monitor the status of managed arrays.

The Storage System Manager window shows array health in the following ways:

- Array Status button that shows the health of all managed arrays at a glance.
- Array icons that show the health of the arrays they represent.

The Array Status button provides the overall status of all arrays managed by the current ssmgui session. The button is located in the lower left corner of the window. When it appears as a grey box with the word Normal, the agent has not detected any errors in any managed arrays.

If the array status button is amber and displays the word Fault, it has detected a failure in one or more arrays, or one or more arrays is inaccessible.

The array icons are displayed on the screen in grey when they are functional and accessible. An amber indicator partially covers the array if a fault exists or the array is not accessible.

C.7.1 Automatic Monitoring Status

This section describes how to update the status, statistics, and event log information that the ssmagent automatically maintains for managed arrays. It outlines procedures you can use to detect and isolate problems that may arise with managed arrays. You can vary the procedures to suit your particular array environment.

C.7.1.1 Automatic Polling

Before you use the ssmgui to monitor the health of managed arrays, you need to make sure that the ssmagent has up-to-date information on the state of these arrays. The GUI updates its information for each array by polling the *ssmagent* on the managed hosts for changes in the state of the disk, SP, LUN, and read and write cache (when applicable). You can have ssmgui automatically poll one or more arrays at specified intervals that can vary for each array and/or you can manually poll selected arrays whenever you want. Automatic polling and manual polling are completely independent procedures; that is, one does not affect the other.

C.7.1.2 Enabling and Disabling Automatic Polling

By default, automatic polling is disabled for the ssmgui session and for each array because the polling process uses array resources and, as a result, affects performance. You enable automatic polling for an ssmgui session by specifying an automatic polling interval, and for an array, by specifying an automatic polling priority for the array.

When automatic polling is enabled for the GUI session and a selected array, the frequency at which it automatically polls the array equals the automatic polling interval multiplied by the automatic polling priority for the array. For example, if the automatic polling interval is 5 minutes and the automatic polling priority for the array is 3, ssmgui polls the array every 15 minutes.

To set the automatic polling interval:

1. From the Storage Management window, follow the menu path Options > Automatic Polling Interval.
2. In the submenu that opens, either select the desired number of seconds or minutes or Other to specify a number that is not listed.

To specify the automatic polling priority:

1. In the Storage Management window, select the arrays whose automatic polling priorities you want to specify.
2. From the Storage Management window, follow the menu path Array > Automatic Polling Priority.
3. In the submenu that opens, select the desired priority.

To disable automatic polling:

1. From the Storage Management window, follow the menu path Options > Automatic Polling Interval.
2. From the Storage Management window, follow the menu path Array > Automatic Polling Priority.
3. In the submenu that opens, select "Disable."

C.7.1.3 Monitoring Managed Arrays

If any of the arrays you want to monitor do not appear in the Storage Management window, follow these steps:

1. From the Storage Management window, follow the menu path File > Select Hosts. The Host Administration dialog box opens.
2. For each server whose arrays are missing from the Storage Management window enter the server host name in the Host Identifier field.
3. Click the Add button. The new host names appear in the list of hosts below the field.
4. When the list contains just the servers you want, click the OK button.
5. In the popup that opens, click the Yes button to save the new list to the */ssmhosts* file; otherwise, click the No button.

You should either manually poll the managed arrays periodically or enable automatic polling.

If the Array Status button displays Normal, all managed arrays are operating normally.

If the Array Status button displays Fault, or any array icon is amber , one or more managed arrays are faulted. If you need more information about the fault, follow these steps, click the Array Status button.

A popup opens that tells you which arrays failed and gives some information about the failures. If you want to determine the state of each CRU in the faulted arrays, follow these steps:

1. Select the faulted array(s).
2. Open the Components dialog box for the selected array(s) by either clicking the Monitor button on the array toolbar or following the menu path Array > Monitor.
3. A Components dialog box opens.
4. In the Components dialog, select the faulty enclosure; that is, any enclosure represented by an amber icon, see Figure C-7.

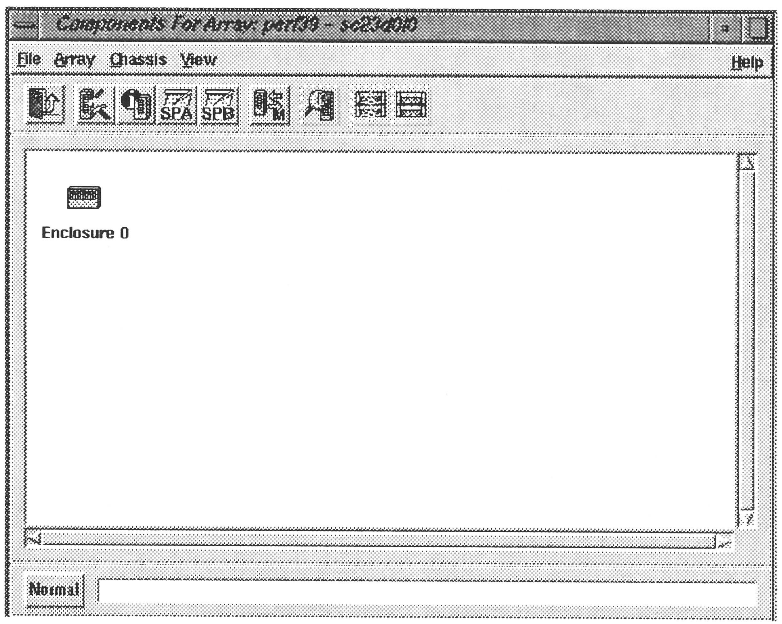

Figure C-7 Components Dialog Box

5. You open the Equipment View dialog box for the selected enclosure(s) by following the menu path Chassis > Monitor.
6. In the Equipment View for each enclosure, look for amber CRU icons.

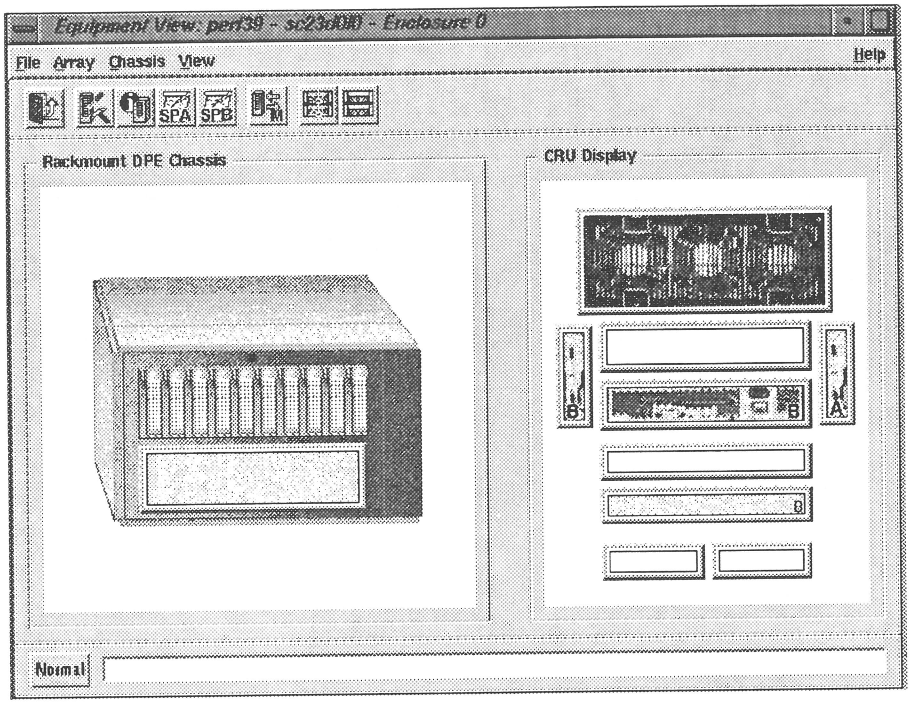

Figure C-8 Equipment View Display Box

If you do not know the name of CRU represented by an amber icon, you can display it by positioning the cursor over the CRU for a few seconds without moving it. The GUI displays a description of the CRU near the icon.

You should look at the Array Status button or the array icons periodically. You can update the information they reflect by selecting an array you are monitoring and clicking the Manual Poll button.

C.7.1.4 The Amber Disk Icon

An amber disk module indicates that the disk module it represents is in one of the following states:

- Either powered up but inaccessible or powered down by the SP, (this can happen when someone inserts a replacement disk module of the wrong capacity).
- Removed from the enclosure after the agent started running.

If the disk module is part of a RAID-3 LUN, the LUN can continue functioning without interrupting user applications. In many cases, you can determine if the LUN is using a hot spare by clicking the button for the LUN.

The system operator can also replace or reinsert the disk module in a RAID-3, LUN without powering down the array or interrupting user applications. Section C.6 describes how to replace a disk module.

Caution: Removing the wrong disk module can introduce an additional fault that shuts down the LUN containing the failed module. Before removing a disk module, be sure to verify that the suspected module has actually failed by checking the error log and the module's check or fault light.

Before anyone replaces the faulty disk module, read the event log for the SP that owns the LUN containing the faulty disk module for a message about the disk module. Also check for any other messages that indicate a related failure, such as a failure of a fibre channel loop connection or a general shutdown of an enclosure. Such a message could mean the disk module itself has not failed. A message about the disk module will contain its module ID. If the disk module button does display the ID, then in the Array Configuration dialog, follow this menu path: View > Show Disk IDs.

Although you can remove a disk module within a RAID-3 LUN without damaging the data on the LUN, you should do this only when the disk module has actually failed.

C.7.1.5 RAID-3 LUN Rebuilding Process

If the array contains a hot spare on standby, the SP automatically rebuilds the failed module on the hot spare. When you replace a disk module in a RAID-3 LUN, the SP formats and checks the new module, and then begins to reconstruct the data. While rebuilding occurs, you have uninterrupted access to information on the LUN.

Use the Disk Information dialog box to monitor the status of the new disk module. In the Array Configuration window, double click on the disk icon at the bottom of the window (next to the enclosure number).

A new disk module's state changes as follows:

- It goes through a powering up stage.
- Rebuilding - The array is reconstructing the data on the new disk from the information on the other disk modules in the LUN. If the disk module is the replacement for a hot spare that is being integrated into a redundant LUN, the state is Equalizing instead of Rebuilding. In this situation, the array is simply copying the data from the hot spare onto the new disk module.
- Enabled - The disk module is bound and assigned to the SP being used as the communication channel to the enclosure. If the array has two SPs, the state is Ready instead of Enabled if the LUN is owned by the SP not being used as the communications channel.

A hot spare's state changes as follows:

- Rebuilding - The SP is rebuilding the data on the hot spare.
- Enabled - The hot spare is integrated fully into the LUN, or the failed disk module has been replaced with a new module and the SP is copying the data from the hot spare onto the new module.
- Ready - The copy is complete. The LUN consists of the disk modules in the original slots and the hot spare is on standby.
- Rebuilding occurs at the same time as user I/O. The time specified for rebuilding determines the duration of the rebuild process and the amount of SP resources dedicated to rebuilding. A short rebuild time consumes many resources and may degrade performance significantly. A long period consumes fewer resources with less effect on performance. You can determine the rebuild period by looking at the configuration section of the LUN Information dialog box.

C.7.1.6 The Amber Storage Processor Icon

An amber SP icon indicates that the SP has failed. When an SP fails, one or more LUNs may become inaccessible. In addition, the SP's check or service light turns on, along with the check or service light on the front of the enclosure. If the enclosure has a second SP and a number of LUNs are inaccessible, you may want to transfer control of the LUNs to the surviving SP.

The system operator or service person can replace the SP under power, without interrupting applications. If array write caching is enabled, disable it before replacing an SP.

C.7.1.7 The Amber LCC Icon

An amber LCC icon indicates that the LCC has failed. In addition, the LCC's fault light turns on, along with the service light on the front of the array. When an LCC fails, the SP it is connected to loses access to its LUNs and the array's performance may decrease. If the array has a second LCC and a number of LUNs are inaccessible, you may want to transfer control of the LUNs to the SP that is connected to the working LCC.

The system operator or service person can replace the SP under power, without interrupting applications. This manual describes how to replace a faulty LCC.

C.7.1.8 The Amber Fan Icon

An amber fan icon indicates that the fan module it represents is in one of the following states:

- The fan module was removed
- The fan module was removed before the agent started running
- The fan module has a failed component (fan)

If one fan fails in a fan pack or module, the other fans speed up to compensate so that the array can continue operating.

If a second fan fails and the temperature rises, the array shuts down after about two minutes.

If you see an amber fan icon, the system operator or a service person should replace the entire fan pack or module as soon as possible.

Caution: Removing the fan module from the enclosure for more than two minutes causes the array to shut down to prevent damage to the disk modules from overheating. You should do this only long enough to replace a failed CRU.

C.7.1.9 The Amber Power Supply Icon

An amber power supply icon indicates that the power supply it represents is in one of the following states:

- The supply failed, was turned off, or was removed after the ssmagent started.
- The supply failed or was turned off before the *ssmagent* started.

A Fibre Channel RAID enclosure has two power supplies: A and optionally B. An array with two power supplies can recover from the failure of one power supply and provide uninterrupted service while the defective power supply is replaced. If a second power supply fails or is removed, the entire enclosure shuts down immediately.

C.7.2 Displaying a Storage Processor Event Log

Each SP maintains a log of event messages. These events include hard errors, startups, and shutdowns involving disk modules, fans, SPs, LCCs, power supplies, and other options. The messages are ordered by the time the event occurred. As a result, the most recent messages are at the end of the log. Periodically, the SP writes this log to disk to maintain it when SP power is off. The log can hold up to 2,000 event messages.

You can display all the events in the log or only the events that occurred from a date and time you specify to the current time. You can also save the contents of the log to a file you specify. The Storage System Manager window and its icons are shown in Figure C-3. "Checking and Saving System Event Messages" on page C-6 also provides additional overview information on the event messages.

The following procedures allow you to customize the SP Event Log for your particular needs.

To open the Event Log dialog for an SP:

1. In the Storage Management window, click on and select the array whose SP event log you want to view.
2. Display the SP A or SP B event log for the selected array by either clicking the History SP A (or B) button or follow the menu path Array > History > SP A (or SP B).

The Event Log dialog for SP A opens and displays all the events in the SP log since the date and time specified in the Log Entries "As Of" field.

To display only events from a date you specify:

1. In the Log Entries As Of field of the SP Event Log dialog, specify the date and time from which to display events.
2. As follows: either enter the date and time in the field or click the field list button and select the desired time from the list that opens.

   The new date and time appears in the field.

   The list of events is updated to show only events from the new date and time.

To clear the list of displayed events:

1. Click the Reset button.
2. You see only events added to the SP log as of the time you clicked the Reset button.

To display all the events:

1. Click the Show All button.
2. The Event Log dialog to shows all the events in the SP log.

To display only events from the date specified in the Log Entries As Of field:

1. Click the Refresh button.
2. The Event Log dialog is reinitialized.

It displays only the events in the SP log since the date and time specified in the Log Entries "As Of" field.

To save the contents of the log to a file:

1. Click the Save button in the Event Log dialog box.
2. A Save SP Log Events file section dialog opens.
3. Select the name of the file where you want the contents of the log written.
4. Select the file type.
5. Select the drive and directory where you want the file written.
6. Click the OK button. If no errors occur, an information popup opens stating that the log has been saved.
7. Click the OK button in the popup to close it.

To exit the log, click the Close button in the Event Log dialog box.

Unsolicited event log messages appear in the log in the form of hexadecimal codes and messages. The 0x600-series codes are informational codes that do not indicate an error condition. The 0x900-series and 0xa00-series codes indicate a serious error condition. See the system release notes for a listing of the unsolicited event log messages.

C.8 Displaying Status and Statistics

This section describes how to display all the Information dialogs that provide status and statistics information about an array, SP, LUN, disk module, or caching. It also describes how to turn on and off the statistics log that supplies the information for these dialogs. For information on how to display and save the event log for an SP, see "Checking and Saving System Event Messages" on page C-6.

Many of these dialogs display statistics information. Normally, statistics information is not logged, since logging it may affect performance. For meaningful statistics information, you must enable statistics logging as explained in the next section. If statistics logging is already enabled or you do not care about statistics information, skip the next subsection.

C.8.1 Statistics Logging

Each SP maintains a log of statistics for the LUNs, disk modules, and array caching that you can turn on and off. When statistics logging is on, it affects array performance, so you may want to leave it off unless you have a reason to monitor performance.

The statistics log uses a 32-bit counter to maintain the statistics numbers, so when the counter is full, the statistics numbers start over at zero. As a result, you will see a sudden decrease in a statistics number if you view it shortly before the counter is full and shortly after the counter restarts at zero. If you want to keep the log turned on for more than two weeks, we recommend that you reset the log about every two weeks, so you know when the numbers start at zero.

To turn SP statistics logging on and off:

1. In the Storage Management window, select the arrays whose logging state you might want to change.
2. Either click the Configure button on the array toolbar or follow the menu path Array > Configure. An Array Configuration dialog opens for each array selected.
3. In the Array Configuration screen, select the menu path: Array > Statistics Logging State
4. Select the SP whose statistics logging state you want to change. A popup appears to ask for confirmation.
5. Confirm by clicking Yes or to leave the setting unchanged click No.
6. To change the logging status of the other SP, repeat steps 3 through 5.

To display the Array Information dialog:

1. In the Storage Management window, select the arrays whose status you want.
2. Either click the Array Information button on the toolbar or follow the menu path Array > Information. An Array Information dialog for each selected array opens indicating a state of "Faulted" or "Normal".

When automatic polling is enabled, the automatic polling priority and polling interval determine the frequency at which ssmgui automatically polls the array as follows. If the automatic polling priority is 1, *ssmgui* polls the selected array each time the polling interval elapses; if the priority is 2, *ssmgui* polls the selected arrays after 2 intervals elapse; if the priority is 3, *ssmgui* polls the selected array after 3 intervals elapse; and so on. The automatic polling interval has no effect on manual polling.

When ssmgui polls the selected arrays for status, it retrieves information about array status changes and updates the status information for the polled arrays in any open dialog. An automatic poll of an array functions the same way as a manual poll of the array.

Automatic Polling Count
: Number of priority cycles that elapsed since this array was polled. For example, if the priority is 3, a count of 2 indicates that 2 priority cycles have elapsed; on the next (third) cycle the array will be automatically polled.

Automatic Disk Formatting
: Controls whether the array firmware will automatically format any disk module that has an unrecognizable format. You cannot change this setting.

Host Connection
: The host name followed by the operating system device entry and SP used to communicate with *ssmgui*. The device entry is from the *ssmgui* agent configuration file on the array's server. The SP is the default owner of the LUN. With two SPs, two host connections are required to provide information on each SP.

C.8.2 Storage Processor Information Dialog

The SP Information dialog displays either configuration, statistics, or cache information. When you open the dialog, the SP configuration information is displayed.

To display SP configuration information:

1. In the Storage Management window, select the arrays whose SP's status you want.
2. Either click the Configure button on the toolbar, or follow the menu path Array > Configuration. An Array Configuration dialog opens for each selected array.
3. In the Array Configuration dialog, click the SP A Information or SP B Information button.
4. An SP Information dialog with configuration information opens, it shows the current operating state of the SP.

An SP can be in any of the following states:

- Present: The SP is the communication channel you are using to communicate with the array.
- Not Connected: The agent cannot talk to the SP because a communication channel specifying the SP is not in the agent's configuration file for the selected host. For example, the SP is connected to a different host than the SP in the communications channel for the array.
- Not Present: The SP was not present in the enclosure when the ssmagent started.
- SP Removed: The SP was present in the enclosure when the ssmagent started but has since been removed.

Firmware Revision
: The revision of the Licensed Internal Code (LIC) that the SP is running. Each SP in the array will be running the same revision of LIC. You can update an SP's LIC as described in the "Updating Licensed Internal Code and PROM Code" section.

  Note that if the firmware revision and/or the PROM revision are zero or blank, the host connection is missing or invalid.

PROM Revision
: The revision of the SP's PROM code. Each SP in the array will be running the same revision of PROM code. PROM code is updated automatically when you update an SP's LIC as described in the "Updating Licensed Internal Code and PROM Code" section.

Model
: The product model number for the SP.

Total Memory
: The size in Mbytes (MB) of the SP memory you can configure (total less memory needed for cache buffers or RAID-3 usage). Each SP in an array must have the same amount of memory to make full use of the memory. Total memory can be 8, 16, 32, 64, or 128 Mbytes.

To display SP statistics information:

1. If you are already at the SP Information dialog, skip to step 4.
2. In the Storage Management window, select the arrays whose SPs status you want.
3. Either click the Configure button in the toolbar or the menu path Array > Configuration. An Array Configuration dialog opens for each selected array.
4. In the Array Configuration dialog, click the SP A Information or SP B Information button. An SP Information dialog opens.

5. In the SP Information dialog, click the Statistics button. The SP Information dialog with statistics information opens.

Statistics Logging
: Enabled if SP statistics logging is enabled; Disabled if statistics logging is disabled.

Number of Reads, Number of Writes
: Total number of read and write requests made to the SP.

Number of Blocks Read, Number of Blocks Written
: Total number of data blocks read from and written to the SP.

Percent Busy
: The percentage of time that the SP was busy (processing requests) and not idle. The number shows the relative load on the SP.

C.8.3 LUN Information Dialog

The LUN Information dialog displays either configuration, statistics, or cache information. When you open the dialog, the configuration information is displayed.

To display LUN configuration information:

1. In the Storage Management window, select the arrays whose LUN status you want.
2. Either click the Configure button in the toolbar or follow the menu path Array > Configuration. An Array Configuration dialog opens for each selected array.
3. In the Array Configuration dialog, either double-click the LUN you want to examine or select multiple LUNs and follow the menu path LUN > Display Information > Configuration.

A LUN Information dialog with configuration information opens for each selected LUN, similar to the following.

RAID Type
: The RAID type of the LUN. Available RAID types are:
  - RAID 3 parallel access array
  - Hot spare

State SP A , State SP B
: The SP state, as follows.

Present	The SP is the communication channel you are using to communicate with the array.
Not Connected	The agent cannot talk to the SP because a communication channel specifying the SP is not in the agent's configuration file for the selected host. For example, the SP is connected to a different host than the SP in the communications channel for the array.
Not Present	The SP was not present in the enclosure when the ssmagent started.
SP Removed	The SP was removed after the ssmagent started.

Element Size	The stripe element size is the number of sectors that the array can read or write to a single disk module in the LUN without requiring access to another disk module in the LUN.
Rebuild Time	Maximum number of hours that the array can take to integrate a hot spare or new disk module. This time was specified when the LUN was bound. You can change it at any time without affecting the data on the LUN by using the following menu path from the Array Configuration dialog: LUN > Change Rebuild Time.
Default SP	The SP that owned the LUN when the array was powered up.
Auto Assignment State	Enabled or disabled. The original default is disabled. Auto assignment controls the ownership of a LUN when an SP fails in an array with two SPs. With auto assignment enabled, if the SP that owns a LUN fails and the server tries to access that LUN through the second SP, the second SP assumes ownership of the LUN so the access can occur. The second SP continues to own the LUN until the array is power cycled (turned off and then on again). When the power is cycled, ownership of each LUN returns to its default SP. If auto assignment is disabled in the previous situation, the other SP does not assume ownership of the LUN, so the access to the LUN does not occur.
Disks	The disk modules that make up the LUN. For example, 0:0 - 0:4 means disk modules in slots 0 through 4 in enclosure 0 in a fibre channel array.
Percent Rebuilt	If nonzero, this indicates that a disk in the LUN is being rebuilt and shows the percentage of the rebuild completed.
Percent Bound	If nonzero, this indicates that the LUN is being bound and shows the percentage of the bind completed.

To display the LUN statistics information:

1. If you are already at the LUN Information dialog, skip to step 6.
2. In the Storage Management window, select the arrays whose LUN status you want.
3. Either click the Configure button in the toolbar or follow the menu path Array > Configuration. An Array Configuration dialog opens for each selected array.
4. In the Array Configuration dialog, either double-click the LUN you want to examine, or select multiple LUNs and from the toolbar, select LUN > Display Information > Statistics.
5. A LUN Information dialog with statistic information opens for each selected LUN.
6. In the LUN Information dialog, click the Statistics button. A LUN Statistics dialog opens.

Average LUN Request Service Time
: Average amount of time in milliseconds that all the disk modules in the LUN required to execute an I/O request after the request reached the top of the queue.

Number of Reads, Number of Writes
: Total number of read and write requests made to the LUN.

Number of Blocks Read, Number of Blocks Written
: Total number of data blocks read from and written to the LUN.

C.8.4 Using the Disk Information Dialog Box

C.8.4.1 Disk Information Dialog

The Disk Information dialog displays either configuration, statistics, or error information. When you open the dialog, the configuration information is displayed.

To display the disk configuration information:

1. In the Storage Management window, select the arrays whose disk module status you want.
2. Either click the Configure button in the toolbar or follow the menu path Array > Configuration. An Array Configuration dialog opens for each selected array.
3. In the Disk Field, double-click the disk module you want to examine.

The Disk Information dialog with configuration information opens, similar to the following;

Capacity — Storage capacity of the disk module in Kbytes.

State — Current operational state of the disk module. A disk module can be in any of the states in the table that follows.

State	Meaning
Binding	Being bound into a LUN.
Enabled	Either a hot spare on standby or part of a bound LUN that is assigned to (owned by) the SP you are using as the communication channel to the enclosure. If the array has another SP, this module's status is Ready when you use the other SP as the communication channel to the array.
Equalizing	Data from a hot spare is being copied onto a replacement disk module.
Failed	Powered down or inaccessible.
Formatting	Being hardware formatted. Generally, modules do not need formatting.
Not Present	No disk module is in the slot.
Off	Powered off by the SP, which can happen if a wrong size module is inserted.

State	Meaning
Powering Up	Power is being applied to the disk module.
Ready	Module is part of a broken LUN or a LUN that is bound and unassigned. This can mean that the disk module is part of a LUN that is not owned by the SP that you are using as the communication channel to the enclosure. If the disk module is part of a LUN assigned to an SP other than the one you are using as the communication enclosure, the module's status is either Enabled or Ready. It is Enabled when you use the other SP as the communication channel to the array.
Rebuilding	Module is either a hot spare or replacement disk module that replaced a failed module in a LUN. The data is being rebuilt on a hot spare or a replacement disk module.
Removed	Removed from the enclosure; applies only to a disk module that is part of a LUN.
Standby	Module is either a hot spare or replacement disk module that replaced a failed module in a LUN. The data is being rebuilt on a hot spare or replacement disk module.
Synchronizing	Write cache on the disk module is being synchronized.
Unbound	Ready to be bound into a LUN.

LUN ID — The hexadecimal number identifying the LUN to which the disk module is bound. You specified this number in the Bind window when you bound the LUN.

LUN Type — The RAID type of the LUN to which the disk module is bound. You specified this type in the Bind window when you bound the LUN. Available RAID types are:

- RAID-3 User Sectors - The number of user-accessible sectors (blocks) on the disk module.
- Hot spare

To display the disk statistics information:

1. If you are already at the Disk Information dialog, skip to step 5).
2. In the Storage Management window, select the arrays whose disk module status you want.
3. Either click the Configure button in the toolbar or follow the menu path Array > Configuration. An Array Configuration dialog opens for each selected array.
4. In the Disk Field, double-click the disk module you want to examine. The Disk Information dialog opens.
5. In the Disk Information dialog, click the Statistics button.

The Disk Information dialog with statistics information opens.

Average Disk Request Service Time
: Average amount of time in milliseconds that the disk module required to execute an I/O request after the request reached the top of the queue.

Number of Reads, Number of Writes
: Total number of read and write requests made to the disk module. You might find the LUN read or write information displayed in the LUN Statistics more useful because it is for the entire LUN, and not just for one of the disk modules in the LUN.

Number of Blocks Read, Number of Blocks Written
: Total number of data blocks read from and written to the disk module.

Number of Read Retries, Number of Write Retries
: Total number of times read and write requests to the disk module were retried because of soft or hard errors.

To display the disk errors information:

1. If you are already at the Disk Information dialog, skip to step 5.
2. In the Storage Management window, select the arrays whose disk module status you want.
3. Either click the Configure button in the toolbar or follow the menu path Array > Configuration.

   An Array Configuration dialog opens for each selected array.
4. In the Disk Field, double-click the disk module you want to examine.

   The Disk Information dialog opens.
5. In the Disk Information dialog, click the Errors button.

The Disk Information dialog with error information opens.

Number of Hard Read Errors, Number of Hard Write Errors
: Total number of read or write errors for all the disk modules in the LUN that persisted through all the retries. A growing number of hard errors may mean that one or more of the LUN's disk modules is nearing the end of its useful life.

Number of Soft Read Errors, Number of Soft Write Errors
: Total number of read or write errors for all the disk modules in the LUN that disappeared before all the retries. A growing number of soft errors may indicate that one of the LUN's disk modules is nearing the end of its useful life.

Remapped Sectors
: Total number of disk sectors on all the disk modules in the LUN that were faulty when written to, and thus were remapped to different parts of the disk modules.

Glossary

Note that the following glossary entries are derived from the FC-PH ANSI standard FC-PH REV 4.3, June 1, 1994.

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.

Origin FibreVault

A Fibre Channel chassis with disks only (disk attach enclosure) and no controllers. Available as a standalone unit or as a chassis mountable in the Fibre Channel 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

H

I

J

K

L

P

R

T