Silicon Graphics® Prism™ Visualization System User's Guide

CONTRIBUTORS

Written by Mark Schwenden

Illustrated by Chrystie Danzer

Production by Karen Jacobson

Additional contributions by Dick Brownell, Mark Cabrales, Matt Marchese, Jim Passint, Keith Rich, Victor Riley, Andrew Spray, Armando Serrato, Dave Shreiner, and Lyle Stoll.

Record of Revision

Version	Description
001	October 2004 First Release

Contents

Figures

Tables

About This Guide

This guide provides an overview of the Silicon Graphics Prism Visualization System components, and it describes how to set up and operate this system.

The most basic Silicon Graphics Prism consists of a base compute module combined with an extensible graphics 2-processor capable node module (XG2N module). Larger systems use a base compute module, XG2N module(s), and other optional modules that are rackmounted together. The following optional modules are available:

- A CPU expansion compute module provides zero, one or two additional 64-bit Intel processors and up to 24 GB of memory. No PCI or disk options are available in this module.

- A CPU, memory and PCI/PCI-X expansion (CMPX) module provides zero, one or two optional processors, up to 24 GB of memory, and four PCI and PCI-X card slots.

- The SGI TP900 storage module provides additional disk storage to the system. For detailed information about this module, see *SGI Total Performance 900 Storage System User's Guide* (007-4428-00*x*). Note than an optional SCSI PCI card is required to use this option in a Silicon Graphics Prism.

This guide is written for owners, system administrators, and users of the Silicon Graphics Prism. It is written with the assumption that the reader has a good general knowledge of computer graphics and computer operations.

Important Information

Your SGI system support engineer (SSE) should perform the addition or replacement of parts, cabling, and service of your Silicon Graphics Prism, with the exception of the following tasks that you may perform yourself:

- Installing your system in a rack.

- Cabling the system modules to each other.

- Using your system console to enter commands and perform system functions such as powering on and powering off.

- Using the On/Off, reset, and non-maskable interrupt (NMI) switches on the front panel of your system.

- Removing and replacing XG2N graphics boards.

- Installing and removing PCI and PCI-X cards.

Caution: Exception: Because the installation and removal of an IO10 PCI card, which installs in the lowermost slot, is more complicated and can cause damage to your system if not performed properly, it can only be installed and removed by a trained SGI system support engineer.

- Installing and removing disk drives.

- Installing and removing power supplies.

- Installing and removing DIMMs.

- Installing and removing the L1 controller display.

Warning: **To ensure your safety and protect your system, do not add or replace any components that this guide does not designate as customer replaceable. Contact your SGI system support engineer (SSE) to install any hardware components that are not designated as customer replaceable in this guide.**

Chapter Descriptions

The following topics are covered in this guide:

- Chapter 1, "Getting Started with the Silicon Graphics Prism," provides introductory information, instructions for cabling, and quick start operation of the Silicon Graphics Prism.

- Chapter 2, "System Overview and Options," provides a general overview of the Silicon Graphics Prism and its topology. This chapter also includes a quick description of the various components and modules that can be part of a Silicon Graphics Prism.

- Chapter 3, "Installing and Removing Customer-replaceable Units," describes how to install and remove the following customer-replaceable units (CRUs):

  - PCI and PCI-X cards

  - Graphics boards

  - Disk drives

  - Power supplies

  - Memory (DIMMs)

  - L1 system controller display

- Chapter 4, "Troubleshooting," describes how to troubleshoot your system by using the L1 controller and your system LEDs. Information on reconfiguring the system's XF86Config file for changes in hardware usage is included.

- Appendix A, "Technical Specifications," contains environmental and physical specifications for the Silicon Graphics Prism, as well as pin assignments for non-proprietary connectors for the various modules that can be part of a Silicon Graphics Prism.

- Appendix B, "Installing Silicon Graphics Prism Modules in a Rack," provides instructions for rack mounting modules with optional slide rails.

- Appendix C, "Regulatory Specifications and Safety Information," contains regulatory specifications and safety information related to the Silicon Graphics Prism.

- An index completes this guide.

Related Publications

This section lists various information sources and explains how to access them. The following SGI publications relate to the Silicon Graphics Prism:

- *SGI L1 and L2 Controller Software User's Guide* (007-3938-00x) (available online). This guide describes the L1 and L2 controller functions, commands, and error messages that you may need to operate and maintain your system.

- *SGI Total Performance 900 Storage System User's Guide* (007-4428-00x) (optional). This guide describes the SGI TP900 storage module, which provides additional storage for the Silicon Graphics Prism.

- *SGIconsole Hardware Connectivity Guide* (007-4340-00x) (optional). This guide describes how to connect an optional SGIconsole to SGI systems. You can use an optional SGIconsole to manage and monitor multiple Silicon Graphics systems.

- *SGI Altix 350 User's Guide* (007-4660-00x). This guide contains detailed information on many of the same optional modules that you may use with your Silicon Graphics Prism.

You can obtain SGI documentation, release notes, or man pages in the following ways:

- See the SGI Technical Publications Library at http://docs.sgi.com. Various formats are available. This library contains the most recent and most comprehensive set of online books, release notes, man pages, and other information.

- SGI ProPack for Linux documentation, and all other documentation included in the RPMs on the distribution CDs can be found on the CD titled "SGI ProPack v*X.X* for Linux - Documentation CD." To access the information on the documentation CD, open the index.html file with a web browser. Because this online file can be updated later in the release cycle than this document, you should check it for the latest information.

- The release notes, which contain the latest information about software and documentation in this release, are on the SGI ProPack for Linux Documentation CD in the root directory, in a file named `README.TXT`.

Note: There are no command line `grelnotes` or `relnotes` available on an SGI Linux system. The **InfoSearch** tool is not available with Linux.

Conventions

The following conventions are used throughout this document:

Convention	Meaning
Command	This fixed-space font denotes literal items such as commands, files, routines, path names, signals, messages, and programming language structures.
variable	The italic typeface denotes variable entries and words or concepts being defined. Italic typeface is also used for book titles.
user input	This fixed-space font denotes literal items that the user enters in interactive sessions. Output is shown in nonbold, fixed-space font.
[ ]	Brackets enclose optional portions of a command or directive line.
...	Ellipses indicate that a preceding element can be repeated.
man page(*x*)	Man page section identifiers appear in parentheses after man page names.
GUI element	This font denotes the names of graphical user interface (GUI) elements such as windows, screens, dialog boxes, menus, toolbars, icons, buttons, boxes, fields, and lists.

Product Support

SGI provides a comprehensive product support and maintenance program for its products, as follows:

- If you are in North America, contact the Technical Assistance Center at 1 (800) 800 4SGI or contact your authorized service provider.

- If you are outside North America, contact the SGI subsidiary or authorized distributor in your country.

Reader Comments

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

You can contact SGI in any of the following ways:

*   Send e-mail to the following address:

    techpubs@sgi.com

*   Use the Feedback option on the Technical Publications Library website:

    http://docs.sgi.com

*   Contact your customer service representative and ask that an incident be filed in the SGI incident tracking system.

*   Send mail to the following address:

    Technical Publications
    SGI
    1500 Crittenden Lane, M/S 535
    Mountain View, California 94043

SGI values your comments and will respond to them promptly.

Getting Started with the Silicon Graphics Prism

The Silicon Graphics Prism Visualization System platform is a high-performance highly-scalable system running the SGI Linux + ProPack operating system. The platform is available in a number of configurations to meet your visualization and compute requirements.

This chapter introduces the two primary modules used in the system, and provides an abbreviated overview (quick start) on how to begin using your system. The following chapters elaborate on various topics relating to larger systems, configuration and cabling questions, component replacement and upgrades, basic troubleshooting, rack mounting, and connector functionality.

XG2N and Compute Modules

The most basic Silicon Graphics Prism uses two modules for graphics output:

- A graphics 2-processor and memory capable node module (called an XG2N module).

- A base compute module that provides base I/O as well as processors and memory.

The XG2N module is a 2U rack mountable enclosure which, in addition to the two high-performance graphics pipes, adds to the host system a node board with two CPUs and four, eight, or twelve memory DIMMs. This module connects as an integral part of the system's compute/memory fabric. Though it contains CPUs and memory, the XG2N module does not have boot I/O functionality, and therefore may not be used as a standalone system.

The rear panel of the XG2N module is shown in Figure 1-1 on page 2.

The XG2N module is a node-connected module, and therefore connects to the host system using NUMAlink, either directly or through an optional NUMAlink module (router). This way the XG2N module becomes an integral part of the system, and can contain both CPUs and memory, in addition to graphics output capabilities. Note that NUMAlink routers may be available only with specific system configurations. Check with your SGI sales or service representative for information on available configurations.

Each XG2N module contains two graphics pipes, each capable of supporting two display devices. In systems with only two pipes, the pipe and channel designations are defined as in Figure 1-1. In systems with additional XG2N graphics modules the pipes are allocated on an n+1 basis (see Figure 1-2).
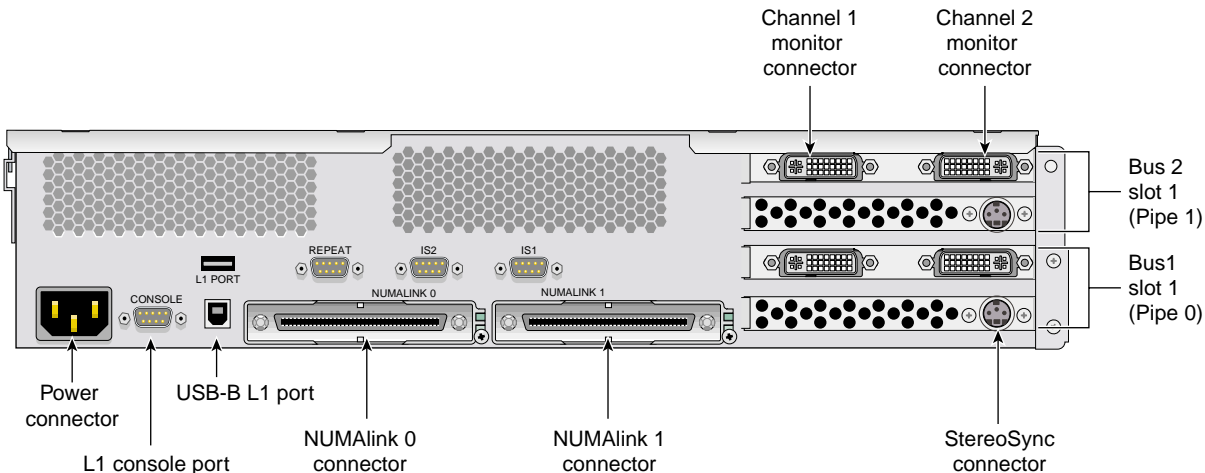
Figure 1-1   Rear View of XG2N Graphics Module
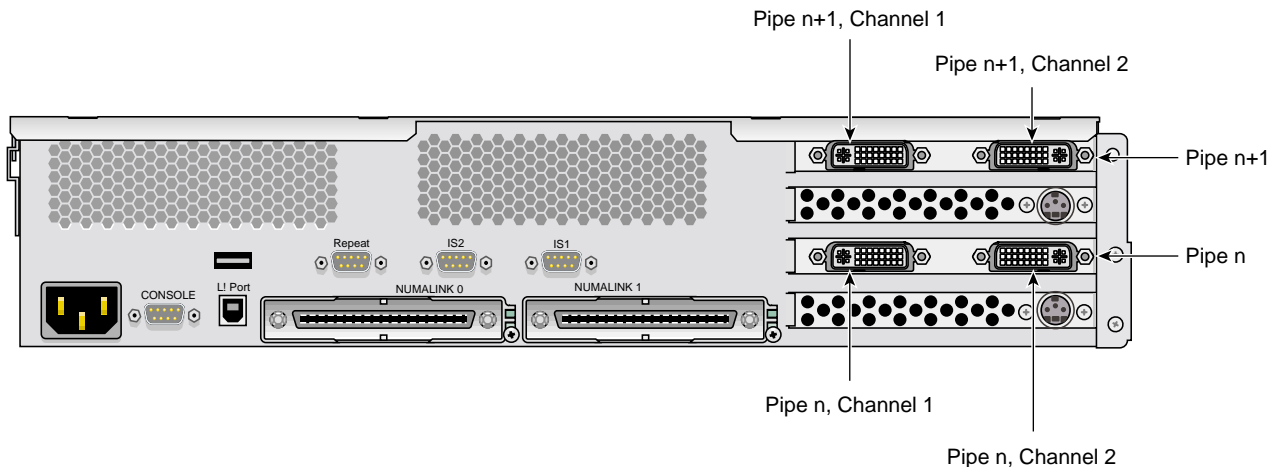
Figure 1-2   Second XG2N Graphics Module

Front Panel Controls

This section describes the front panel controls and indicators of the modules as shown in Figure 1-3.
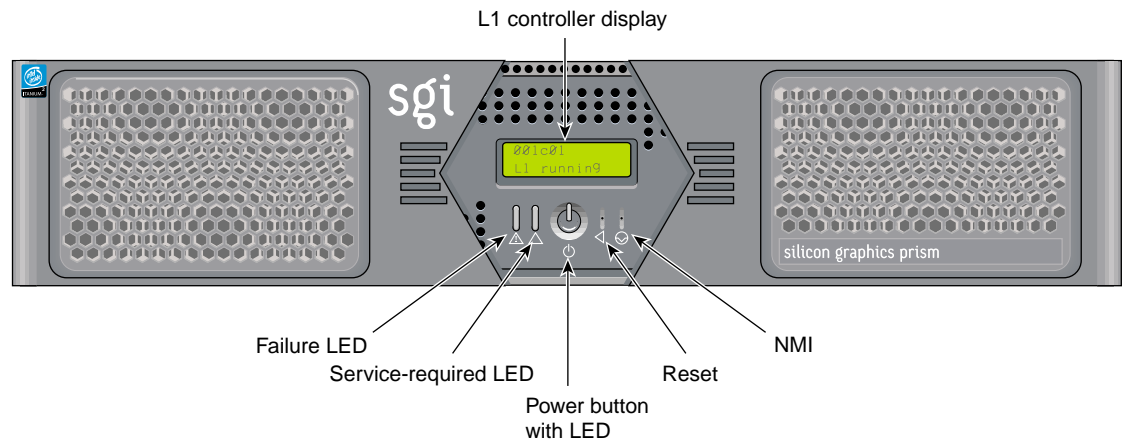
Figure 1-3   Silicon Graphics Prism Module Front Panel Controls

The front panels of the modules have the following controls:

- **L1 controller display.** A liquid crystal display (LCD) that shows status and error messages generated by the L1 controller.

  Note: Refer to the *SGI L1 and L2 Controller Software User's Guide* (007-3938-00x) for more information on the L1 controller.

- **Status LEDs.** The front panel has the following LEDs:
  - Power button LED. This LED illuminates green when the internal components are on.
  - Service-required LED. This LED illuminates yellow to indicate that an item is not functioning properly (for example, a fan is off), but the system is still operating.
  - Failure LED. This LED illuminates red to indicate that a failure has occurred and the system or module has shut down.

- **Power button.** Press this button to power on the individual module. Alternatively, you can power on all the system modules at once from an optional system console.

- **Reset button.** Press this button to reset the internal processors and ASICs. The reset will cause a memory loss.

- **NMI button**. Pressing this button issues a non-maskable interrupt command to a module. When the system hangs, you can send the affected module an NMI interrupt. The interrupt goes to PROM and causes the CPU state to be captured for that module. This information is saved in flash PROM and in the system log, and can assist SGI technicians in debugging system hangs and customer problems.

Compute Module Rear Panel Items

Figure 1-4 shows a view of the rear panel of a Silicon Graphics Prism compute module.
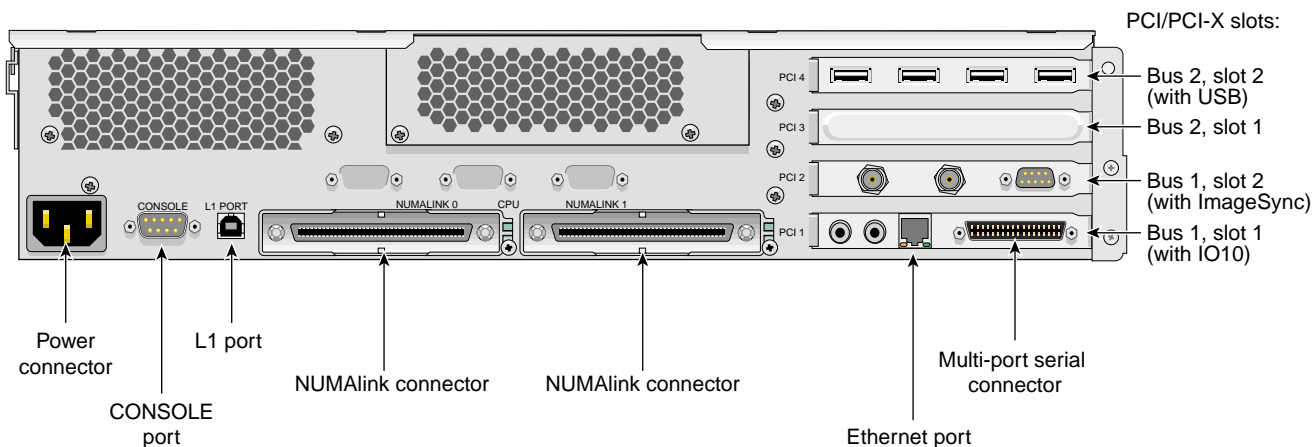
Figure 1-4   Silicon Graphics Prism Compute Module Rear Panel With ImageSync

The rear panel of the module has the following items:

- **AC power input.** This connector connects the graphics module to an AC power outlet.

- **Console port.** This DB-9 serial port (console and diagnostic port) enables you to connect a system console to the L1 controller on the graphics module.

- **L1 port (USB type B).** This universal serial bus (USB) type B connector connects the graphics module L1 controller to an optional L2 controller.

- **NUMAlink connectors.** These NUMAlink connectors connect the XG2N-module or other optional modules to the compute module. This connection is made with a NUMAlink cable at 3.2 GB/s in each direction.

- **NUMAlink LED.** Each NUMAlink connector has a yellow LED and a green LED (located to the right of the NUMAlink connector). The yellow LED indicates that both the graphics module and the host system module to which it is connected are powered on. The green LED indicates that the host system has established a link to the XG2N graphics or other optional module.

Internal Components

Figure 1-5 shows an internal view of the XG2N module.
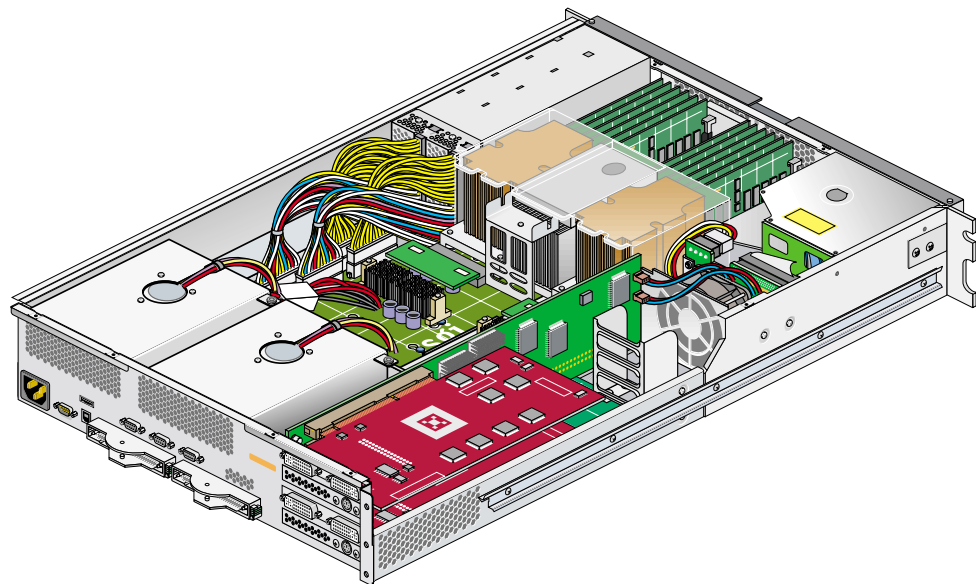
Figure 1-5   Silicon Graphics Prism XG2N Module Internal View

Safety Precautions

Before you install a Silicon Graphics Prism, you should familiarize yourself with the safety precautions discussed in the following subsections:

- "Hazard Statements" on page 6
- "ESD Precautions" on page 6
- "Safety Measures" on page 7

Hazard Statements

During the installation of the computer system, be alert for hazard advisory statements with icons, which signify the following:

- **Caution** Indicates a potentially hazardous situation that, if not avoided, can result in minor or moderate injury. A caution statement also alerts you to unsafe practices that can result in equipment damage and/or data corruption. A caution message is accompanied by an icon as shown in the following example:

  Caution:

- **Warning** indicates a potentially hazardous situation that, if not avoided, could result in death or serious injury. A warning message is accompanied by icon as shown in the following example:

  Warning:

ESD Precautions

Observe electrostatic discharge (ESD) precautions during the entire installation process to eliminate possible ESD damage to the equipment. Wear an SGI-approved wrist strap when you handle an ESD-sensitive device. Connect the wrist strap cord directly to earth ground.

Caution: Observe all ESD precautions. Failure to do so can result in damage to the equipment.

Safety Measures

Observe the following safety measures when you install the system:

- Use caution when you remove the system from the shipping crate. Failure to handle the system carefully can result in personal injury or property damage.

Warning: Ensure that the shipping crate is positioned close to its destination before you unpack the crate.

Warning: Employ a minimum of two people to lift the module or modules off the shipping pallet, to move the module(s) from one location to another, and to install the module(s) in a rack. Otherwise, someone could be seriously injured.

- Do not move the system while it is connected to power.

Warning: Keep fingers and conductive tools away from high-voltage areas. Failure to follow these precautions will result in serious injury or death. The high-voltage areas of the system are indicated with high-voltage warning labels.

- Ensure that a qualified electrician has properly installed the power receptacles.
- Set all circuit breakers to the OFF (**O**) position before you plug in the system power cord.

Warning: Use the following guidelines to prevent the rack from toppling over. Otherwise, people could be seriously injured and/or equipment could be damaged.

Follow these guidelines to prevent the rack from toppling over:

- Make sure that only one module is extended out of the rack at one time.
- Install all equipment in the lowest available position in the rack.
- Ensure that the tip tray is bolted to the front of the rack.

Quick Start Information

The following sections and illustrations are intended to guide a knowledgeable user through the installation, setup, and simple configuration of a basic Silicon Graphics Prism visualization system. For more detailed information on system components or maintenance, go on to the chapter that covers your detailed requirements.

Positioning and Power for Your Silicon Graphics Prism

Depending on the system ordered, your visualization system may have been shipped as two or more individual modules, or pre-mounted in a rack. If the system was shipped as a number of individual chassis, they may be placed on a flat surface, as shown in Figure 1-6 on page 9, or placed in a rack. For instructions on mounting modules in a rack see Appendix B, "Installing Silicon Graphics Prism Modules in a Rack".

Plug your system into a suitable power outlet, as shown in Figure 1-6 (for non-racked chassis) or Figure 1-7 (for racked chassis). Refer to Appendix A, "Technical Specifications" for detailed power requirements for the system modules.
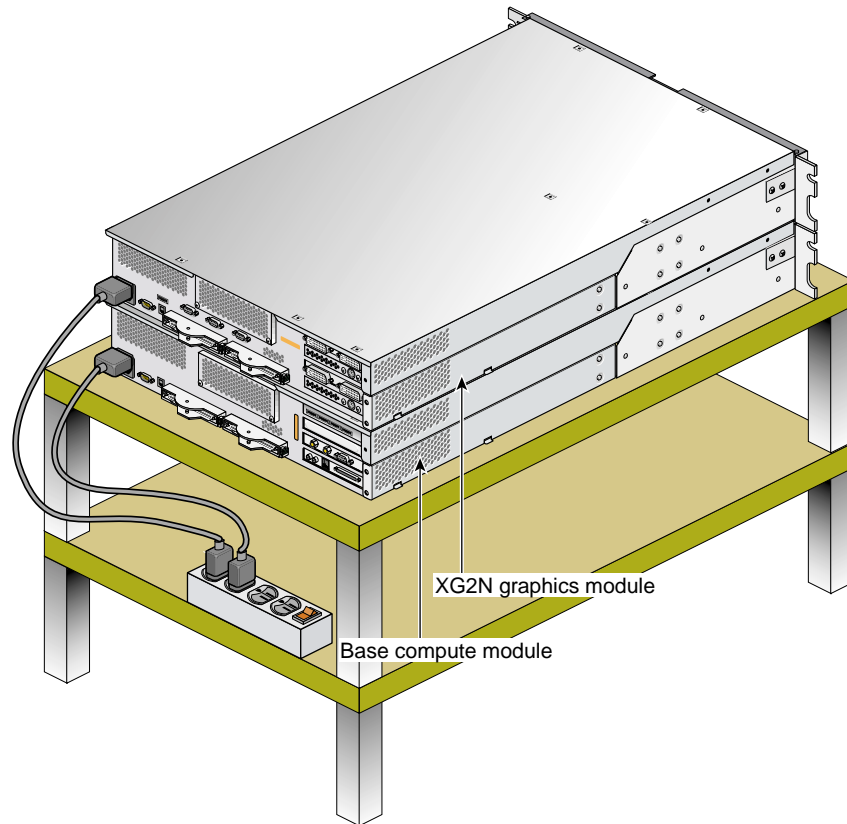
Figure 1-6   Example Connection of a Non-Racked System to Power

CMPX module
(optional)

XG2N graphics
module

XG2N graphics
module

Base compute
module

TP900 storage
module (optional)

Power
source

Power
distribution
unit (PDU)

Components within rack
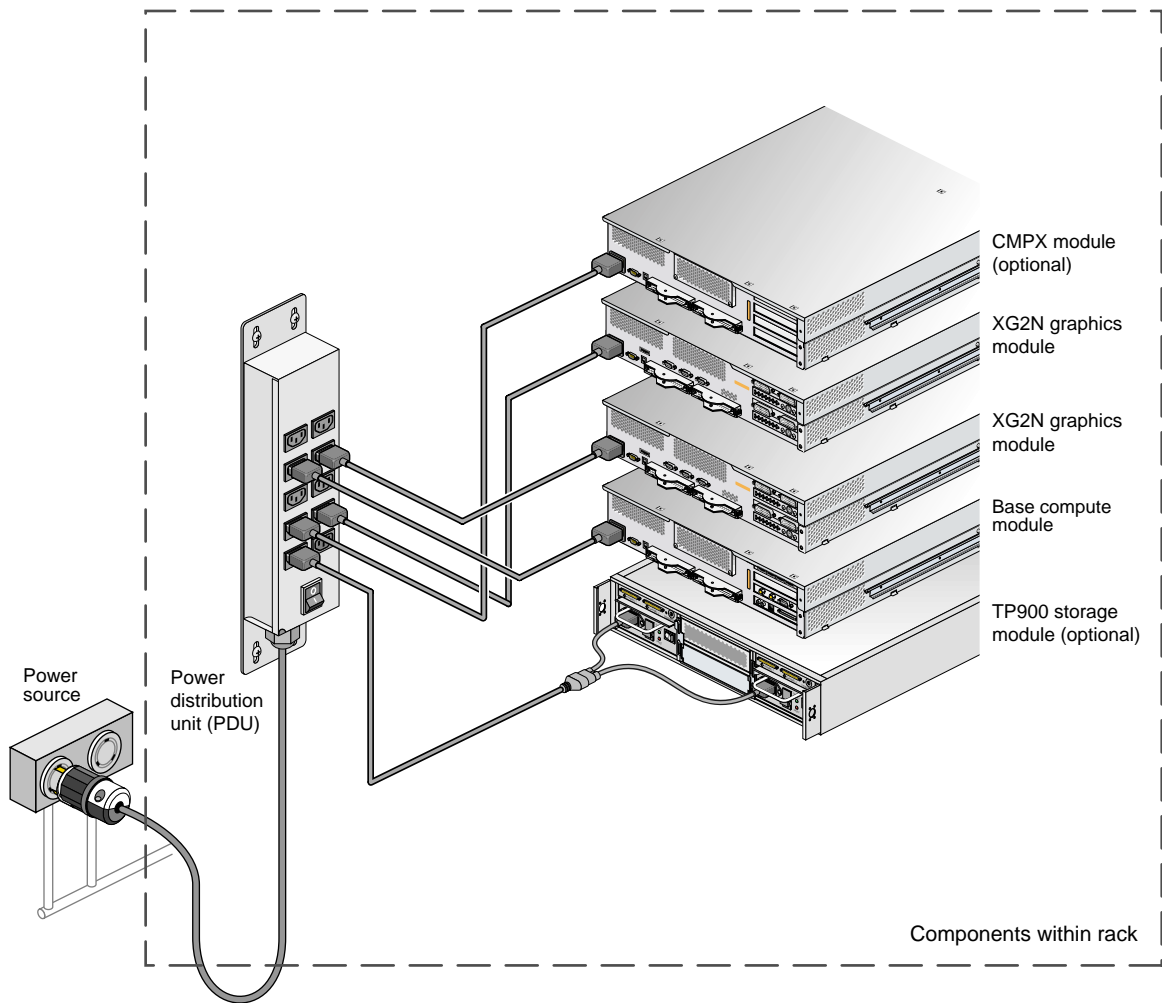
Figure 1-7   Connection Example of a Racked System to Power

NUMAlink Cabling for a Silicon Graphics Prism

If your system was shipped from SGI with the modules already mounted in a rack and the NUMAlink cables pre installed, you should proceed to "Optional SGI ImageSync Cabling" on page 11 to continue.

This section describes the NUMAlink cabling for basic configuration. For additional NUMAlink cabling information, see "System Configurations" on page 26.

Figure 1-8 shows a system with a base compute module and one XG2N graphics module.
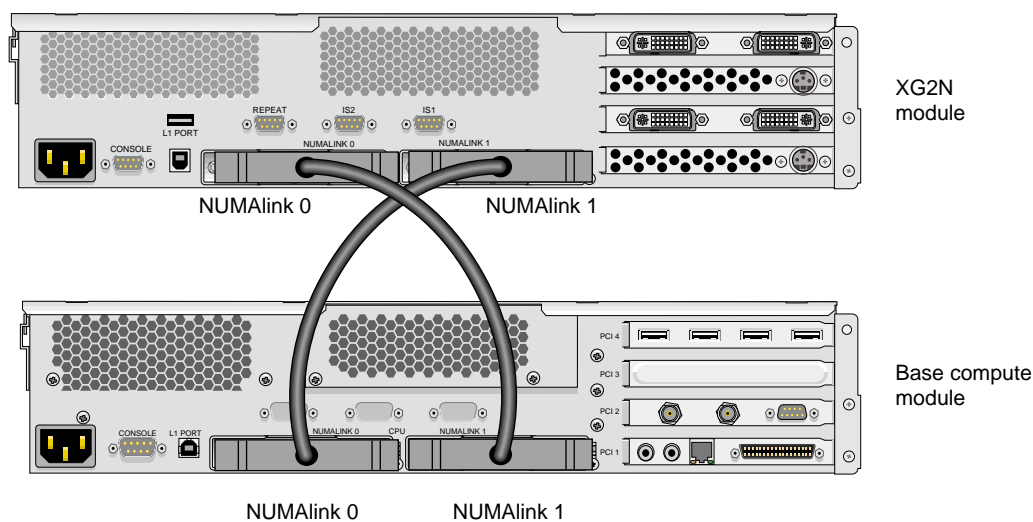
Figure 1-8   NUMAlink Cabling in a Two-module System

Optional SGI ImageSync Cabling

If your system was shipped from SGI with the modules already mounted in a rack and the SGI ImageSync (IS) cables pre installed, you should proceed to the next section "Connecting a Monitor" on page 13 and continue with the setup.

If your system did not come with SGI ImageSync cables pre installed, use the following steps to install them:

1. Connect the DB9 connector on the IS card in your base compute module to the SGI ImageSync DB9 connector (labeled **IS1**) located above the **NUMAlink 1** connector on the back of the XG2N using an SGI ImageSync DB9-to-DB9 cable.

2. Daisy-chain the SGI ImageSync cable from this first XG2N to any additional graphics modules in your system by connecting an ImageSync DB9-to-DB9 cable from the **Repeat** connector (located above the **NUMAlink 0** connector) and routing it to the ImageSync connector (labeled **IS1**) on the back of the next XG2N graphics module, as shown in Figure 1-9.

Note: Figure 1-9 is intended only to show an example of how the ImageSync cables connect. Your system may be configured differently.
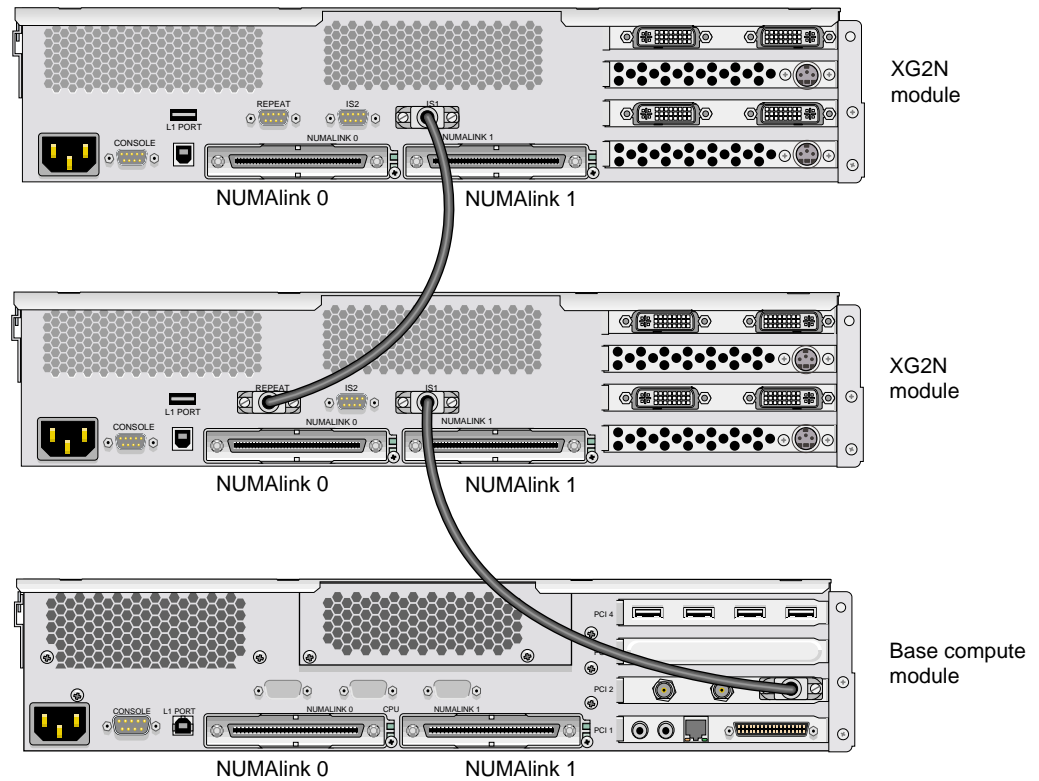
Figure 1-9   Cabling Example of SGI ImageSync Card to Graphics

Connecting a Monitor

Select the monitor on which you want the Linux OS console to appear, and connect this monitor to the pipe 0, channel 1 connector on your graphics module. Pipe 0, channel 1 is located in the lower left section of the PCI graphics area on the back of your primary system graphics module. See Figure 1-1 for reference.

If you have a single XG2N graphics module, pipe 0 will be located in that module.

If you have multiple graphics bricks, pipe 0 will be located in the XG2N module located directly above the base compute module. The Silicon Graphics Prism platform comes with a choice of optional monitor types and sizes.

⚠ **Caution:** Ensure that the electrical rating on the monitor label matches the outlet voltage range (100–120 VAC or 220–240 VAC) before you plug in the monitor.

If you use a monitor in locations that do not have the appropriate outlets, contact your SGI system support engineer (SSE) before plugging in the monitor power cable.

Plug in and turn on your monitor or display as described in the documentation that accompanies your monitor.

Connecting a Keyboard and Mouse

Connect a USB keyboard and a USB mouse to your system, as shown in Figure 1-10. The USB port closest to the side of the chassis is port 1.

Note: Some system configurations may be limited to the use of one keyboard and mouse set. Check with your SGI sales or service representative if you have questions regarding the use of multiple keyboards and mice with your system.

Figure 1-10 shows the USB PCI card in bus 2, slot 2 (the top PCI slot) of the base module, which is the preferred location for this card. Some configurations may have the USB card in a different PCI slot in the host system.

Note: If your keyboard/mouse pair will be further than 10 feet (3 meters) away from the USB ports on the system, please refer to the section "Optional Keyboard Extender" on page 34 for details about how to connect these devices through a USB extender.
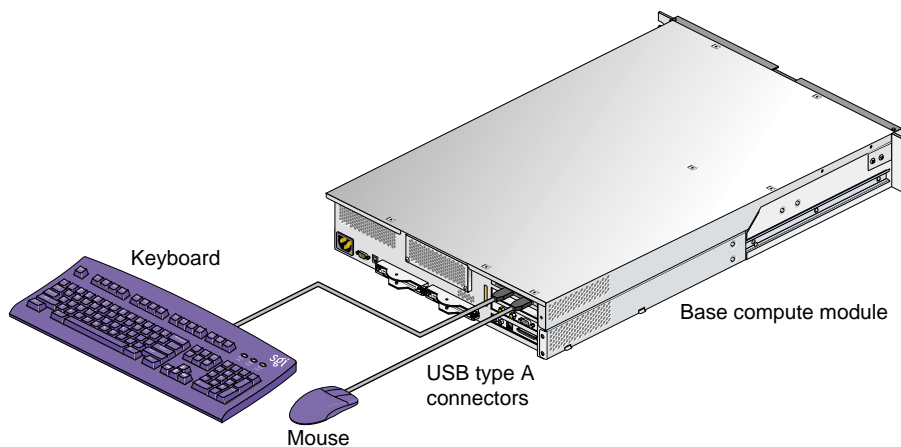
Figure 1-10   Connecting the USB Keyboard and Mouse to the System Compute Module

Connecting the Multi-port Serial Cable

A multi-port serial cable that converts the serial output on the IO10 board into four DB-9 serial cables is included with each base compute module containing an IO10.

"Connecting the Multi-port Serial Cable" on page 14 shows an example connection.

Note: Leave the multi-port serial adapter cable unplugged if you are not using it for optional serial connections.
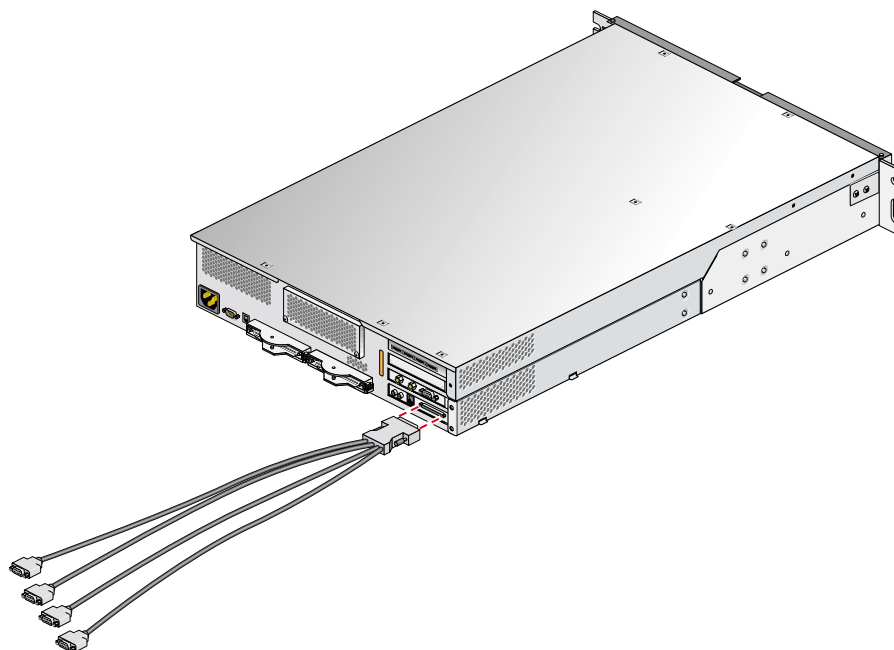
Figure 1-11   Multi-port Adapter Serial Cable

Connecting an L1 System Console

Connect a serial terminal to the Level 1 console connector on the base compute module as shown in Figure 1-12.

Base compute
module

DB-9
connector

Serial terminal

DB-9 serial port
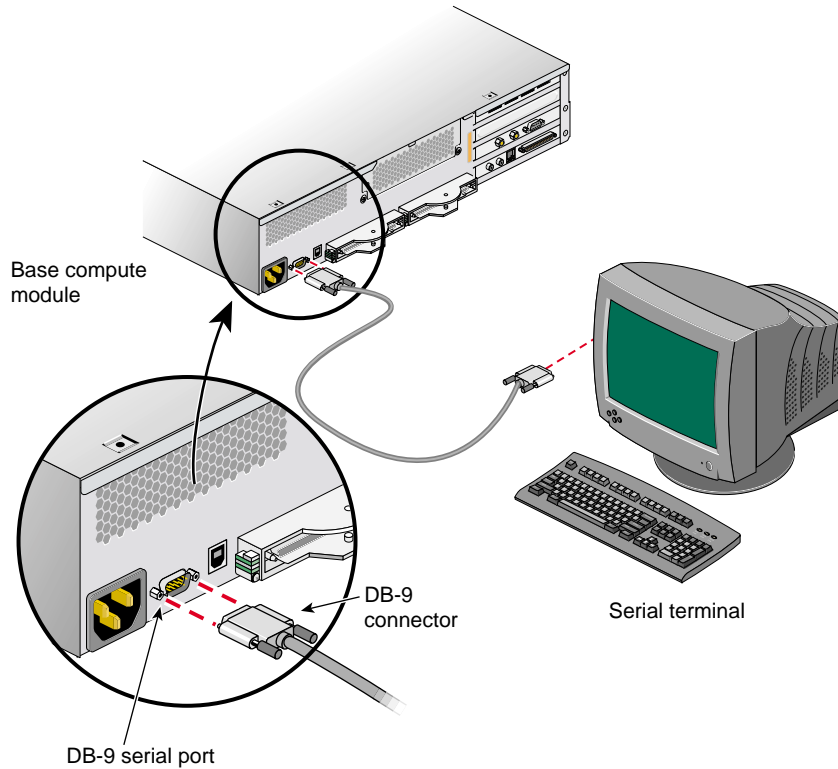
Figure 1-12   Connecting a Serial Terminal to the Base Compute Console Port

Power-On the System

If your Silicon Graphics Prism does not have a a local system console, you can manually power it on and off by using the power buttons. To power on your system or an individual module manually, follow these steps:

1. If the monitors and other peripherals are equipped with voltage select switches, verify that they are set for the appropriate AC voltage and plug them in. Note that they are normally plugged into power sources outside a rack system.

2. Turn on the circuit breaker switch of the PDU if applicable.

3. Press the power buttons on each of the modules that you want to power on in the following order:

    a. For the optional TP900 storage module, press the rear-panel power button to ON (**I**).

    b. Power on all the XG2N graphics modules.

    c. Power on any optional compute expansion or CMPX modules.

    d. Power on the base (system) compute module last.

To boot the system from the L1 interface you must have a console connected to the system via the console port on the base compute module. The serial connection parameters are:

*   38400 baud

*   1 stop bit

*   no parity

Use the following command at the L1 prompt to bring up all the system modules:

001c01-L1> * power up

If the system does not boot, recheck all power, and cable connections to be sure they are properly plugged in. See Chapter 4, "Troubleshooting" for additional trouble shooting tips and L1 messaging.

To contact the SGI Customer Service Center, call 1-800-800-4SGI, or visit http://www.sgi.com/support/customerservice.html. From outside the United States contact your local SGI sales office.

Verifying System Connections

Once your Silicon Graphics Prism Visualization System is installed in a rack (or otherwise situated), is cabled together via NUMAlink, and is powered on, you should verify that all modules are being seen by Linux + ProPack. To do this, follow these steps:

1. From a Linux prompt, type **hinv <Enter>**

2. To verify the connection of all the graphics modules check the output for lines similar to the following examples:
   VGA compatible controller: ATI Technologies Inc Radeon R350 NG [FireGL X2] (rev *xxx*). on pci11.01.0
   Display controller: ATI Technologies Inc Radeon R350 [FireGL X2] (Secondary) (rev *xxx*). on pci11.01.1
   VGA compatible controller: ATI Technologies Inc Radeon R350 NG [FireGL X2] (#2) (rev *xxx*). on pci12.01.0
   Display controller: ATI Technologies Inc Radeon R350 [FireGL X2] (Secondary) (#2) (rev *xxx*). on pci12.01.1

3. To verify all graphics boards in the system are seen, count the number of entries in the output similar to the example output in step 2. Each two-line entry represents one graphics pipe.

4. If the output does not show all the graphics pipes present, recheck all power, and cable connections to be sure they are properly plugged in. See Chapter 4, "Troubleshooting" for additional trouble shooting tips and L1 messaging.

To confirm the presence of the graphics boards in the system using a Linux line command, use the following:

1. Enter **lspci** from a Linux prompt.

2. Check the output for information similar to the following:
   11:01.0 VGA compatible controller: ATI Technologies Inc Radeon R350 NG [FireGL X2] (rev *xx*)
   11:01.1 Display controller: ATI Technologies Inc Radeon R350 [FireGL X2] (Secondary) (rev *xx*)
   12:01.0 VGA compatible controller: ATI Technologies Inc Radeon R350 NG [FireGL X2] (rev *xx*)
   12:01.1 Display controller: ATI Technologies Inc Radeon R350 [FireGL X2] (Secondary) (rev *xx*)

3. If the output does not show all the graphics pipes present, recheck all power, and cable connections to be sure they are properly plugged in. See Chapter 4, "Troubleshooting" for additional trouble shooting tips and L1 messaging.

Powering Off Manually

To power off your system manually, follow these steps:

⚠️ **Caution:** If you power off the system before you halt the operating system, you can lose data.

1. Shut down the operating system by entering the following command:

   `# init 0`

2. Press the power buttons or power switches on each of the modules that you want to power off. You may power off the modules in any order:

   • To power off the TP900 storage module, press the power button on its rear panel to the OFF (**O**) position.

To power off the base compute, XG2N, expansion compute, or CMPX module, press the power button on the front panel of each module.

System Overview and Options

This chapter provides an overview of the physical and architectural aspects of your Silicon Graphics Prism platform. System configurations and components are described and illustrated. This chapter includes the following sections:

- "Physical Features" on page 22
- "Functional Architecture" on page 24
- "System Configurations" on page 26
- "System Components" on page 29
- "Optional Tape Devices" on page 39
- "Optional DM8 Audio Board" on page 40
- "Optional L2 Controller" on page 43
- "SGI Racks for Silicon Graphics Prism" on page 44

Physical Features

The Silicon Graphics Prism platform is the latest advancement in the SGI NUMAflex approach to Linux based modular visualization and computing. It is designed to deliver maximum sustained performance in a compact system footprint. Independent scaling of graphics pipes, computational power, I/O bandwidth, and in-rack storage lets you configure a system to meet your unique visualization and computational needs.

The system can be expanded from a two-module system with two graphics pipes, up to 48 GB of memory, and 4 PCI/PCI-X slots to a higher-performance system that contains eight processors, 96 GB local DIMM memory, and eight PCI/PCI-X slots. Future configurations will offer larger amounts of graphics pipes, processors, memory and PCI expansion. Note that many of the available PCI/PCI-X slots may be used to support I/O, USB, ImageSync, or optional audio or SCSI cards needed for the system.

For most configurations, the visualization system is housed in one 17U rack or one 39U rack as shown in Figure 2-1; however, for small system configurations, the system modules can be placed on a table top.

Systems that are housed in 17U racks have a maximum weight of approximately 610 lb (277 kg). The maximum weight of systems that are housed in 39U racks is approximately 1,366 lb (620 kg). The racks have casters that enable you to remove the system from the shipping container and roll it to its placement at your site.

Check with your SGI service representative for additional physical planning information or documentation that may be available.

For more information about the technical specifications of your system and individual modules, see Appendix A, "Technical Specifications," in this manual.
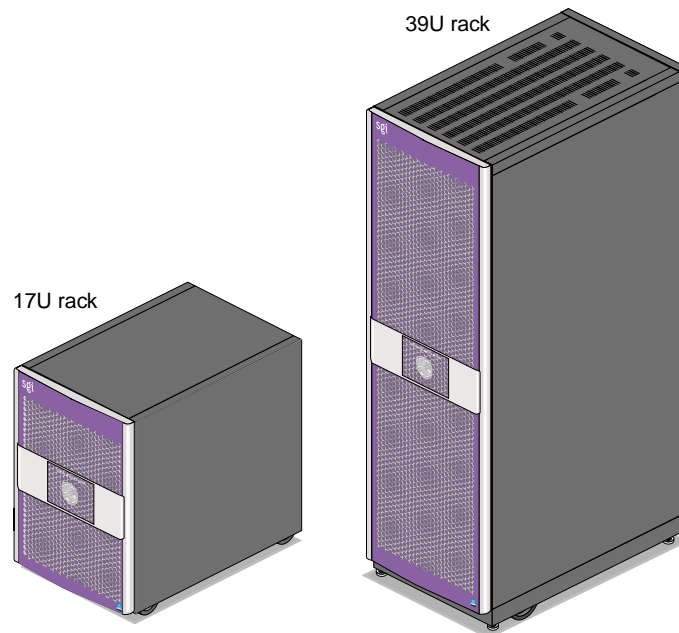
Figure 2-1   Example of Silicon Graphics Prism Rack Systems

Functional Architecture

The Silicon Graphics Prism is based on the SGI NUMAflex architecture, which is a shared-memory system architecture that is the basis of SGI HPC servers and supercomputers. The NUMAflex architecture is specifically engineered to provide technical professionals with superior performance and scalability in a design that is easy to deploy, program, and manage. It has the following features:

Shared access of processors, memory, and I/O. The Super Hub (SHub) ASICs and the NUMAlink interconnect functions of the NUMAflex architecture enable applications to share processors, memory, and I/O devices.

- Each SHub ASIC in the system acts as a memory controller between processors and memory for both local and remote memory references.

- The NUMAlink interconnect channels information between all the modules in the system to create a single contiguous memory in the system of up to 96 GB and enables every processor in a system direct access to every I/O slot in the system.

Together, the SHub ASICs and the NUMAlink interconnect enable efficient access to processors, local and remote memory, and I/O devices without the bottlenecks associated with switches, backplanes, and other commodity interconnect technologies.

System scalability. The NUMAflex architecture incorporates a low-latency, high-bandwidth interconnect that is designed to maintain performance as you scale system computing, I/O, and storage functions. For example, the computing dimension in some system configurations can range from 4 to 8 processors in a single system image (SSI).

Efficient resource management. The NUMAflex architecture is designed to run complex models and, because the entire memory space is shared, large models can fit into memory with no programming restrictions. Rather than waiting for all of the processors to complete their assigned tasks, the system dynamically reallocates memory, resulting in faster time to solution.

The components shown in Figure 2-2 are 2U modules that can be configured as one system using the following devices:

- **Base compute module.** All Silicon Graphics Prism systems contain at least one base compute module that contains the following components:

  - Two Intel Itanium 2 processors; each processor has integrated L1, L2, and L3 caches

- Up to 24 GB of local memory

- Four PCI/PCI-X slots

- One four-port USB card that comes factory-installed in the top PCI/PCI-X slot

- One IO10 PCI card that comes factory-installed in the lowermost PCI/PCI-X slot

---

Note: Each system or partition requires one IO10 PCI card. This card provides the base I/O functionality for the system. Additional IO10 cards are required if you want additional serial ATA hard drives and/or DVD-ROM drives in additional modules. These cards must reside in additional base compute modules (one card per module).

The IO10 card has a real time interrupt input port and output port, an Ethernet port, and a multi-port serial adapter connector. The IO10 card is also needed to support a base module's serial ATA disk drive(s), and DVD-ROM.

Note: The RT interrupt input and RT interrupt output functionality of the IO10 PCI card is not supported under SGI Linux + ProPack.

- One Super-Hub (SHUB) ASIC (the crossbar between the processors, local memory, the network interface, and the I/O interface).

- **XG2N graphics module.** All systems contain at least one graphics module that has the following components:

  - Two graphics pipes

  - A node board with two Intel Itanium 2 processors; each processor has integrated L1, L2, and L3 caches

  - Up to 24 GB of optional local memory

- **CPU Expansion module.** The difference between the base compute module and the CPU expansion module is that the CPU expansion module does not contain any usable PCI/PCI-X slots, removable media, or disk drives.

- **CMPX module.** The only difference between the base compute module and the CMPX module is that the CMPX module does not contain an IO10 PCI card, removable media, or disk drives.

When the system consists of a base compute module and XG2N graphics module only, the maximum number of processors is 4 and the maximum amount of memory is 48 GB. To increase the number of graphics pipes, processors and/or memory in the system, the base compute module can connect to additional XG2N, CPU expansion and/or CMPX modules via direct NUMAlink.

System Configurations

Table 2-1 lists the minimum and maximum ranges of the configurable items for the Silicon Graphics Prism.

Table 2-1   Silicon Graphics Prism Configuration Ranges

	Configuration Minimum	Configuration Maximum[a]
Compute/expansion modules	1	3
Graphics modules	1	2
Processors	4	8
1.4 GHz 1.5 GHz 1.6 GHz	~ 5.6 GFLOPS ~ 6.0 GFLOPS ~ 6.4 GFLOPS	(See note a.)
Memory capacity	4 GB	96 GB
Internal disk storage	One or two serial ATA disks	6 serial ATA disks
Internal PCI/PCI-X slots[b]	1	7
Storage device	None	Customer-configurable
Compute rack	None	1
Optional L2 controller	None	1

a. Maximum configurations and peak performance will change over time based on new technology available; check with your SGI sales or service representative for the latest available options and configurations.

b. Each base compute module has four internal PCI/PCI-X slots; however, one slot is required for the base IO PCI card, one for the ImageSynch card, and one for USB. Therefore, the number of available slots in the base compute module is reduced.

The basic Silicon Graphics Prism uses a type of network configuration that is referred to as a ring topology. As the name implies, the network connection between the base compute module, graphics module, and optional CPU expansion or CMPX modules forms a ring. A message is passed around the ring until it reaches its destination. The data flow of this ring topology flows in both directions, enabling the modules to have direct connection to two other modules and providing an alternative path when a connection fails between two modules.

The bisection bandwidth of the ring depends on the number of modules on the ring; the bisection bandwidth is greatest when there are only two modules on the ring. See Table 2-2.

Table 2-2   Bisection Bandwidth of Ring Topology

Module Count	Bisection Bandwidth
2	3.2 GB/s
3	2.13 GB/s
4	1.6 GB/s

Figure 2-2 provides an example of the two module ring topology.
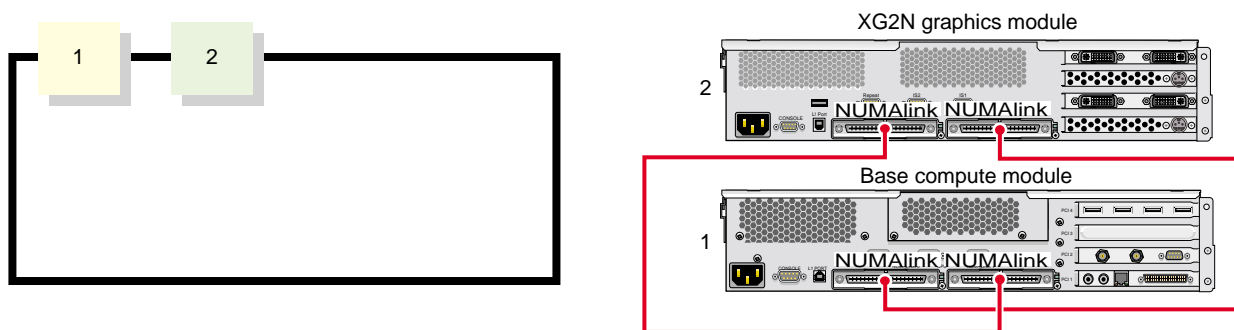
Figure 2-2   Two Module Ring Topology Example

Figure 2-3 on page 28 shows an example of a three module ring topology.

Figure 2-4 on page 28 shows an example of a four module ring topology.
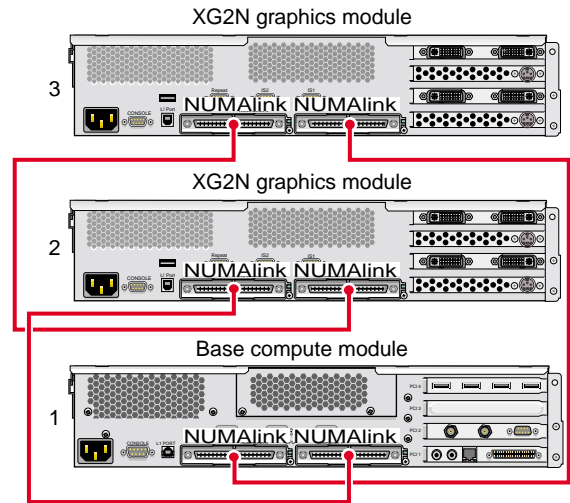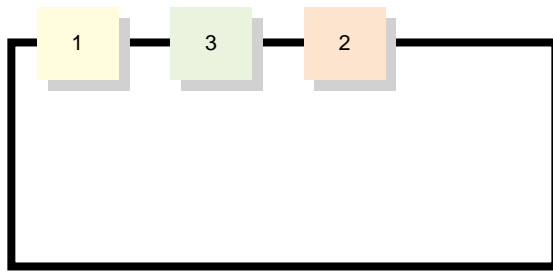
Figure 2-3   Three Module Ring Topology Example
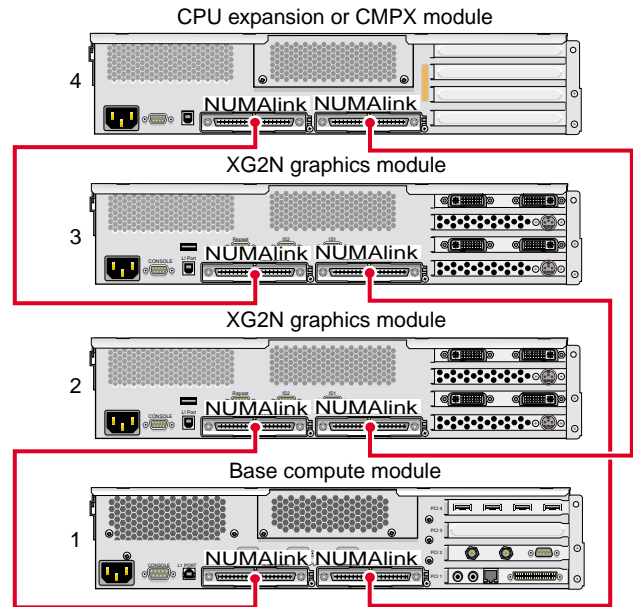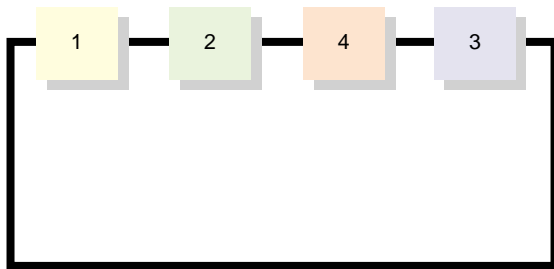
Figure 2-4   Four Module Ring Topology Example

System Components

This section briefly describes the standard and optional major system components of a Silicon Graphics Prism, in the following subsections:

Base Compute Module

The base compute module is a 2U AC-powered device that consists of the following:

- Two Intel Itanium 2 processors; each processor has integrated L1, L2, and L3 caches

- Up to 24 GB of DIMM memory.

- One to four PCI/PCI-X cards.

    Note: At least one base compute module comes factory-installed with an IO10 card in the bottom PCI/PCI-X slot and a four-port USB card in the upper slot.

- One or two sled-mounted serial ATA disk drives (at least one disk drive is required in the system). The serial ATA disk drive(s) and the DVD-ROM require an IO10.

- DVD-ROM.

Each base compute module also contains an L1 controller that provides the following:

- Controls and sequences power.

- Controls and monitors the environment.

- Initiates a reset.

- Stores identification and configuration information.

Figure 2-5 shows the front and rear views of a base compute module.

Front view
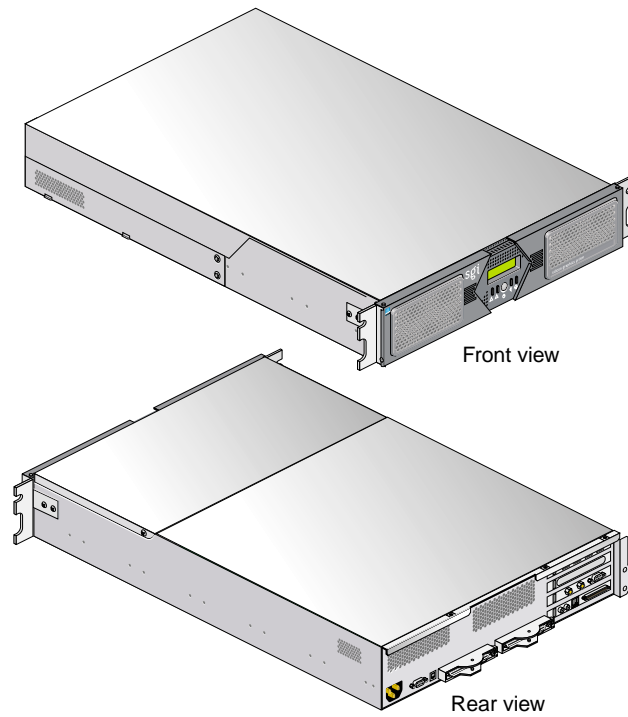
Rear view

Figure 2-5   Front and Rear Views of Base Compute Module with IO10 PCI Card

XG2N Graphics Module

The XG2N graphics module (see Figure 2-6) is a 2U AC-powered device that consists of the following:

- Two graphics pipes

- Node board with two 64-bit processors

- Up to 24 GB of optional local DIMM memory (4, 8, or 12 DIMMs)

- One L1 controller that provides the following services:

    - Controls and sequences power.

    - Controls and monitors the environment.

    - Initiates a reset.

    - Stores identification and configuration information.

Front view
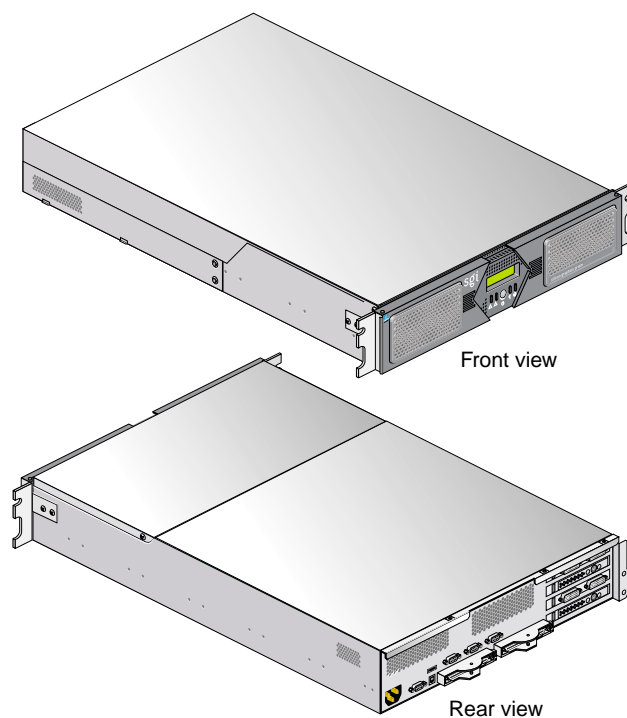
Rear view

Figure 2-6   Front and Rear Views of XG2N Module

CPU Expansion Module

The CPU expansion module is a 2U AC-powered device that consists of the following:

- Zero or two Intel Itanium 2 processors; with integrated L1, L2, and L3 caches

- Up to 24 GB of optional DIMM memory.

- One L1 controller that provides the following services:

    - Controls and sequences power.

    - Controls and monitors the environment.

    - Initiates a reset.

    - Stores identification and configuration information.

Figure 2-7 shows the front and rear views of a CPU expansion module.
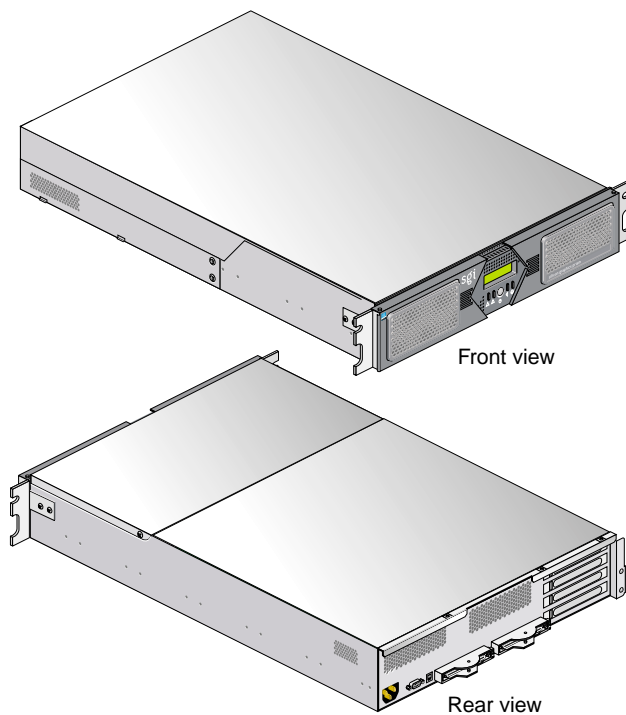
Figure 2-7   Front and Rear Views of CPU Expansion Module

CMPX Module

The CMPX module (see Figure 2-8) is a 2U AC-powered module that offers:

- Zero, or two optional processors
- Up to 24 GB of optional DIMM memory
- Four PCI/PCI-X slots (see Figure 2-8)
- One L1 controller that provides the following services:
  - Controls and sequences power.
  - Controls and monitors the environment.
  - Initiates a reset.
  - Stores identification and configuration information.

Front view
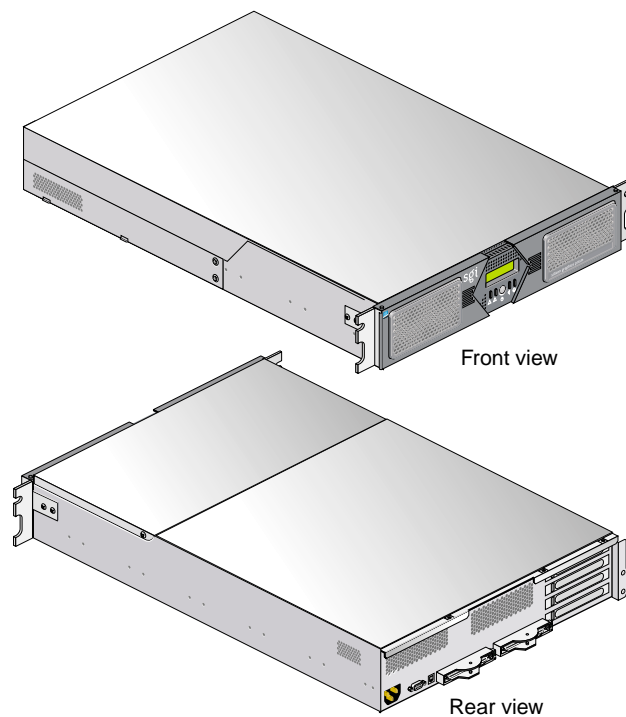
Rear view

Figure 2-8   CMPX Option Module Front and Rear Views

Optional Keyboard Extender

You can plug the keyboard and mouse directly into the USB connectors of a compute module (see "Connecting a Keyboard and Mouse" on page 13), or you can use an optional USB extender that allows you to place the keyboard and mouse up to 328.68 feet (100 m) from the compute module, see Figure 2-9 on page 35.

Note: The optional local extender (LEX) receives AC power from the system power source or power distribution strip (PDS) in the rack. A 6-ft. (1.82-m) adapter cable connects the extender's power adapter to the PDU. The REX requires AC power from a source within a 6-ft. (1.82-m) range from the keyboard/mouse connection point.
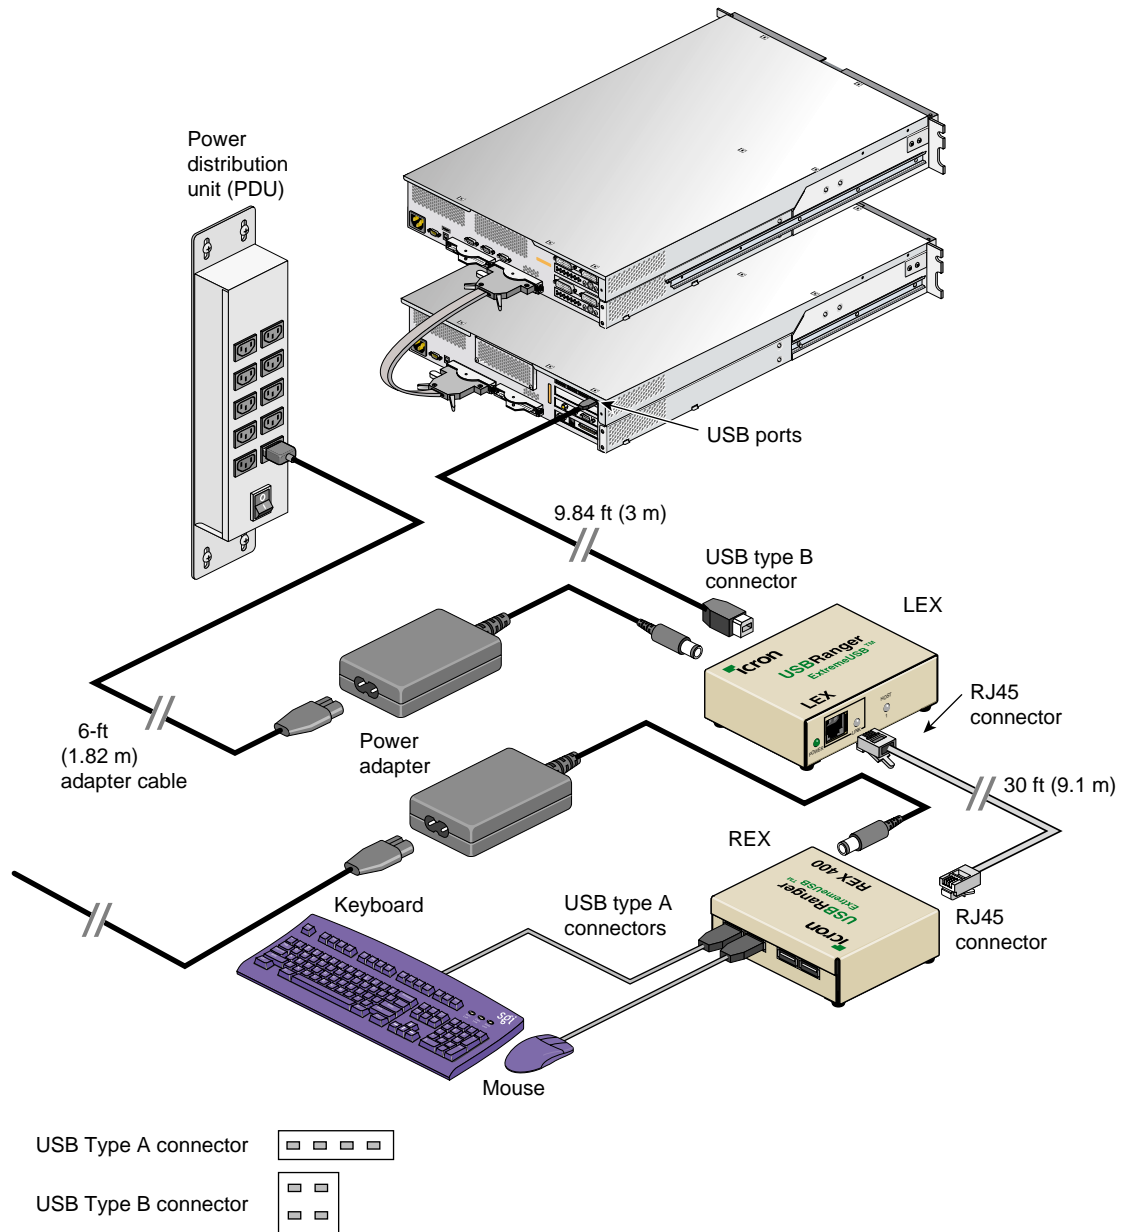
Figure 2-9   Keyboard and Mouse Connected via Optional USB Extender

Storage Expansion

A base compute module contains an IO10 card and two disk-drive bays. You can add additional storage to the system as follows:

- For a SCSI (small computer system interface) JBOD (just a bunch of disks) solution, SGI offers the TP900 storage module, that can be added to base compute modules or CMPX expansion modules (with optional SCSI PCI card).

- For a Fibre Channel solution that supports both JBOD and RAID configurations, SGI offers the 2Gb SGI TP9100 storage system.

- The Silicon Graphics Prism also supports a number of tape devices; check with your SGI sales or support representative for available options.

The various storage devices are discussed in the subsections that follow.

SGI TP900 Storage Module

The TP900 storage module, shown in Figure 2-10, is a 2U-high 8-drive storage system that provides compact, high-capacity, high-availability JBOD storage. The enclosure backplane connects the 8 drives on one SCSI bus. As an option, the storage module can also be configured on two SCSI buses (2 strings of 4 drives).

This storage module has the following features:

- It mounts in a standard 19-inch rack; it is available in factory-installed configurations.

- It uses SCSI Parallel Interface 3 (SPI-3) capable Low Profile (1-inch high) 3.5-inch disk drives.

- Its drive carriers accept SGI-qualified 10,000- or 15,000-RPM SCSI disk drives.

For more information about the TP900 storage module, see *SGI Total Performance 900 Storage System User's Guide* (007-4428-00*x*).

Front view

Rear view
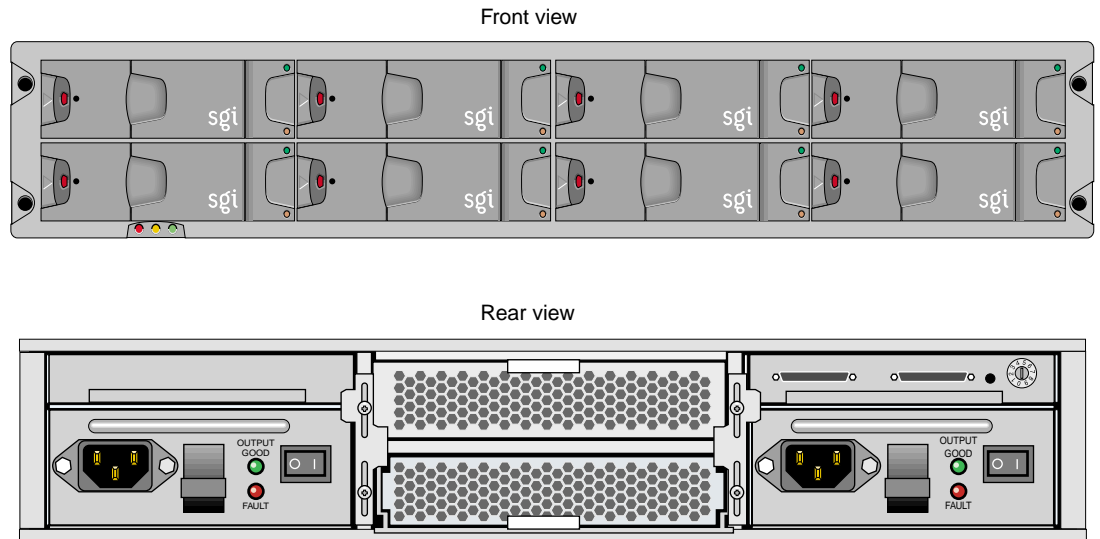
Figure 2-10   SGI TP900 Storage Module

2Gb SGI TP9100 Optional Storage System

The 2Gb SGI TP9100, shown in Figure 2-11, is an affordable, entry-level RAID storage array that is easily expandable and comes in either a deskside tower or a rackmounted configuration. You can start with a basic JBOD configuration and later add RAID controllers, or you can start with a RAID configuration.

The 2Gb SGI TP9100 storage system connects to base compute and/or CMPX modules via an optional Fibre Channel PCI card. For more information about the SGI TP9100 storage system, see *SGI Total Performance 9100 (2 Gb TP9100) Storage System User's Guide* (007-4522-00*x*).

Rackmounted configuration

Deskside tower configuration
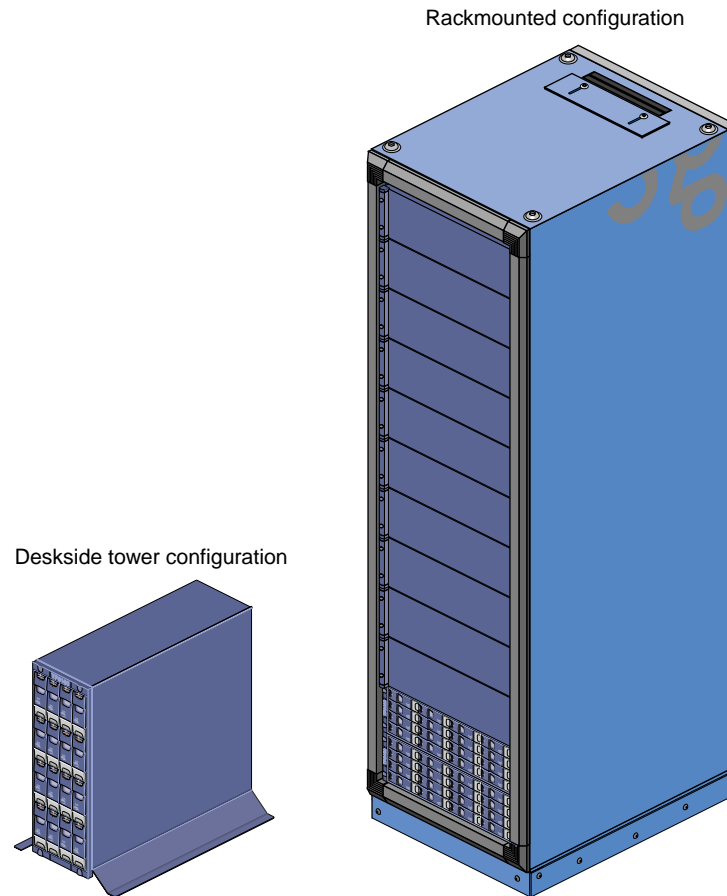
Figure 2-11   2Gb SGI TP9100 Storage System

Optional Tape Devices

The Silicon Graphics Prism supports optional or third-party tape drives. For current tape solutions for your system visit:

http://sales.corp.sgi.com/products/storage/tape.html

or check with your SGI sales or service representative.

Optional DM8 Audio Board

The DM8 is a half size PCI expansion board that provides audio expansion via the PCI bus. It connects to consumer and professional audio and video equipment via industry-standard interfaces. Note that not all connector functionality is supported under Linux. The DM8 interface panel is shown in Figure 2-12.
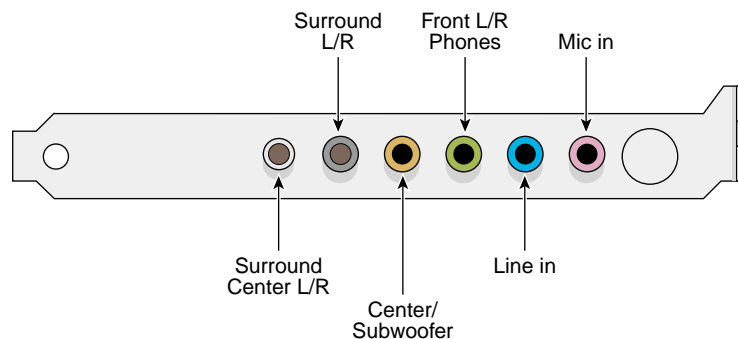
Figure 2-12   DM8 Interface Panel and Connectors

Board Installation

Your DM8 audio card is tested for operation in specific SGI visualization systems. Select SGI Linux based systems support use of a single DM8 audio card (check with your SGI sales or support representative for functional confirmation). It is recommended that it not be removed and installed in other SGI systems.

SGI Linux based visualization systems support only one DM8 option card. Under most circumstances you should always leave the DM8 installed in the original factory-configured slot. If it does become necessary to either move or replace the board, use the information on removing and replacing PCI cards in your system user's guide for step-by-step instructions. If you don't have the printed user guide handy, it is available on the Internet in the SGI Technical Publications Library at the following location:

http://docs.sgi.com

Proceed to the next sections for information on cabling, verifying functionality, and basic operation.

Cables and Connectors

The DM8 audio board has the following connectors on it's rear panel as shown from left to right in Figure 2-12 on page 40:

- Surround Center L/R (silver) a line-level output connector for powered surround Center Left and Right channel speakers (not supported under Linux).

- Surround L/R (black) a line-level output that can connect to powered Surround Left and Right speakers (not supported under Linux).

- Center/Subwoofer (orange) a line-level output that can connect to powered Center and Subwoofer speakers (not supported under Linux).

- Front L/R Phones (green) a jack that connects to powered Left and Right speakers. If you are using a basic two-speaker stereo setup, this is the jack you will use. This is also the jack to use for stereo headphones.

- Line In (blue) a stereo line-level input intended for recording into an audio recording program.

- Mic In (pink) a mic input for connecting a mono microphone for recording, videophone, or voice recognition applications. This jack is a three-conductor TRS-type jack for condenser (two-conductor) and electret (three-conductor) type microphones.

Functional Features

The DM8 PCI audio option has the following functional features:

Line outs:

- Dynamic range of 106 dB (typical, -60 dB input, a-weighted)

- Maximum line output level of 1.26 Vrms (+2 dBV)

- Frequency response of +0.8/-3.0 dB, 20Hz to 80 kHz

Mic in:

- Dynamic range of 90 dB (a-weighted)

- Maximum input level of 1.45 Vrms (3.2 dBV)

Line in:

- Dynamic range of 105 dB (a-weighted)

Sampling frequencies (in kHz):

• Supports sampling frequencies 8, 11.025 16, 22.05, 32, 44.1, or 48kHz

Audio playback formats:

• Output formats supported are:

• 24-bit linear PCM/48kHz playback capability for Linux systems

• 24-bit linear PCM/48kHz recording capability for Linux systems

Connector features:

• Analog connector for powered speakers (headphone jack)

• Stereo 1/8-inch mini analog line input

• Mono 1/8-inch mini analog microphone input (with electret mic power)

Troubleshooting Tips

Use the following information to help correct any functional problems with the DM8 audio PCI option card:

If the sound is distorted or a channel is not producing sound, try the following:

• Check all cable connections between the audio card and speakers

• Confirm that any powered speakers are powered on

If the audio is not working at all, try these solutions:

• Reboot the system

• Reseat the PCI audio card

If none of these solutions solve the problem, contact your SGI customer support representative or local service provider for additional help.

Optional L2 Controller

The L2 controller (see Figure 2-13), which is an optional component, is a system-level controller that performs the following functions:

- Controls resource sharing.

- Controls L1 controllers.

- Maintains system configuration and topology information.

- Enables remote maintenance.

- Routes data between upstream and downstream devices, as follows:

    – Upstream devices (for example, the system console) provide control for the system, initiate commands for the downstream devices, and act on the messages that they receive from downstream devices.

    – Downstream devices (for example, L1 controllers) perform the actions specified by the L2 controller commands, send responses to the L2 controller that indicate the status of the commands, and send error messages to the L2 controller.

All components within a given system that have an L1 controller can connect to the L2 controller (see Figure 2-13). For example, base compute, XG2N, CPU expansion, and CMPX modules can connect to the L2 controller directly. Other available components such as compositors or TP900 storage modules do not have L1 controllers and do not connect to the optional L2.
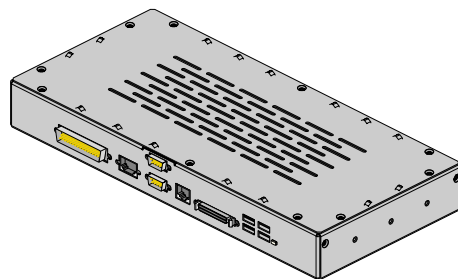
Figure 2-13   Optional L2 Controller Unit

Power Components

The Silicon Graphics Prism platform can consist of the following power components:

- **One or two power distribution units (PDUs)**. The second PDU is added to the system only when more than 10 AC power receptacles are needed within the rack.

  The PDU inputs AC voltage from an external power receptacle and it can output AC voltage to the base compute modules, XG2N graphics modules, CPU expansion modules, CMPX modules and TP900 storage modules.

  See Figure 1-7 on page 10 for an example.

SGI Racks for Silicon Graphics Prism

The system is offered in two rack types: a short rack and a tall rack. The racks are measured in standard units (U); one U is equal to 1.75 in. (4.45 cm). The short rack is a 17U rack (see Figure 2-14). The tall rack is a 39U rack, see Figure 2-17 on page 46 for an example.

The components within the rack are identified by the lowest U number that they occupy. For example, the top (XG2N) module shown in Figure 2-14 is identified as U12 in the short rack.
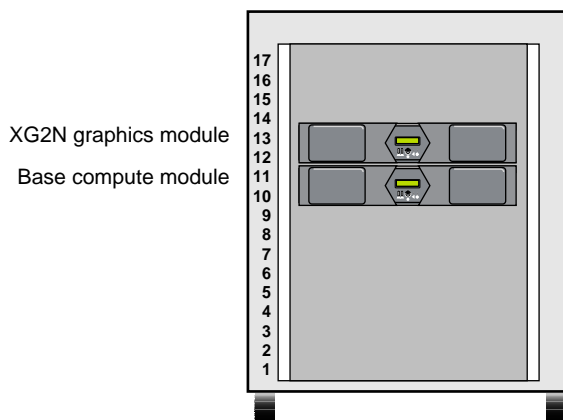
Figure 2-14   Unit Numbering Within Racks

Both rack types are industry-standard 19-inch racks, and they support two types of mounting rails (optional slide rails and shelf rails) that hold the modules within the rack. For example, the base compute, XG2N, CPU expansion, and CMPX modules can use shelf rails or optional slide-mounting rails (see Figure 2-15). The optional TP900 storage modules always use shelf rails, which are two parallel L-shaped mounting rails within the rack (see Figure 2-16).
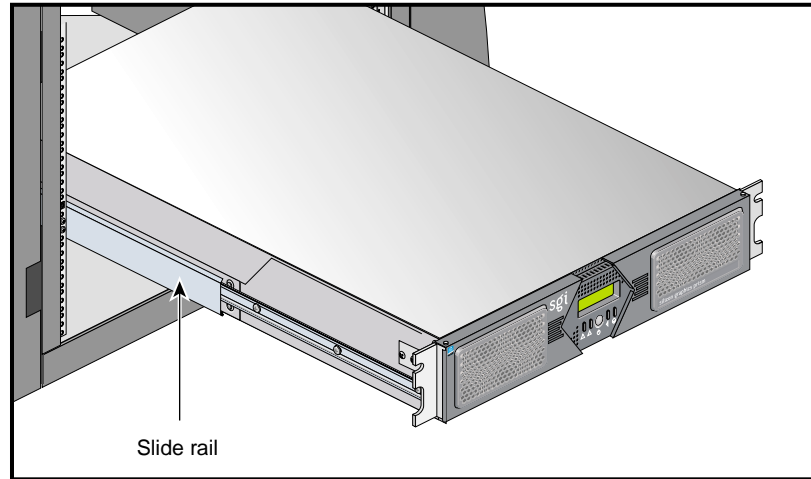
Slide rail

Figure 2-15   Optional Slide Rails

Shelf rails

19 in.
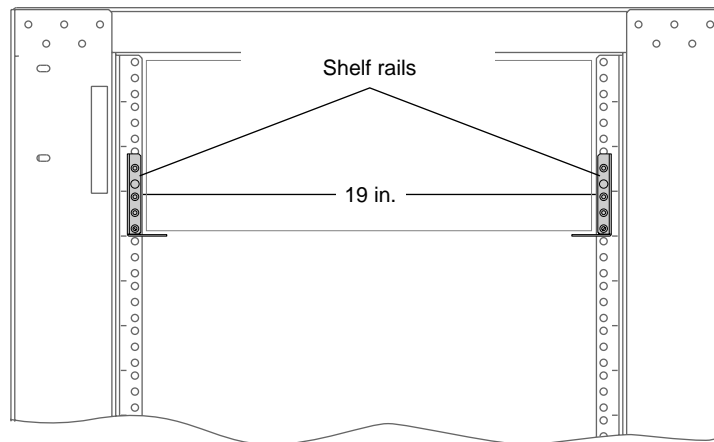
Figure 2-16   Shelf Rails

Both rack types, as shown in Figure 2-17, have front and rear doors that have keylocks to prevent unauthorized access to the system. The racks also have cable entry/exit areas at the bottom of the racks. The 39U racks have cable management hardware in the rear.

Both rack types are mounted on four casters, two of which are swivel casters. The casters enable the rack to be rolled out of a shipping crate and to its placement at your site.

The base of the racks have seismic tie-down attachment points. The base of the tall rack also has leveling pads.
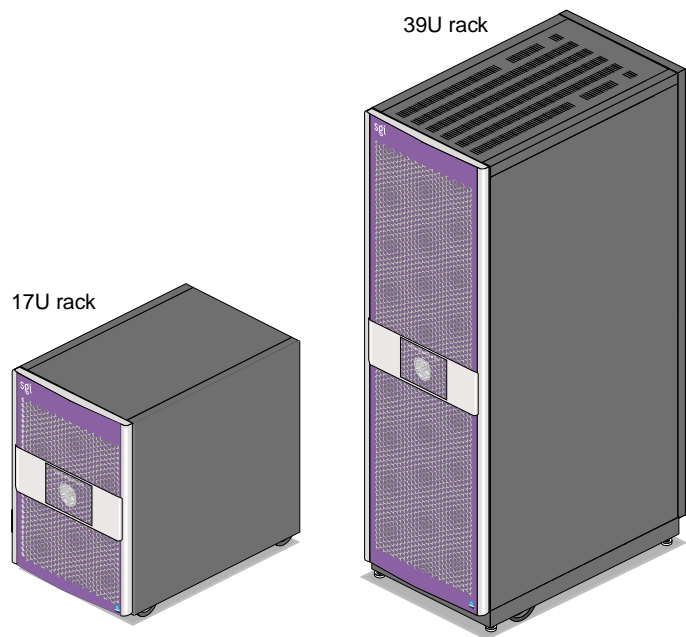
Figure 2-17   Front Views of Short and Tall Racks

Installing and Removing Customer-replaceable Units

This chapter provides safety instructions to follow when using and maintaining your system. It also describes how to install and remove customer-replaceable units (CRUs). This information is covered in the following sections:

- "Safety Instructions" on page 47

- "Accessing the System's Internal Components" on page 49

- "PCI and PCI-X Cards" on page 50

- "Replacing an XG2N Graphics Card" on page 69

- "Disk Drives" on page 76

- "Removing and Replacing Power Supplies" on page 81

- "Memory DIMM Group Placement and Guidelines" on page 85

- "L1 Controller Display" on page 93

Safety Instructions

Before you perform any type of maintenance to your system, read the following safety instructions:

- Follow all warnings and instructions marked on the product and noted in this and other documentation included with the product.

- Unplug this product from the wall outlet before you clean it. Do not use liquid cleaners or aerosol cleaners. Use a damp cloth for cleaning.

- Do not use this product near water.

- Do not place this product or components of this product on an unstable cart, stand, or table. The product may fall, causing serious damage to the product.

- Perforation openings on the cabinet and components are provided for ventilation, reliable operation, and protection from overheating of the product. These ventilation holes must not be blocked or covered. This product should never be placed near or over a radiator or heat register, or in a built-in installation unless proper ventilation is provided.

- This product should be operated with the type of power indicated on the marking label. If you are not sure of the type of power available, consult your dealer or local power company.

- Do not allow anything to rest on the power cord. Do not locate this product where people will walk on the cord.

- Do not use extension cords with your SGI system.

- Never push objects of any kind into this product through cabinet holes because they may touch dangerous voltage points or short out parts that could result in a fire or electric shock.

- Never spill liquid of any kind on the product.

- Do not attempt to service this product yourself except as noted in this guide. Opening or removing covers of internal components may expose you to dangerous voltage points or other risks. Refer all servicing to qualified service personnel.

- Unplug this product from the wall outlet and refer servicing to qualified service personnel under the following conditions:

  - If the power cord or plug is damaged or frayed.

  - If the product has been exposed to rain, water, or other type of liquid.

  - If the product does not operate normally when the operating instructions are followed.

    Note: Adjust only those controls that are covered by the operating instructions, because improper adjustment of other controls may result in damage and will often require extensive work by a qualified technician to restore the product to normal condition.

  - If the product has been dropped or if the cabinet has been damaged.

  - If the product exhibits a distinct change in performance, which indicates a need for service.

- Only qualified service personnel should replace the soldered lithium battery (or batteries) in the Silicon Graphics Prism platform. Please see Appendix C, "Regulatory Specifications and Safety Information" for more information.

- Use only the proper type of power supply cord set (provided with the system) for this unit.

Accessing the System's Internal Components

Most of the procedures in this chapter require the removal of the top cover to access the internal component(s) you are removing or replacing. Use the following steps and illustrations to open and close the cover:

1. To remove the module's top cover and access the internal system components, remove the ten Phillips screws as shown in Figure 3-1.

2. Lift the hinged cover into the fully open position.
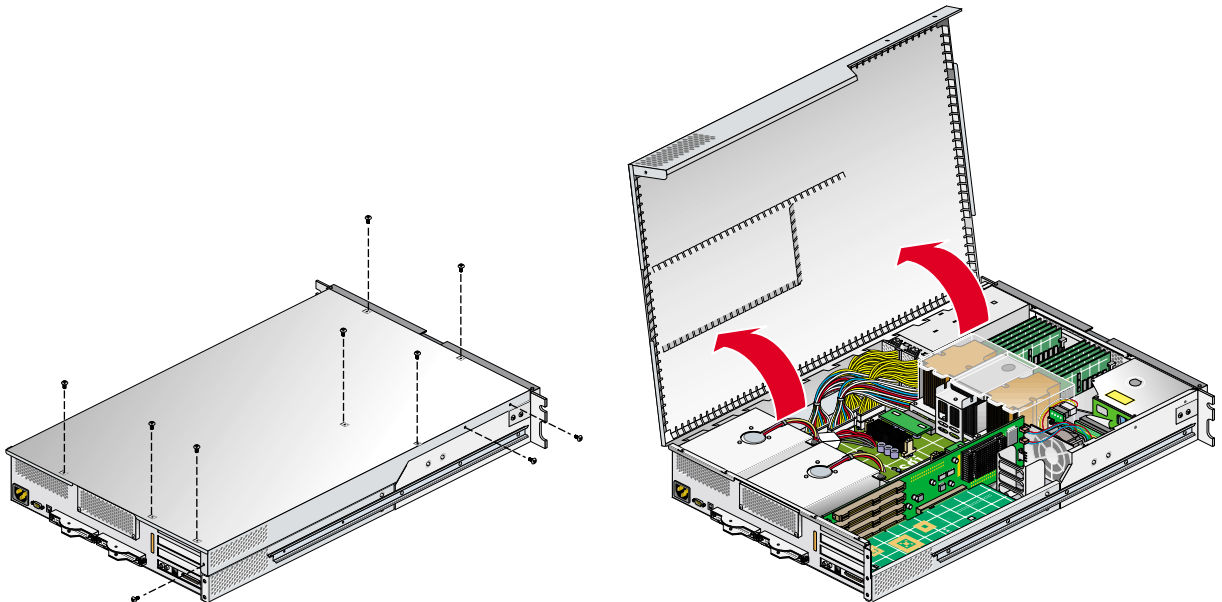
Figure 3-1   Opening the System Cover

3. Close the hinged cover on the system and screw in the ten Phillips screws to secure the cover, as shown in Figure 3-2.
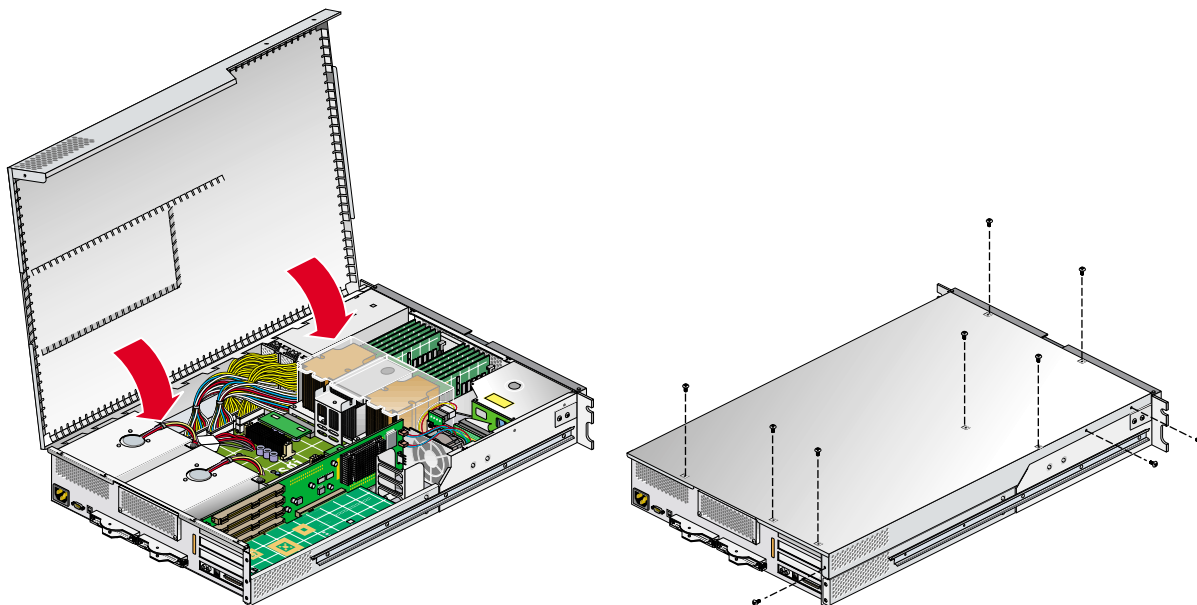
Figure 3-2   Closing the System Cover

PCI and PCI-X Cards

Your base compute module and various optional modules in your system support PCI and PCI-X cards. The following instructions, which describe how to install and remove a PCI or PCI-X card from system base compute module, can be used to install and remove PCI and PCI-X cards from the base compute, expansion compute, or CMPX module. For specific information on removing or replacing XG2N graphics cards, see "Replacing an XG2N Graphics Card" on page 69.

Figure 3-3 shows the location of the PCI and PCI-X card slots in the base compute module.

For an updated list of supported PCI and PCI-X cards, see SGI Supportfolio at http://support.sgi.com.

This section describes the following:

- "Rules for Card Installation and Removal" on page 51
- "Installing a PCI or PCI-X Card" on page 52
- "Removing a Card" on page 60
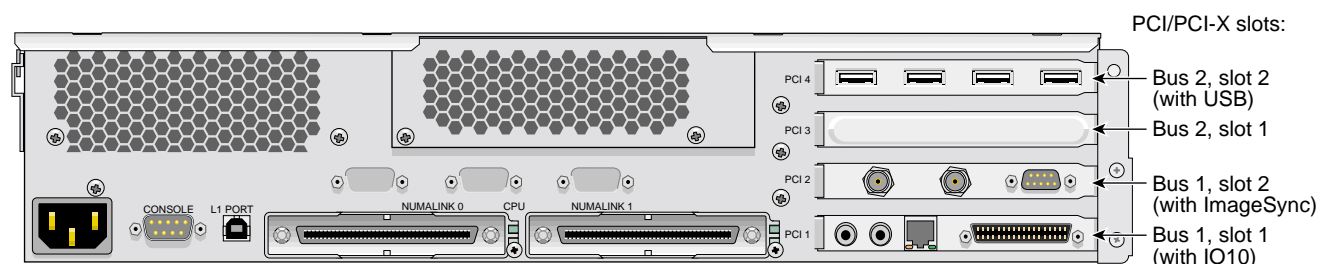
Figure 3-3   PCI and PCI-X Card Slots in the Base Compute Module

Rules for Card Installation and Removal

When installing or removing PCI or PCI-X cards from your system, be aware of the following rules:

- If you operate PCI and PCI-X cards on the same bus at the same time, the PCI-X card will run in PCI mode.

- If you place cards of different speeds on the same bus, the highest-speed card will run at the speed of the slower card. For example, if you have a 100 MHz-capable card in one slot of a bus and a 33 MHz card on the second slot of the same bus, both cards will run at 33 MHz.

- When deciding to install a PCI or PCI-X card, be aware that if an IO10 card (a PCI card that runs at 66 MHz) is installed in your module (in the lowermost slot), and if you want to install a card in the slot immediately above the IO10, the slot will operate only in PCI mode at a speed no faster than 66 MHz.

• Note the following caution when installing or removing a card:

⚠

Caution: Electronic equipment can be irreparably damaged by electrostatic discharge (ESD). Always follow these preventive measures when you handle a system component:

- Remove a component from its antistatic bag only when you are ready to install it.

- If you handle a component before installation, do not place it on surfaces that produce ESD (carpeting, for example) or near devices that create static electricity.

- Attach a static wrist strap to a grounded connection on your system when you install or remove a component.

Installing a PCI or PCI-X Card

To install a PCI or PCI-X card, follow these steps:

1. Power off the Silicon Graphics Prism platform. For power off instructions, see "Powering Off Manually" in Chapter 1.

2. Disconnect all of the cables at the rear of the module.

⚡

Warning: **Components may be hot. To avoid injury, allow the components to cool for approximately five minutes before you proceed with these instructions.**

3. If your module is rackmounted, remove the two screws that secure the module to the front rails of the rack. If the module is mounted with shelf rails, remove the screws at the rear. If your module is not rackmounted, proceed to step 5.

4. If the module is mounted with the optional slide rail kit, pull the module from the rack until it is stopped by the safety latches. If the module uses shelf rails, two people should lift the unit from the rack and place it on a stable work surface.

5. Open the top cover as described in "Accessing the System's Internal Components" on page 49.

6. If you are installing a card into the top two PCI/PCI-X slots, you can install it without completely removing the module from the rack (optional slide rails only).

> **Note:** If you are removing a card from one of the two lower slots (bus 1, slot 1, or bus 1, slot 2), you must remove the module from the rack. For detailed instructions, see "Removing a Rail-mounted Module from a Rack" on page 159.

7. Make sure that you read "Safety Instructions" on page 47, and "Rules for Card Installation and Removal" on page 51 before beginning your card installation.

8. If you are installing a card in one of the two bottommost card slots, you also need to remove the optional chassis rail by unscrewing the five Phillips screws, as shown in Figure 3-4. (If you are not installing a card into one of the bottom-most card slots or your module is not slide-rail mounted, proceed to step 9.)

⚠️ **Caution:** The module may or may not have a factory-installed IO10 card, which always comes installed in the lowermost slot. To prevent damage to your system, only a trained SGI service support engineer can install or remove an IO10 card.
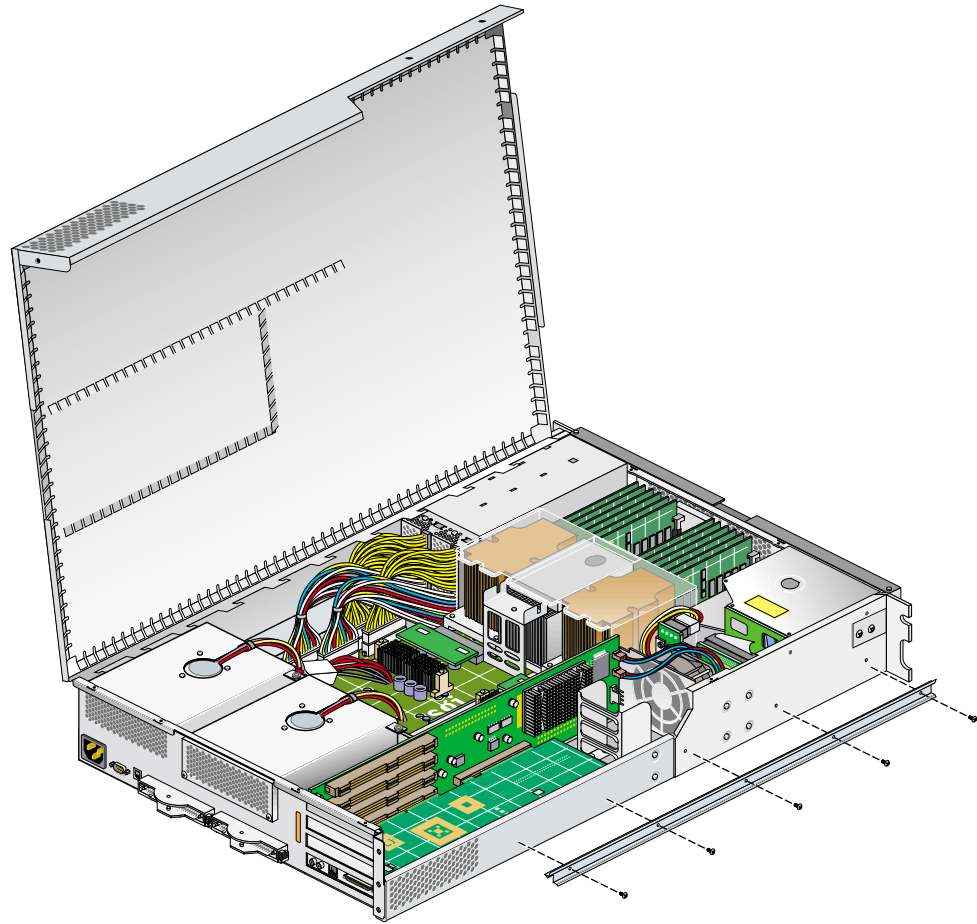
Figure 3-4   Removing an Optional Chassis Rail

9. If you are installing a card in one of the two lowermost card slots, you must remove the lower PCI/PCI-X support bracket that covers the two lowermost slots, as shown in Figure 3-5, by removing the four Phillips screws.
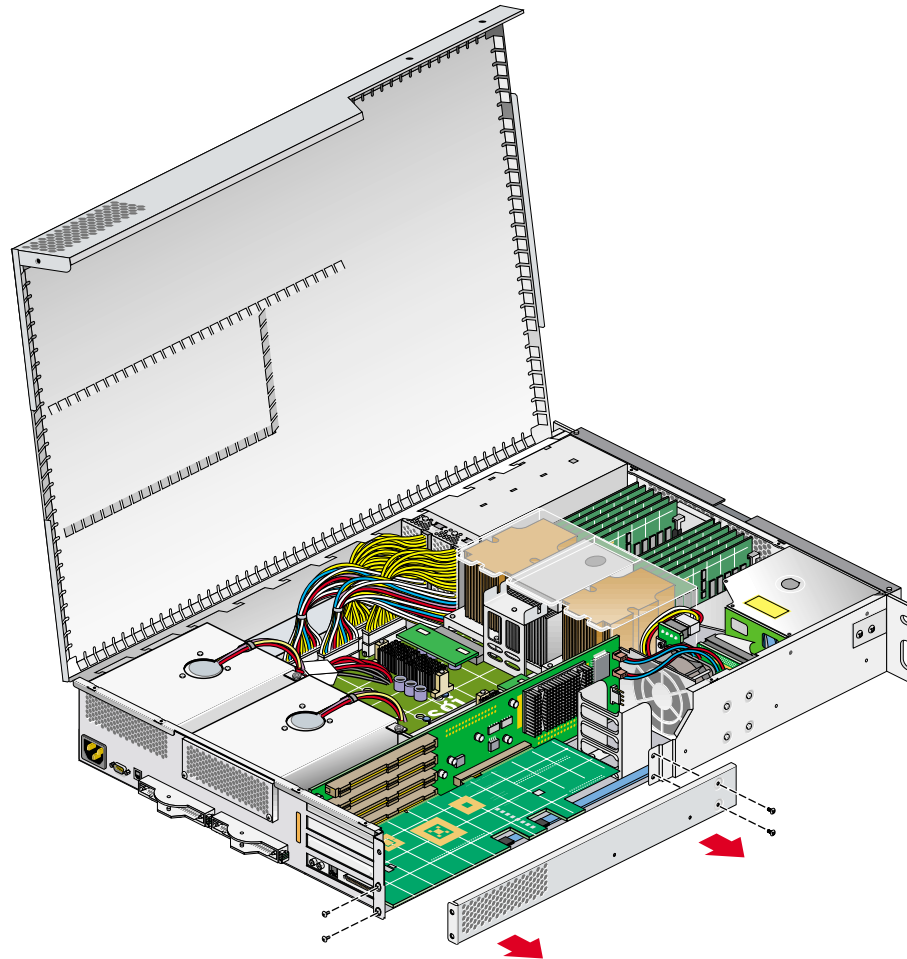
Figure 3-5 Removing the Lower PCI/PCI-X Support Bracket

10. If a blanking plate covers the slot that is needed for the installation, remove the retaining screw, as shown in Figure 3-6, and the blanking plate.
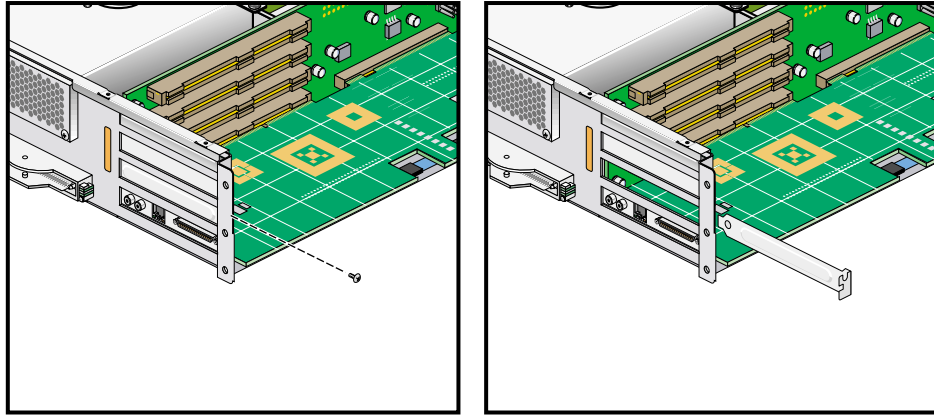
Figure 3-6   Removing the Blanking Plate

11. Insert the card into the slot by pushing the card into the connector until it is properly seated and install the retaining screw, as shown in Figure 3-7. If you have installed the card into one of the upper two slots, proceed to step 13.
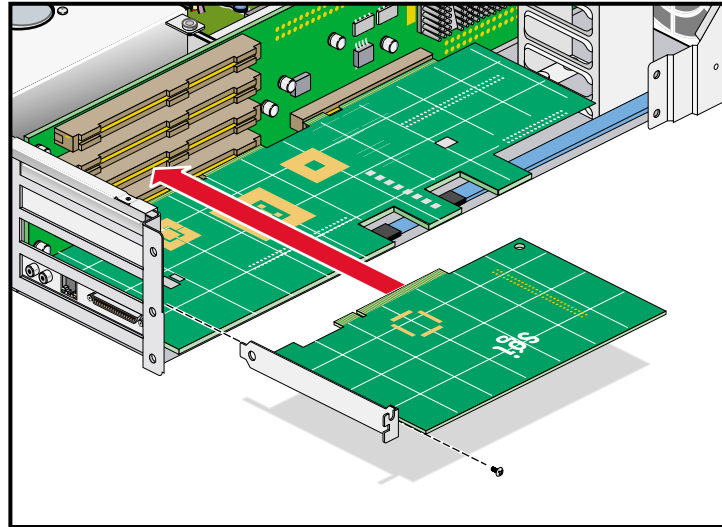
Figure 3-7   Installing the Card and Retaining Screw

12. If you have installed a card in one of the two bottommost card slots, replace the lower PCI/PCI-X support bracket that covers the two bottommost slots and screw in the four Phillips screws, as shown in Figure 3-8.

Figure 3-8   Replacing the Lower PCI/PCI-X Support Bracket

13. Replace the optional chassis rail (if applicable) by screwing in the five Phillips screws, as shown in Figure 3-9.

Figure 3-9   Replacing the Chassis Rail

14. If you removed the module from the rack, perform the following substeps. (If you have not removed the module from the rack, proceed to step 15.)

    a.   Fully extend the left and right slide rails from the rack until they lock into place.

    b.   This step requires two people. With one person holding each side of the module, align the chassis rails of the module with the slide rails of the rack.

    c.   Slide the chassis rails into the slide rails until the chassis rails are stopped by the safety latches.

15. Press the safety latches on both sides of the module and slide the module into the rack.

16. Install the two screws that secure the module to the front rails of the rack. For shelf mounted units, install the two screws at the rear of the module.

17. Install all of the cables at the rear of the module.

18. Power on the Silicon Graphics Prism platform. For power on instructions, see Chapter 1, "Getting Started with the Silicon Graphics Prism".

Removing a Card

To remove a PCI or PCI-X card, follow these steps:

1. Power off the Silicon Graphics Prism platform. For powering off instructions, see Chapter 1, "Getting Started with the Silicon Graphics Prism".

2. Disconnect all of the cables at the rear of the module.

Warning: **Components may be hot. To avoid injury, allow the components to cool for approximately five minutes before you proceed with these instructions.**

3. Remove the two screws that secure the module to the front rails of the rack.

4. If the module is mounted on shelf rails, remove the two screws at the rear of the unit.

5. If your system uses optional slide rails, pull the module from the rack until it is stopped by the safety latches, then press the latches to release the unit.

6. Use two people to move the module from the rack to a sturdy non-static surface.

Note: If you are removing a PCI or PCI-X card from one of the two lower slots (bus 1, slot 1, or bus 1, slot 2), you must remove the module from the rack. For detailed instructions, see "Removing a Rail-mounted Module from a Rack" on page 159.

7. Make sure that you read "Safety Instructions" on page 47, and "Rules for Card Installation and Removal" on page 51, before installing the card.

8. To access the card, remove the ten Phillips screws, as shown in Figure 3-10, and lift the hinged cover.
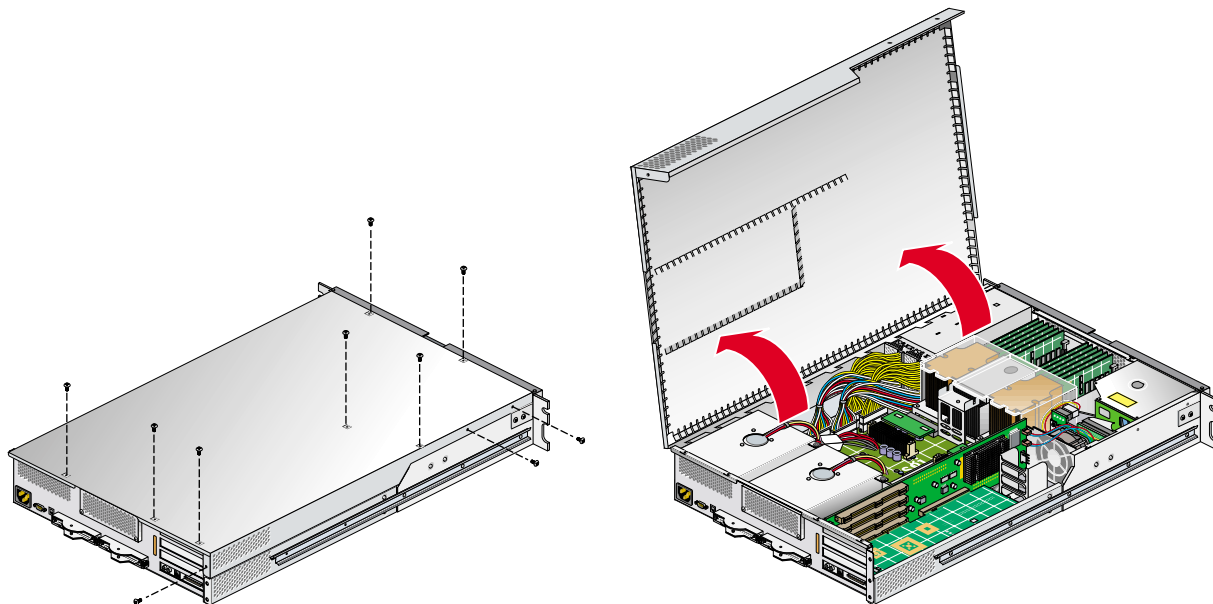
Figure 3-10   Opening the Cover to Remove the Card

9. If you are removing a card from one of the two lowermost card slots and your system uses optional slide rails, remove the chassis rail, by unscrewing the five Phillips screws as shown in Figure 3-11. (If you are not removing a PCI/PCI-X card from one of the two lowermost card slots, proceed to step 9.)

⚠️ **Caution:** Only a trained SGI service support engineer can install and remove an IO10 card. Otherwise, your system could be damaged. Therefore, if an IO10 card is installed in the module, you can remove a card only from the upper three slots.
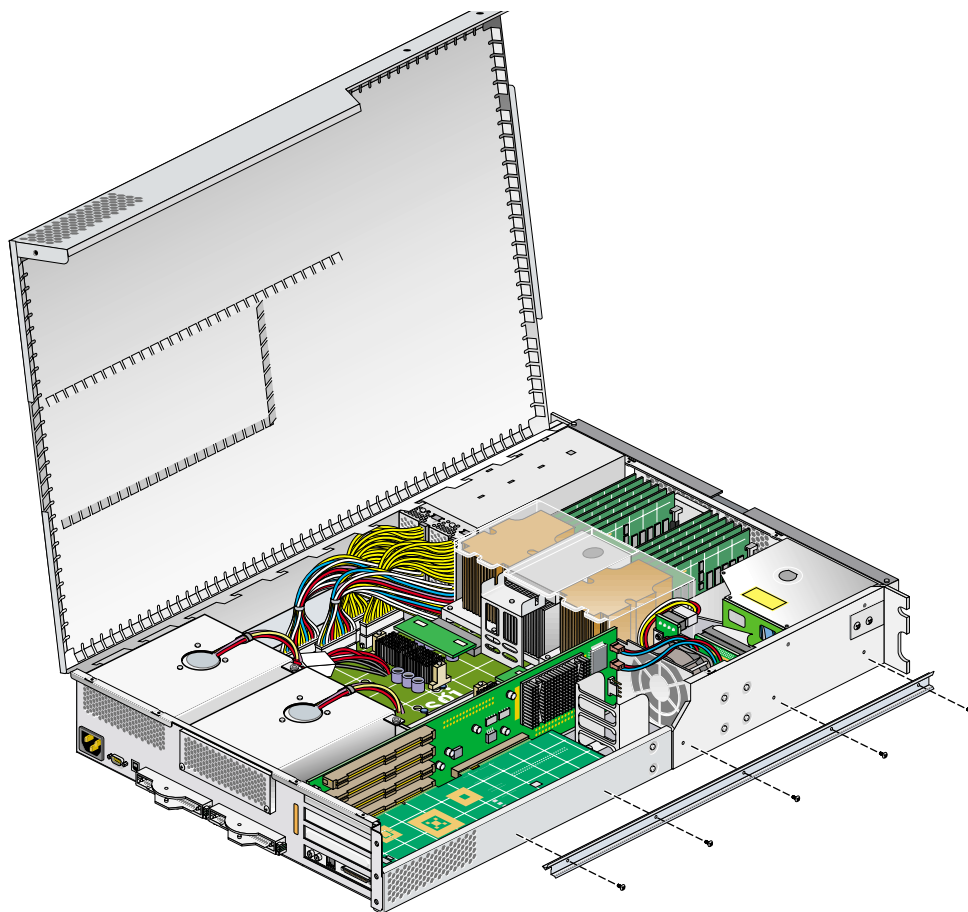
Figure 3-11   Removing the Chassis Rail

10. If you are removing a card from one of the two lowermost slots, you must also remove the lower PCI/PCI-X support bracket that covers the two lowermost slots by unscrewing the four Phillips screws, as shown in Figure 3-12.
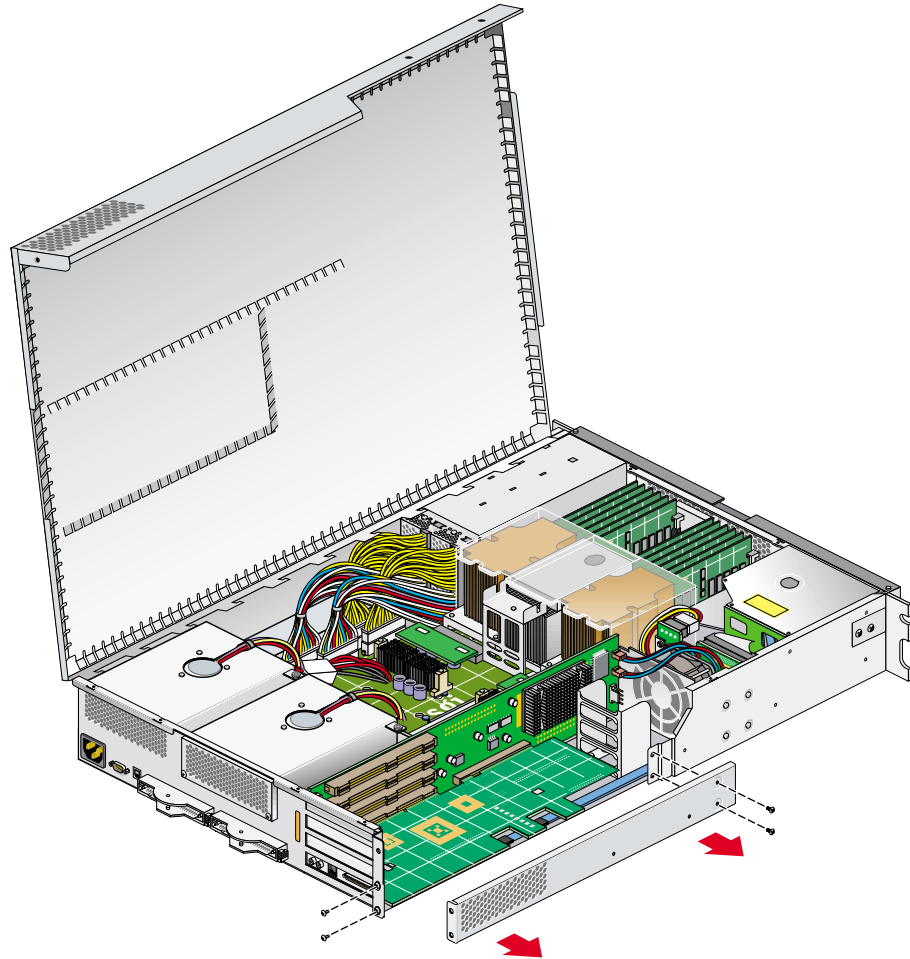
Figure 3-12   Removing the Lower PCI/PCI-X Support Bracket

11. Unscrew the retaining screw from the card that you will remove, and extract the card, as shown in Figure 3-13. Place the card on an ESD-safe surface.
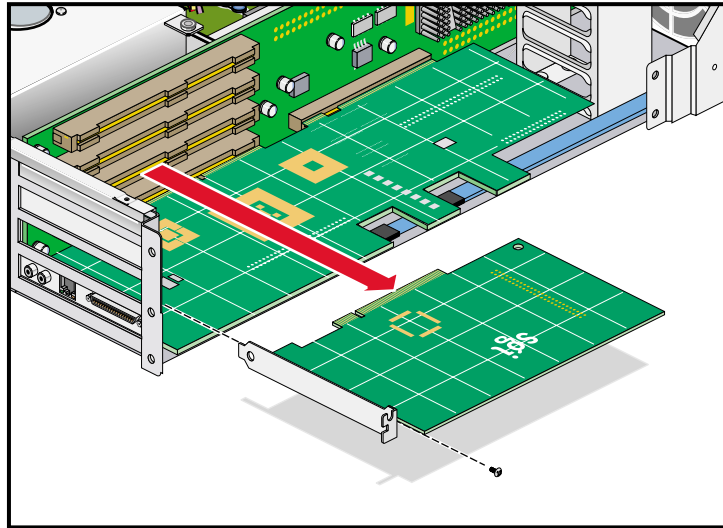
Figure 3-13   Removing the Retaining Screw and Extracting the Card

12. If you are replacing the card that you have removed, reference the instructions in "Installing a PCI or PCI-X Card" on page 52. If you are not replacing the card that you have removed, proceed to the next step.

13. Place a blanking plate to cover the slot where you removed the card, and screw in the retaining screw, as shown in Figure 3-14. After installing the blanking plate, if you have removed a card from one of the two lowermost slots, proceed to the next step. Otherwise, proceed to step 14.
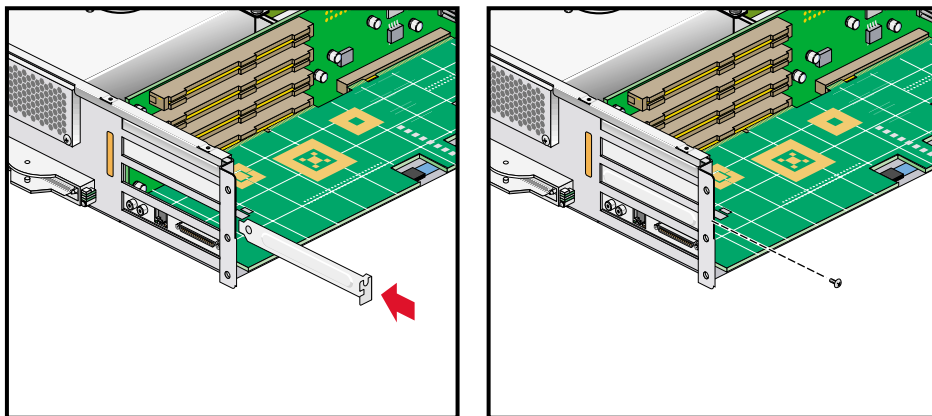
Figure 3-14   Replacing a Blanking Plate

14. If you have removed a card from one of the two bottommost card slots, replace the lower PCI/PCI-X support bracket that covers the two bottommost slots, and screw in the four Phillips screws, as shown in Figure 3-15.

Figure 3-15   Replacing the Lower PCI/PCI-X Support Bracket

15. Replace the optional chassis rail (if applicable), by attaching the five Phillips screws, as shown in Figure 3-16.

Figure 3-16   Replacing the Optional Chassis Rail

16. Close the hinged cover on the system and screw in the ten Phillips screws to secure the cover, as shown in Figure 3-17.
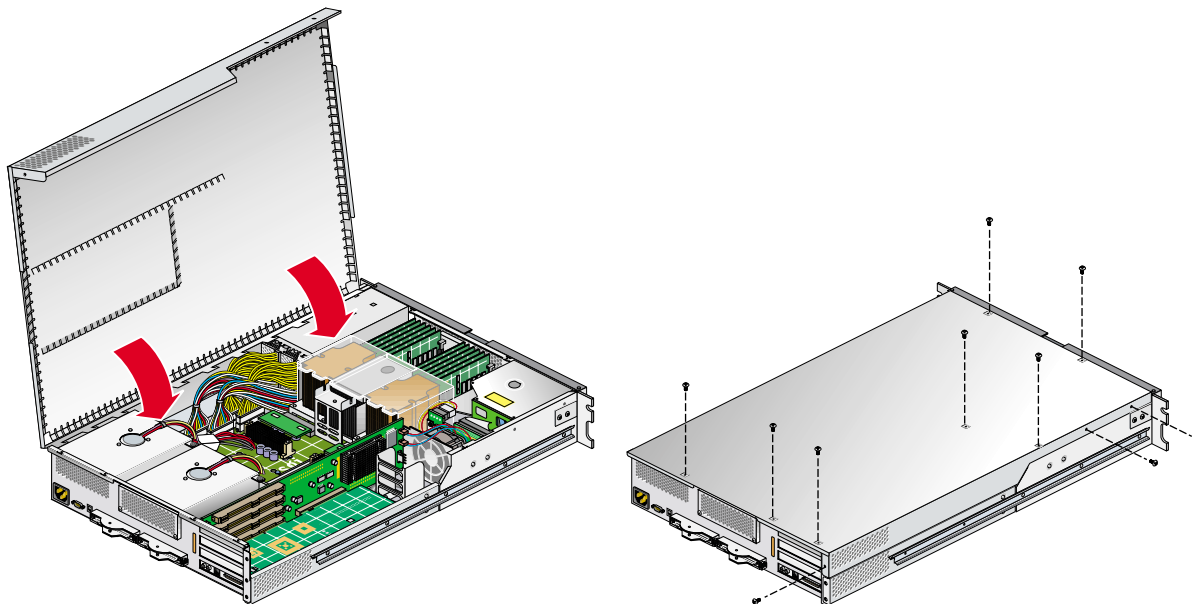
Figure 3-17   Close and Screw Down Cover

17. If your system uses optional slide rails and you removed the module from the rack, follow these substeps; if not, go on to step 19.

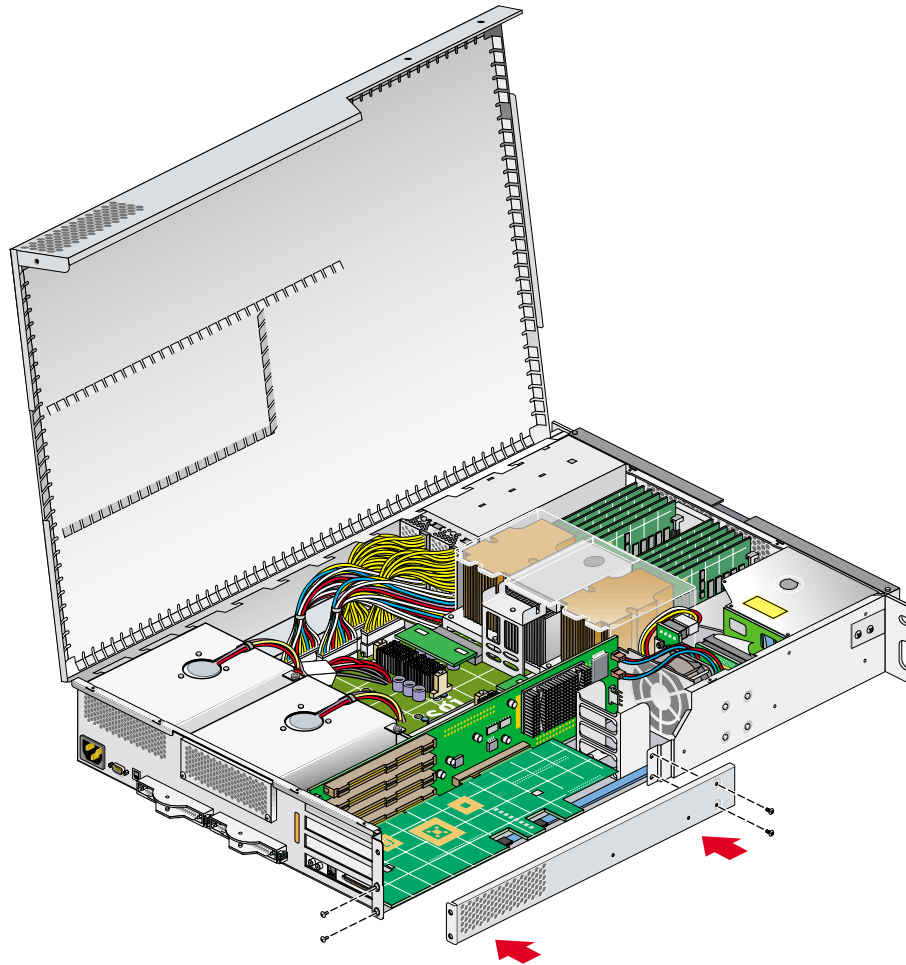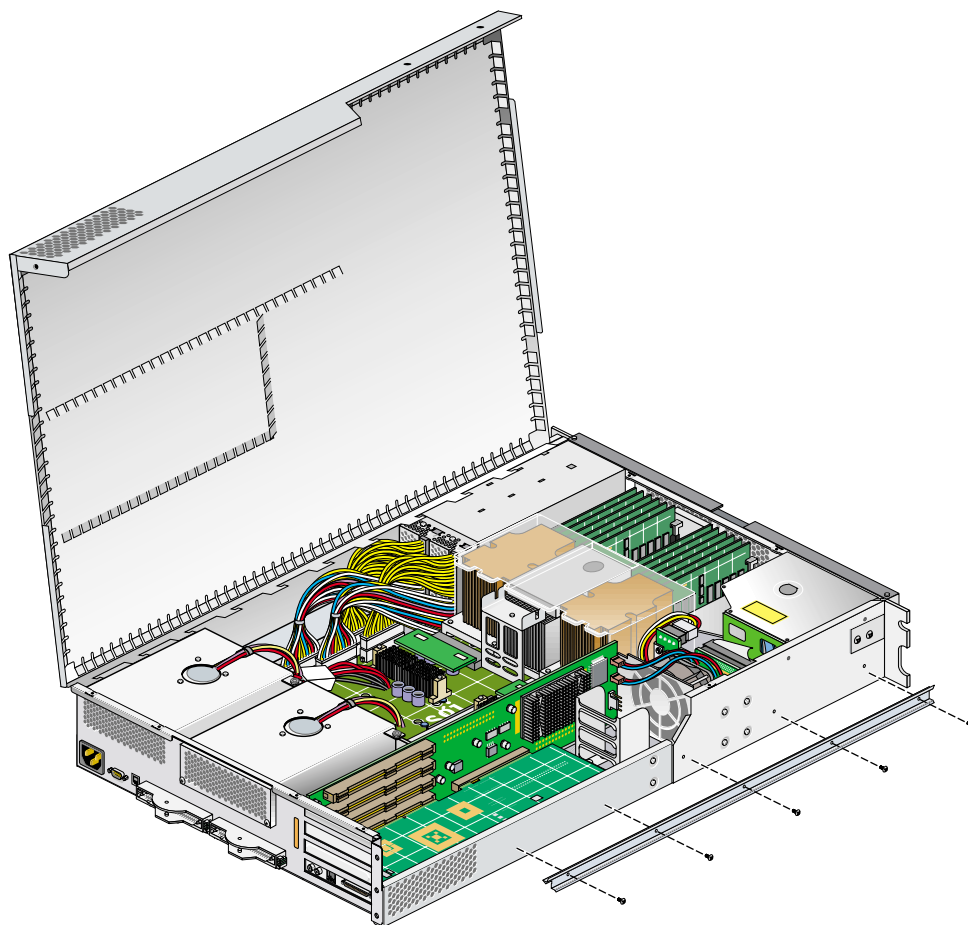   a.   Fully extend the left and right slide rails from the rack until they lock into place.

   b.   This step requires two people. With one person holding each side of the module, align the chassis rails of the module with the slide rails of the rack.

   c.   Slide the chassis rails into the slide rails until stopped by the safety latches.

18. Press the safety latches on both sides of the module, and slide it into the rack.

19. For modules mounted on shelf rails, use two people to slide the unit into the rack.

20. Install the two screws that secure the module to the front rails of the rack. For shelf mounted units, secure the two screws at the back.

21. Install all of the cables at the rear of the module.

22. Power on the Silicon Graphics Prism platform. For powering on instructions, see "Power-On the System" on page 16.

Replacing an XG2N Graphics Card

To replace a graphics card, follow these steps:

1. Power off the Silicon Graphics Prism platform. For power off instructions, see "Powering Off Manually" in Chapter 1.

2. Disconnect all of the cables at the rear of the module.

Warning: Components may be hot. To avoid injury, allow the components to cool for approximately five minutes before you proceed with these instructions.

3. If your module is rackmounted, remove the two screws that secure the module to the front rails of the rack. If the module is mounted with shelf rails, remove the screws at the rear also. If your module is not rackmounted, proceed to step 6.

4. If the module is mounted with the optional slide rail kit, pull the module from the rack until it is stopped by the safety latches. Otherwise, two people should lift the unit from the rack and place it on a stable work surface.

5. Remove the module (using two people) from the optional slide rails by depressing the safety latches and sliding the module outward. Place it on a stable work surface.

6. Open the top cover as described in "Accessing the System's Internal Components" on page 49.

7. Make sure that you read "Safety Instructions" on page 47, before beginning removal of a graphics card.

8. When removing and replacing a graphics card, you may need to remove an optional chassis rail by unscrewing the five Phillips screws, (an example is shown in Figure 3-11).

9. Remove the lower PCI/PCI-X support bracket by unscrewing the four Phillips screws, as shown in Figure 3-18.
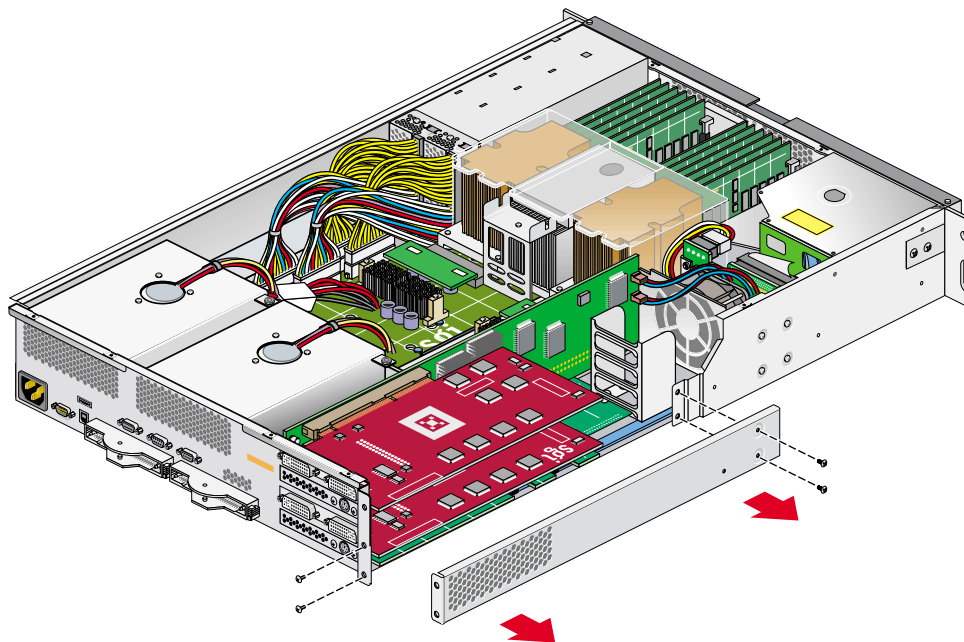
Figure 3-18   Removing the XG2N Graphics Card Bracket

10. Undo and remove the two screws that hold the card to the backpanel, as shown in Figure 3-19.

⚠️ **Caution:** Each graphics card is connected to the riser board by a 7-pin synchronization cable connected from the PCI riser board to the underside of the graphics card. Be careful not to pull the graphics card away from the board too far as you remove it.

11. Pull the card gently out of the main connector on the riser board until it is clear of the backpanel and main connector.

12. Carefully rotate the board over and disconnect the 7-pin keyed synchronization cable from the bottom of the graphics board (it attaches the graphics card assembly to the PCI riser board) See Figure 3-20 on page 72.

13. Place the graphics card on an ESD-safe surface or in a protective bag.

Rotate board downward
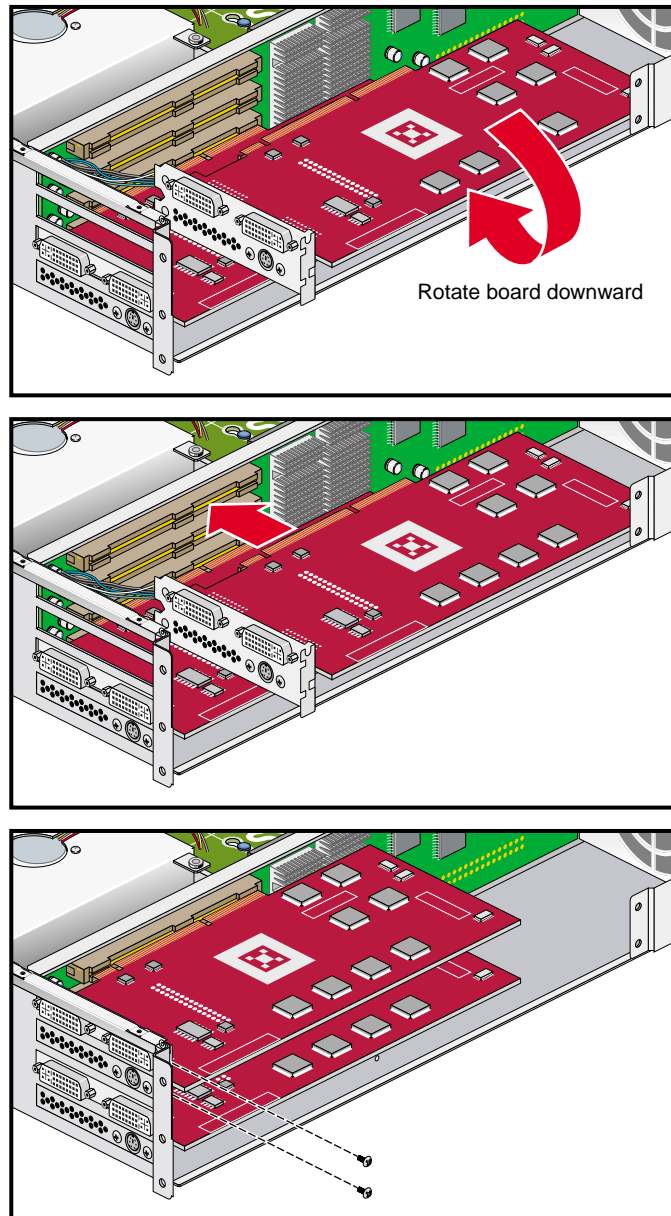
Figure 3-19   Remove XG2N Graphics Board Assembly
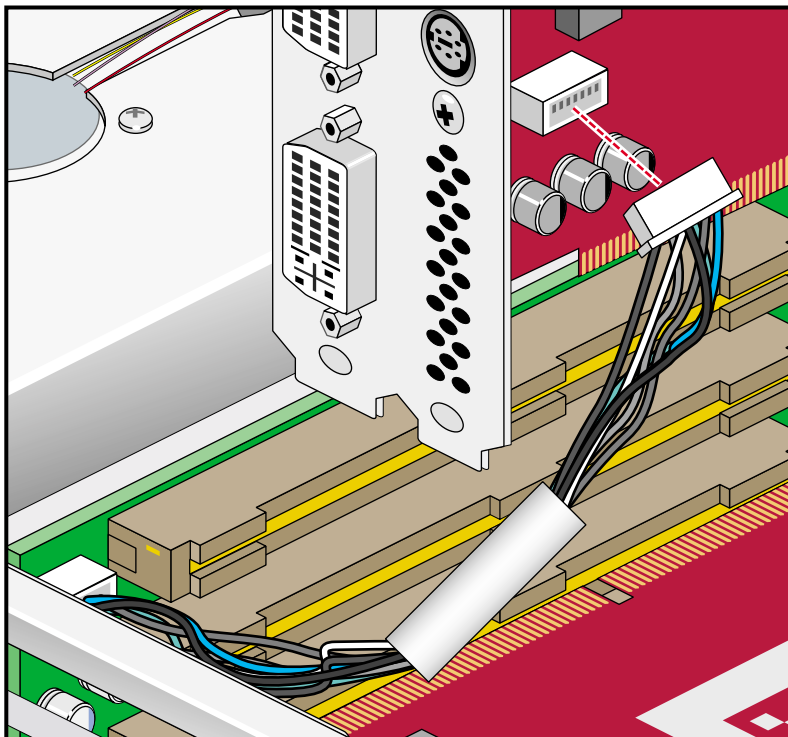
Figure 3-20   Synchronization Cable Connection on Underside of Graphics Board

To install a new XG2N graphics card, use the following steps:

1. Place the board close to the graphics connector on the riser board and tilt it up until you can see the connector location for the 7-pin synchronization cable.

2. Gently insert the 7-pin synchronization cable into the connector on the bottom of the replacement graphics card, see Figure 3-20. The connector is keyed and will only insert in the proper orientation. You should hear or feel a slight snap as the connector is seated.

3. Insert the card into the slot the old graphics card was removed from by pushing the card into the connector until it is properly seated and installing the backpanel retaining screws, as shown in Figure 3-21.

    Be careful not to snag the 7-pin synchronization cable between the graphics card and the riser board connector when you install the new card.

Rotate board downward
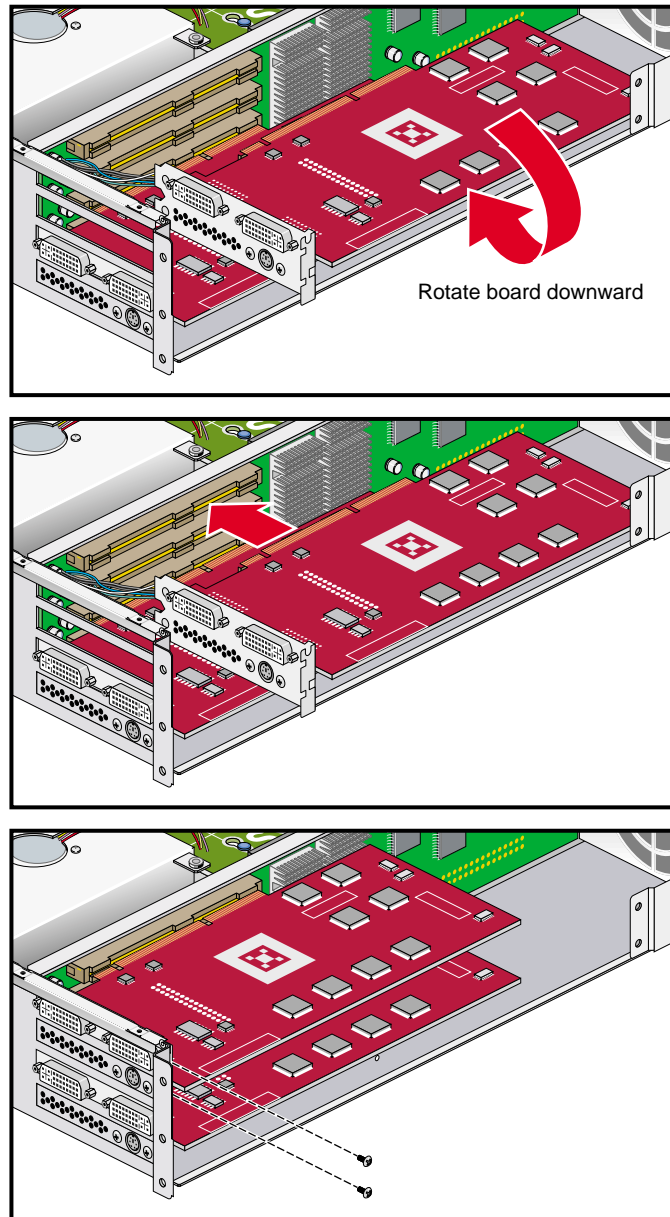
Figure 3-21   Installing the Graphics Card and Retaining Screws

4. Replace the lower PCI/PCI-X support bracket and screw in the four Phillips screws, as shown in Figure 3-22.
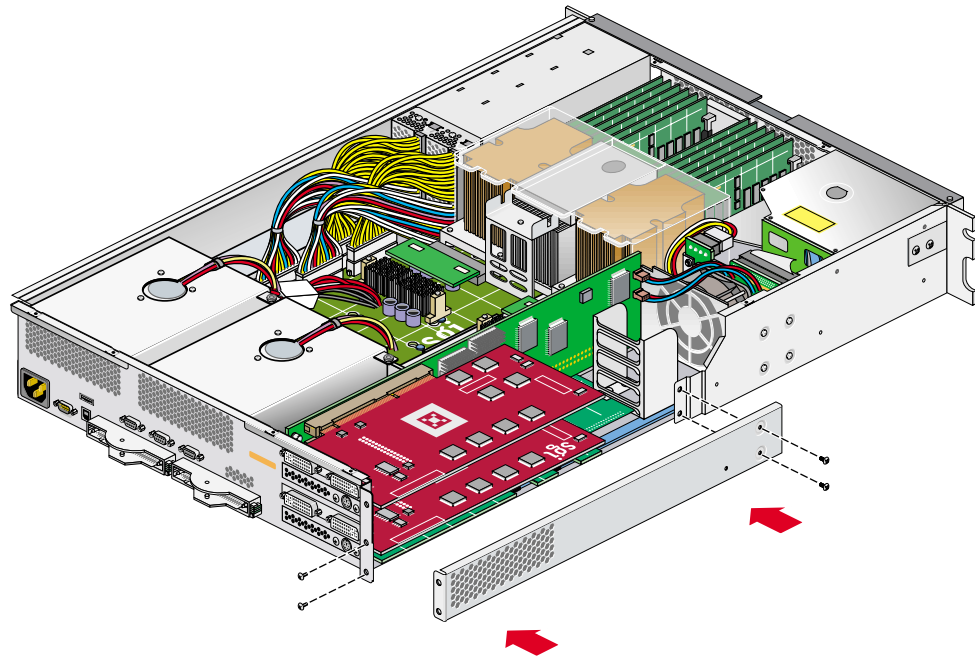
Figure 3-22   Replacing the PCI/PCI-X Support Bracket

5. Replace the optional chassis rail (if applicable) by screwing in the five Phillips screws. Reference Figure 3-16 if necessary.

6. Close the hinged cover on the graphics module and screw in the ten Phillips screws to secure the cover, (reference Figure 3-17 for an example).

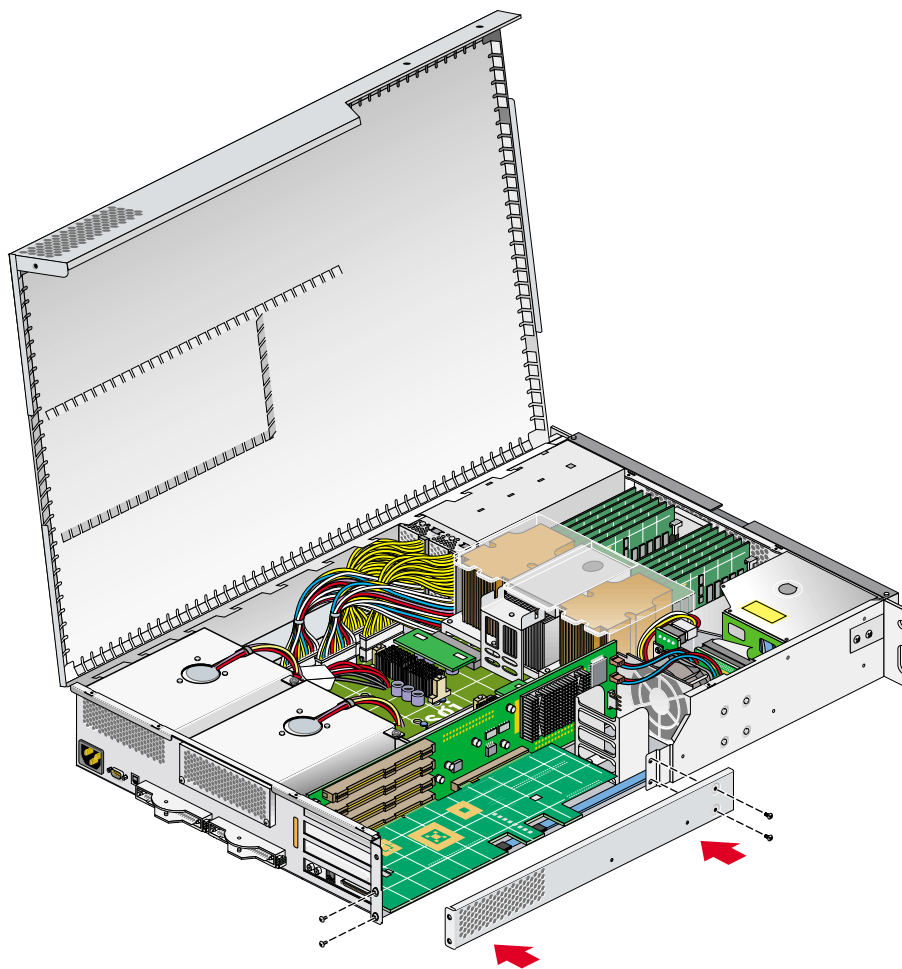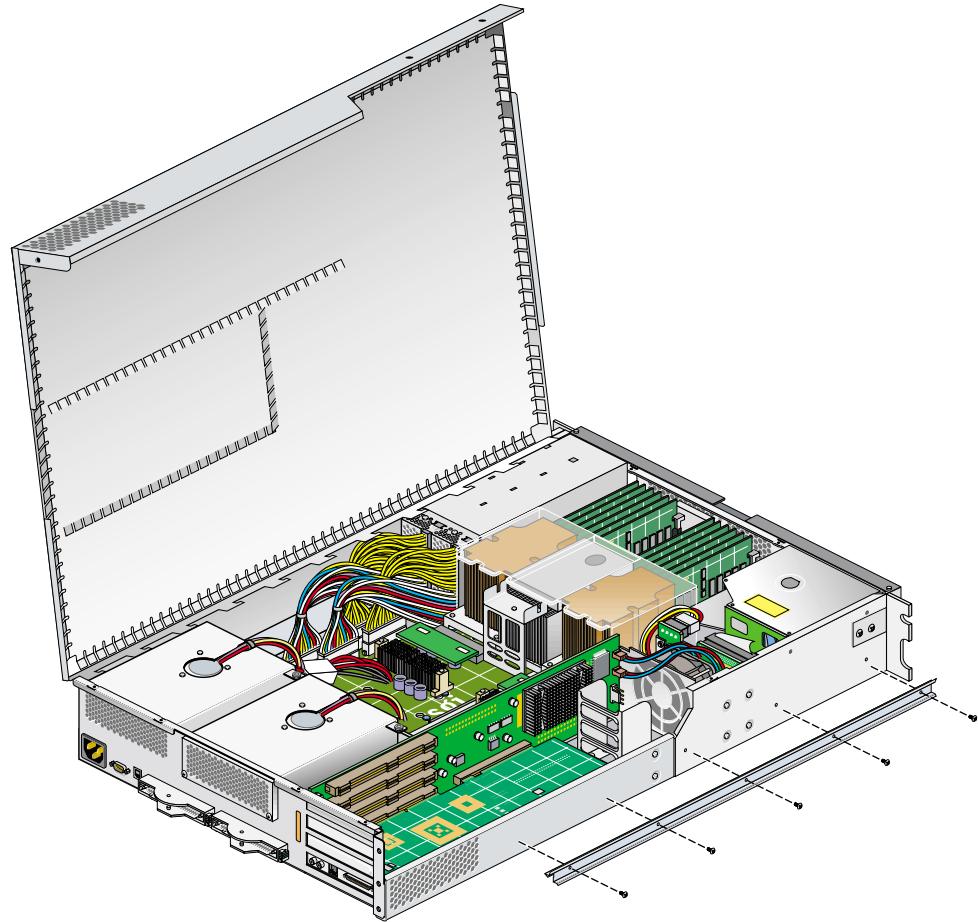7. If your system uses optional slide rails, follow substeps a-d; otherwise go to step 8.

    a. Fully extend the left and right slide rails from the rack until they lock into place.

    b. This step requires two people. With one person holding each side of the module, align the chassis rails of the module with the slide rails of the rack.

    c. Slide the chassis rails into the slide rails until the chassis rails are stopped by the safety latches.

    d. Press the safety latches on both sides of the module, and slide the module into the rack.

8. For modules mounted on shelf rails, use two people to slide the unit into the rack.

9. Install the two screws that secure the module to the front rails of the rack. For shelf mounted units, secure the two screws at the back.

10. Install all of the cables at the rear of the module.

11. Power on the Silicon Graphics Prism platform. For powering on instructions, see "Power-On the System" on page 16.

Disk Drives

Each base compute module contains one or two sled-mounted serial ATA disk drives (see Figure 3-23).
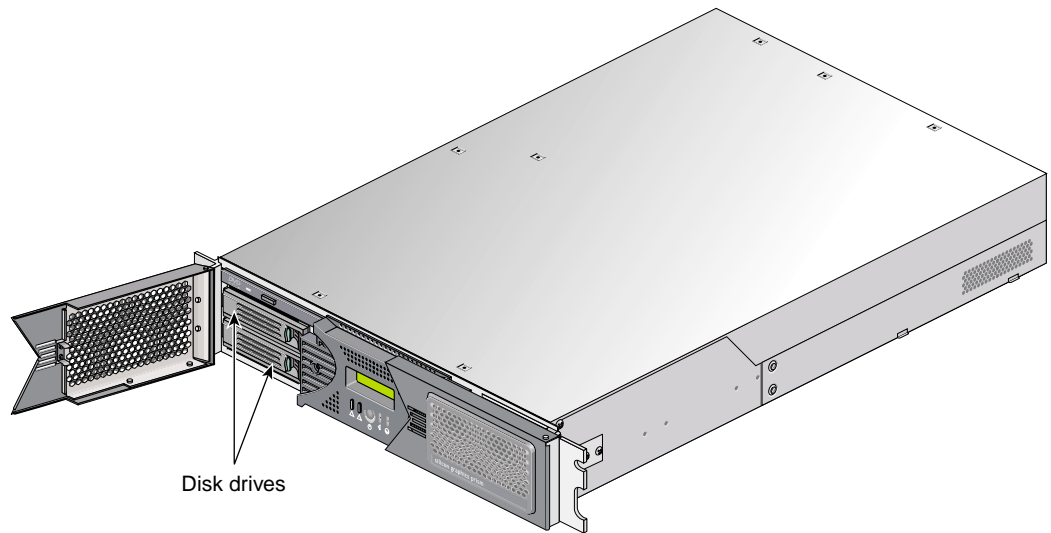
Disk drives

Figure 3-23   Disk Drive Locations

Installing a Disk Drive

To install a disk drive, follow these steps:

1. Open the bezel door as far as it will open. Position the drive assembly so that it engages the bay guide rails and, with the locking handle fully swung open, gently push the drive into the bay until the locking handle engages with left side of the bay opening, as shown in Figure 3-24A.

   Note: If you have only one disk drive in your system, it should be located in the bottommost slot.

2. Swing the locking handle towards the chassis until the locking handle engages the latch, as shown in Figure 3-24B and Figure 3-24C.

3. Close the bezel door, as shown in Figure 3-24D.
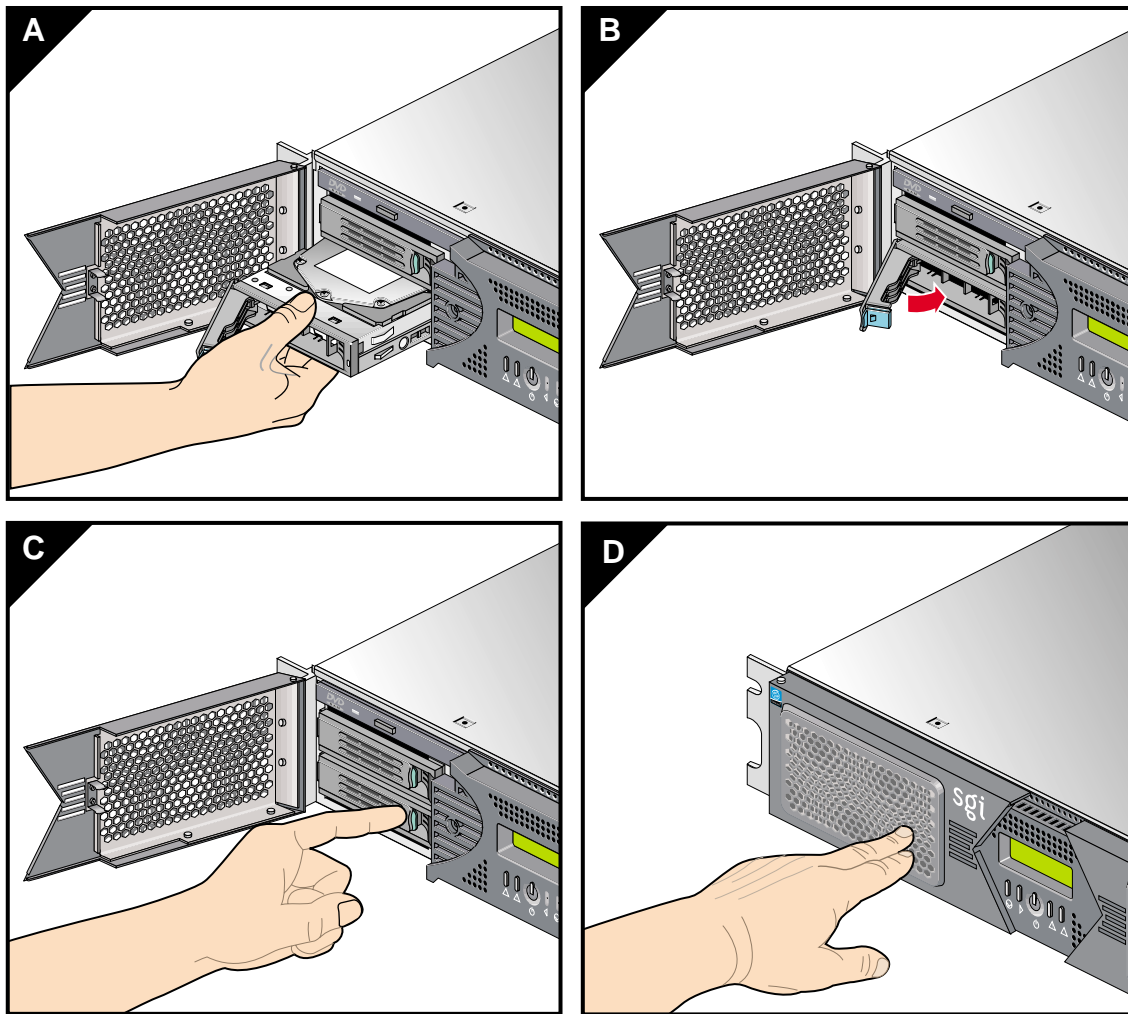
Figure 3-24　Installing a Disk Drive

Removing a Disk Drive

To remove a disk drive, follow these steps:

1. If you are replacing a data drive, ensure that the drive has spun down before you remove it.

2. If you are replacing the system drive, you must first power off the Silicon Graphics Prism platform. For instructions to power off the Silicon Graphics Prism platform, see "Powering Off Manually" on page 19.

3. Open the bezel door located on the left side of the front panel of the module, as shown in Figure 3-25A. (Make sure that you open the door as far as it will open.)

4. Remove the drive by depressing the locking handle with your forefinger (Figure 3-25B). Then swing open the locking handle away from the chassis until the handle disengages the drive connector from the backplane connector (see Figure 3-25C).

   Note: If you will have only one disk drive, it should be located in the bottom-most slot.

5. Carefully slide the drive out of the bay (see Figure 3-25D) and gently place it on a flat ESD-safe surface. (Do not use the handle to pull the drive out of the bay.)

6. If you are replacing the disk drive, proceed to "Installing a Disk Drive" on page 77. (After you have replaced the disk drive, return to step 8 for instructions to power on your module.) If you are not replacing the disk drive, proceed to the next step.

7. Close the bezel door.

8. Power on the Silicon Graphics Prism platform as described in "Power-On the System" on page 16.
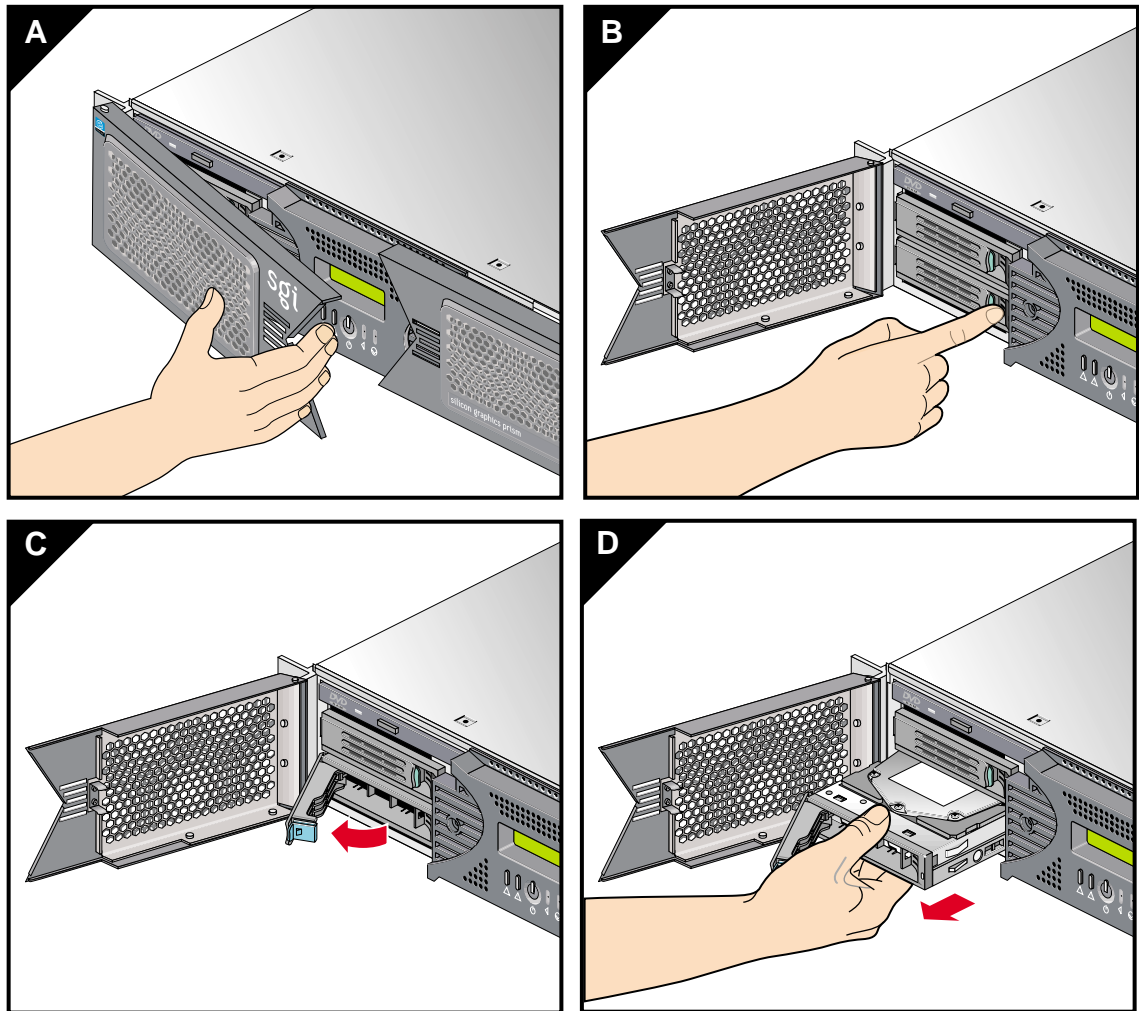
Figure 3-25   Removing a Disk Drive

Removing and Replacing Power Supplies

Each Silicon Graphics Prism module can contain one or two sled-mounted power supplies (see Figure 3-26). For information on determining the operational status of a power supply, see "Reading Power Supply Status LEDs" in Chapter 4.

Note: The XG2N graphics module always requires two power supplies, no redundant power is available. The system should be shut down prior to power supply replacement.
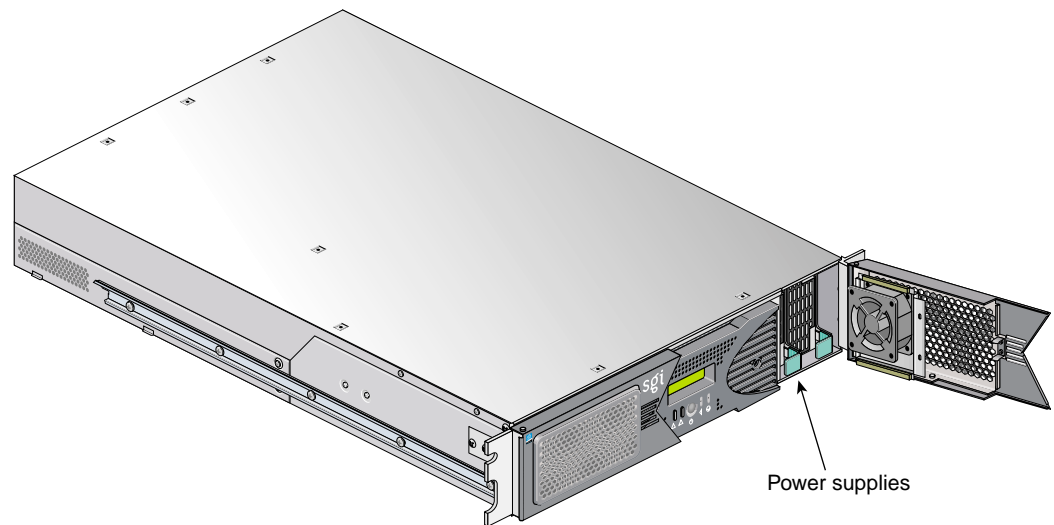
Power supplies

Figure 3-26   Power Supplies Location

To replace a power supply, follow these steps:

1. If you are replacing a power supply in an XG2N module, power off the Silicon Graphics Prism platform as described in the section "Powering Off Manually" on page 19. The XG2N module always uses two supplies. There is no redundant power supply option available.

2. Remove the power supply to be replaced, as follows:

    a. Swing open the bezel door located on the right side of the module front panel. With a Phillips screw driver, unscrew the two screws on the screen cover as shown in Figure 3-27A.

    b.   Swing open the screen cover as shown in Figure 3-27B.

    c.   Disengage the power supply from the power supply bay by pushing the interior release button to the right and pulling up and out on the green-colored handle lock as shown in Figure 3-27C.

    d.   Gently pull out the power supply from the chassis until it clears the power supply bay as shown in Figure 3-27D. Place the power supply on an ESD-safe surface.

> **Caution:** When pulling out the power supply from the chassis, make sure not to disturb the power supply fan's ribbon cable.

3. Install the replacement power supply, as follows:

    a.   Position the power supply in the slot and with the power supply handle pulled up (fully opened), gently push the power supply into the bay as shown in Figure 3-28A.

> **Caution:** When installing the power supply, make sure that the power supply does not clip or pinch the power supply fan's ribbon cable.

    b.   Push in and down on the green-colored handle and snap the power supply into place as shown in Figure 3-28B.

    c.   After you have installed the power supply, swing the screen cover until it closes as shown in Figure 3-28C.

> **Caution:** When closing the screen cover, make sure that the cover does not clip or pinch the power supply fan's ribbon cable.

    d.   Screw in the two Phillips screws that you had removed as shown in Figure 3-28D, and close the bezel door.

4. Power on the Silicon Graphics Prism platform, as described in "Power-On the System" on page 16.
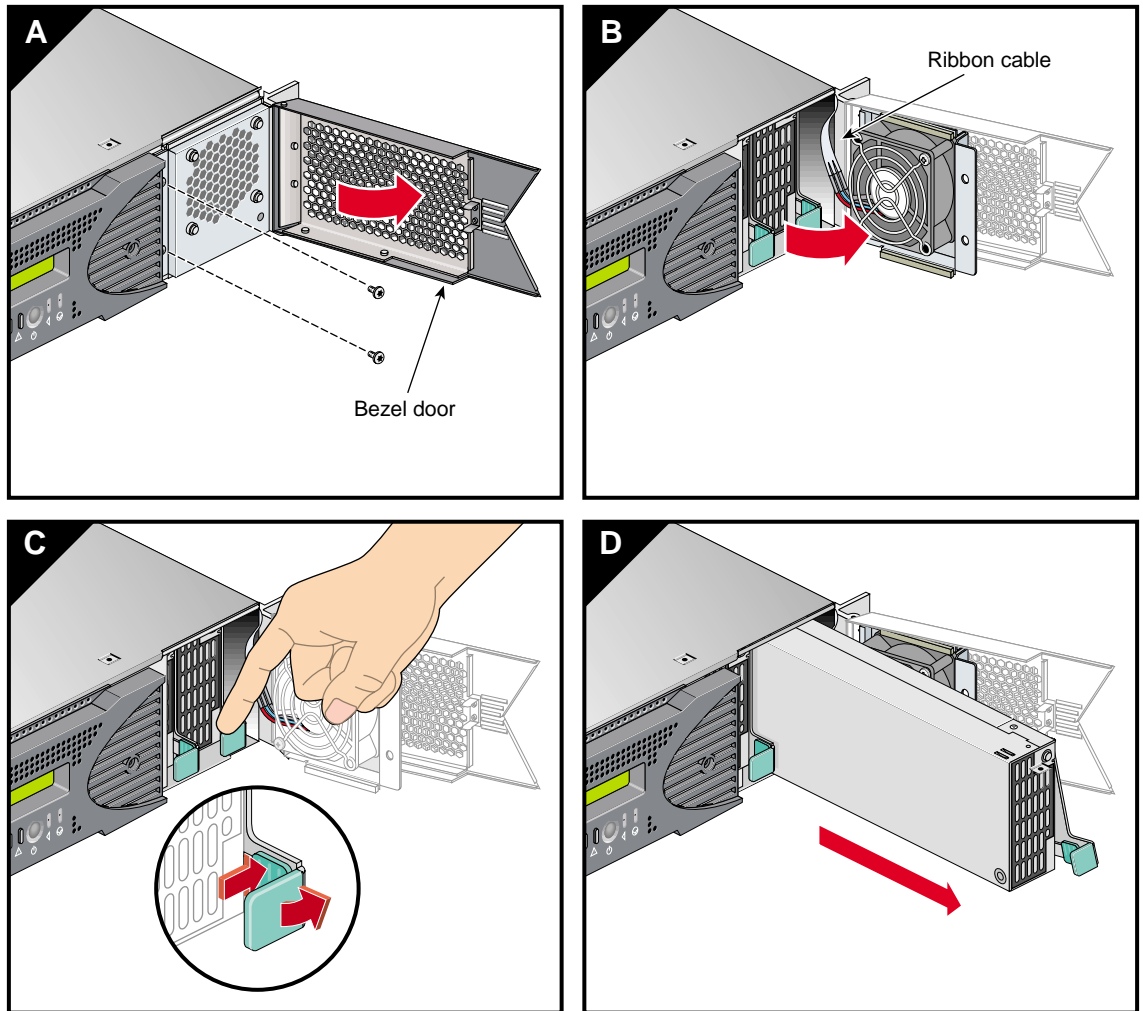
Figure 3-27   Removing a Power Supply
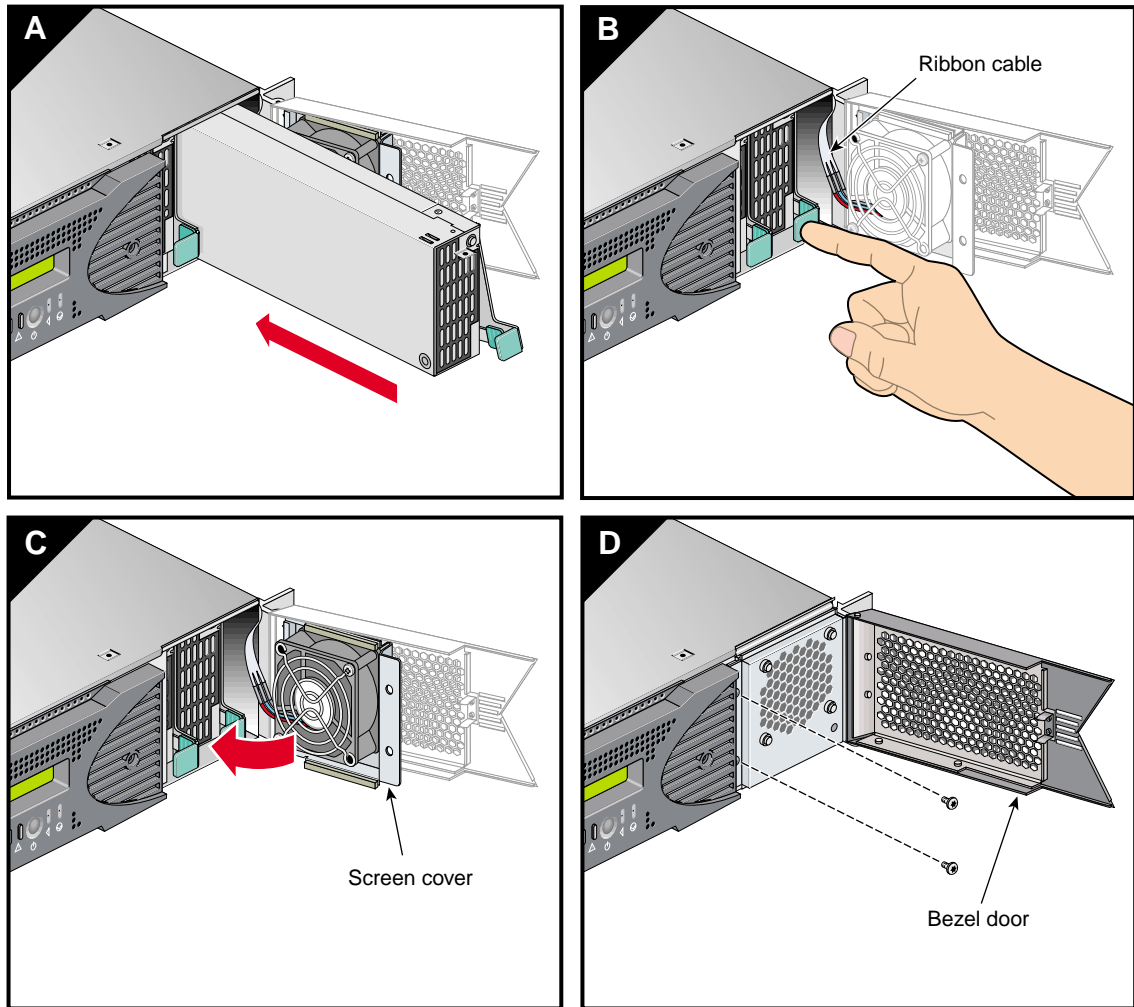
Figure 3-28   Installing the Power Supply

Memory DIMM Group Placement and Guidelines

Memory is contained on cards that are referred to as DIMMs (dual inline memory modules). Each module can contain four, eight, or twelve DIMMs installed in DIMM slots located on the module's node board.

These twelve DIMM slots are organized into three groups of 4 DIMMs each, as shown in Figure 3-29 on page 86.

DIMMs are installed one per DIMM slot, and four at a time (group of four).

Follow these guidelines when installing DIMM groups:

- Memory is increased or decreased in four-DIMM group increments only.

- Each of the four DIMMs that make up a DIMM group must be the same memory size and speed; however, different groups of DIMMs can differ in memory size.

- The first four DIMMs (group 0) must be in place for the base module to operate properly.

- Mixing DIMM groups with different access speeds will cause the memory bus to default to the speed used for the slowest group.

- The maximum bus speed for a completely full set of DIMM memory slots (three installed groups), defaults to 100 MHz or 133 MHz, depending on the type of DIMMs installed.

- The DIMMs used in the Silicon Graphics Prism platform modules are **not** compatible with DIMMs used in the Origin 200, Origin 350 series, SGI 3000 series, SGI 2000 series, Onyx systems, Fuel, Octane or Tezro systems.
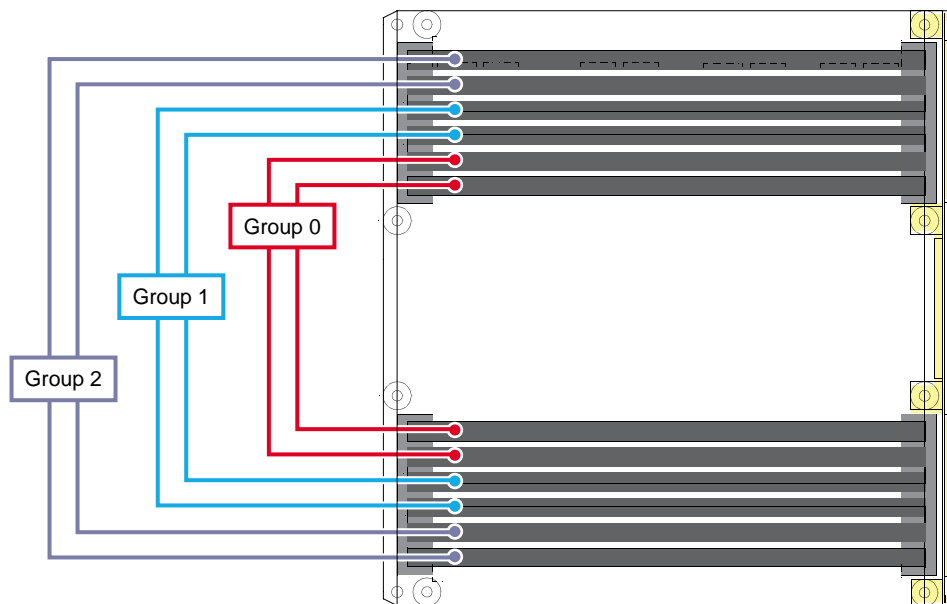
Figure 3-29   Layout of Slots and DIMM Memory Groups

> ⚠️ **Caution:** Electronic equipment can be irreparably damaged by electrostatic discharge (ESD). Always follow these preventive measures when you handle a system component:
>
> - Remove a component from its antistatic bag only when you are ready to install it.
>
> - If you handle a component before installation, do not place it on surfaces that produce ESD (carpeting, for example) or near devices that create static electricity.
>
> - Attach a static wrist strap to a grounded connection on your system when you install or remove a component.

Installing a DIMM Group

To install a DIMM group, follow these steps:

1. Power off the Silicon Graphics Prism platform. For powering off instructions, see "Powering Off Manually" on page 19.

2. Disconnect all of the cables at the rear of the module.

Warning: **Components may be hot. To avoid injury, allow the components to cool for approximately five minutes before you proceed with these instructions.**

3. Remove the two screws that secure the module to the front rails of the rack. For shelf-mounted modules, remove the two screws at the rear of the unit.

4. With shelf-mounted rails, use two people to slide the module out and move it to a stable work surface. With optional slide rails, pull the module from the rack until it is stopped by the safety latches.

5. To access the DIMMs, remove the ten Phillips screws shown in Figure 3-30 and lift and open the hinged cover.
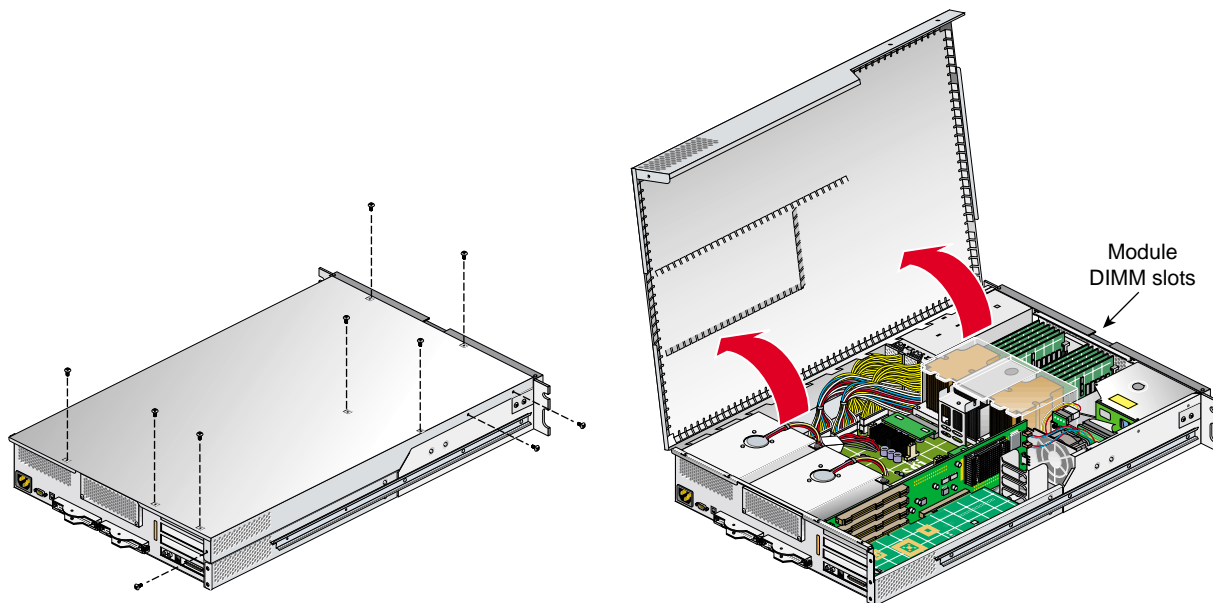
Figure 3-30   Opening Module Cover to Install DIMMs

6. Install the DIMMs, as follows (see Figure 3-31):

Note: If you need to find the correct locations in which to install the DIMMs, make sure to read the introductory material in "Memory DIMM Group Placement and Guidelines" on page 85.

---

a. Ensure the ejector latches are in the open position (leaning away from the slot).

b. Hold the DIMM only by its edges and remove it from its antistatic package.

c. Align the bottom edge of the DIMM with the keyed socket.

d. Insert the bottom edge of the DIMM into the socket, and then press down evenly on the DIMM until it seats correctly. Use extreme care when you install a DIMM. If you apply too much pressure, you can damage the socket.

e. When the DIMM is fully seated in the connector, the ejector latches snap into place flush with each end of the DIMM, see Figure 3-31.

7. Attach the hinged cover and secure it to the module with ten Phillips screws.

8. If using optional slide rails, press the safety latches on both sides of the module and slide the module into the rack. For shelf-mounted units, use two people to slide the module back into the rack.

9. Install the two screws that secure the module to the front rails of the rack. If shelf rails are used, install the two rear screws.

10. Install all of the cables at the rear of the module.

11. Power on the Silicon Graphics Prism platform as described in "Power-On the System" on page 16.

Figure 3-31 Inserting a DIMM

Removing a DIMM

To remove a DIMM, follow these steps:

1. Power off the Silicon Graphics Prism platform. For powering off instructions, see "Powering Off Manually" on page 19.

2. Disconnect all of the cables at the rear of the module.

Warning: **Components may be hot. To avoid injury, allow the components to cool for approximately five minutes before you proceed with these instructions.**

3. Remove the two screws that secure the module to the front rails of the rack. For shelf-mounted modules, remove the two screws at the rear of the unit.

4. With shelf-mounted rails, use two people to slide the module out and move it to a stable work surface. With optional slide rails, pull the module from the rack until it is stopped by the safety latches.

5. To access the DIMMs, remove the ten Phillips screws shown in Figure 3-32 and open the hinged cover.
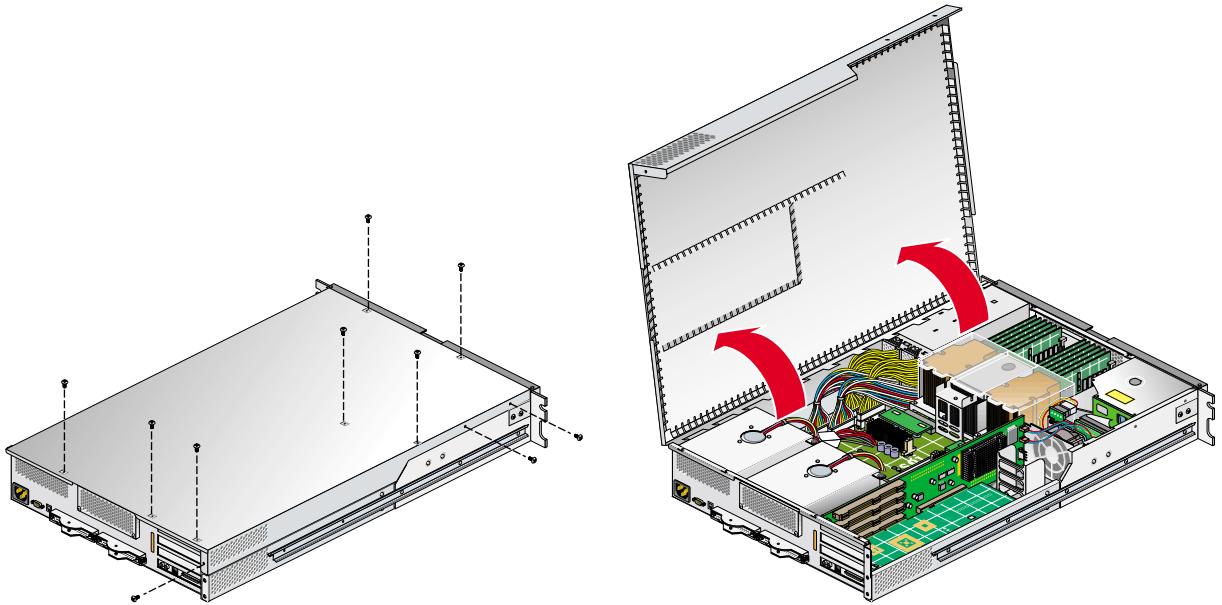
Figure 3-32   Opening the Module to Remove a DIMM

6. Remove the DIMM, as follows (see Figure 3-33):

> **Note:** For guidelines on removing or replacing the DIMMs, make sure to read the introductory material in "Memory DIMM Group Placement and Guidelines" on page 85 and see Figure 3-29 on page 86.

   a. Push out and down on the two ejector latches (located at each end of the DIMM socket) to disengage the DIMM from its connector.

   b. Carefully grasp the DIMM and pull it up and out of the guide rails.

   > **Note:** Hold the DIMM only by its edges. Be careful not to touch its components or gold edge connectors.

   c. Place the DIMM on an ESD-safe surface.

7. If you are installing a new DIMM, proceed to "Installing a DIMM Group" on page 87. If you are not installing a new DIMM, proceed to the next step.

8. Attach or close the hinged cover and secure it to the module with the ten Phillips screws.

9. If using optional slide rails, press the safety latches on both sides of the module and slide the module into the rack. For shelf-mounted units, use two people to slide the module back into the rack.

10. Install the two screws that secure the module to the front rails of the rack. If shelf rails are used, install the two rear screws.

11. Install all of the cables at the rear of the module.

12. Power on the Silicon Graphics Prism platform as described in "Power-On the System" on page 16.
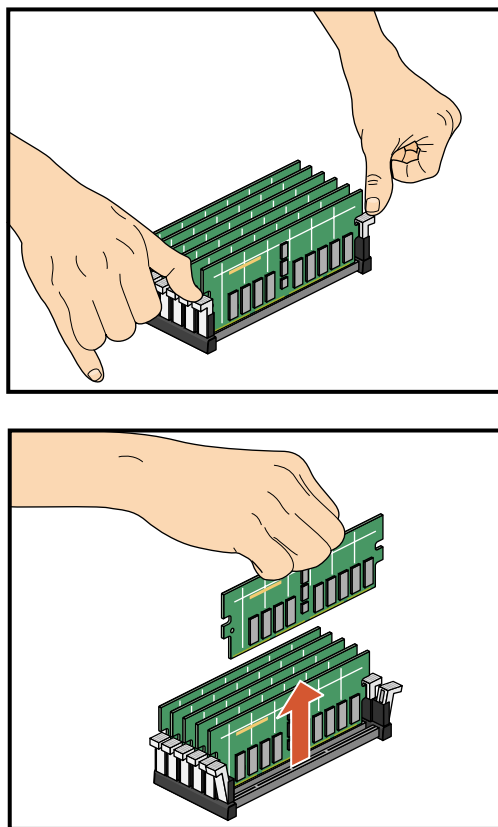
Figure 3-33   Removing a DIMM

L1 Controller Display

The L1 controller, which is used to monitor and manage the base compute module of the Silicon Graphics Prism platform, has a display located on the front panel as shown in Figure 3-34. Every Silicon Graphics Prism platform module is factory- shipped with an L1 controller display. Note that the optional TP900 mass storage module is an exception and does not have an L1.

This section describes how to replace an L1 controller display panel for a base compute module. (You can also use these instructions to replace an L1 controller display panel for an XG2N, system expansion, or CMPX module.)

L1 controller display
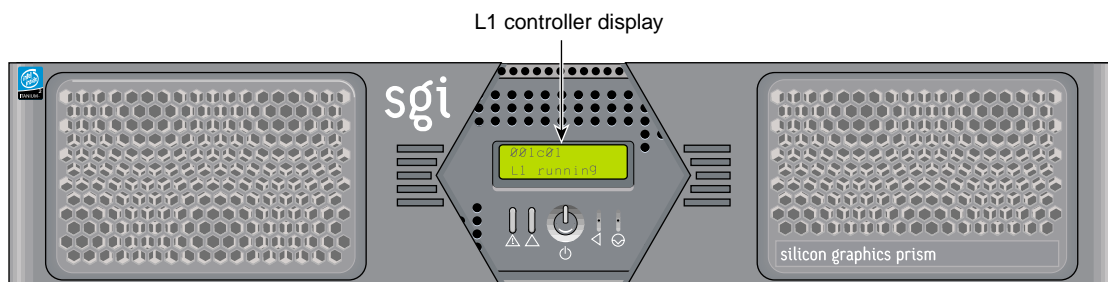
Figure 3-34   L1 Controller Display on Base Compute Module

To replace an L1 controller display, follow these steps:

1. Power off the system as described in "Powering Off Manually" on page 19.
2. Disconnect all of the cables at the rear of the module.

Warning: Components may be hot. To avoid injury, allow the components to cool for approximately five minutes before you proceed with these instructions.

3. Remove the two screws that secure the module to the front rails of the rack. For shelf-mounted modules, remove the two screws at the rear of the unit.
4. With shelf-mounted rails, use two people to slide the module out and move it to a stable work surface. With optional slide rails, pull the module from the rack until it is stopped by the safety latches.

5. To access the area where the L1 display is replaced, remove the ten Phillips screws shown in Figure 3-35, and lift and open the hinged cover.
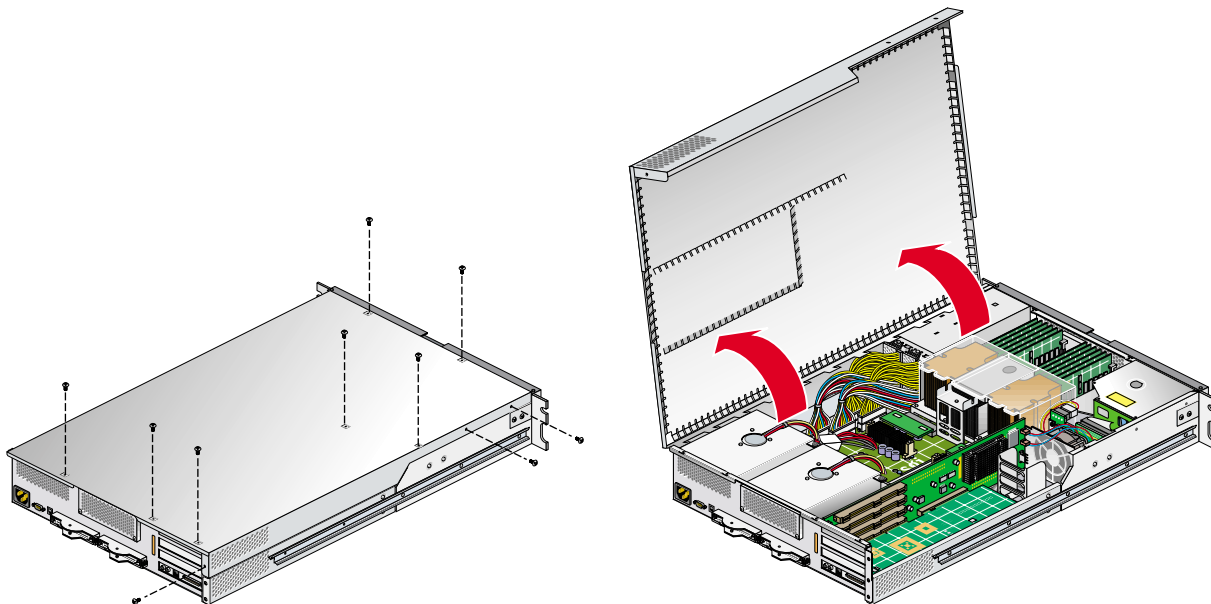
Figure 3-35   Opening the Module to Access the L1 Controller Display

6. On the front panel of your system, remove the front bezel by unscrewing the two Phillips screws holding the bezel to the chassis, as shown in Figure 3-36A.

7. Holding the L1 display cover with one hand, unscrew the single Phillips screw holding the L1 display cover to the chassis, as shown in Figure 3-36B. Gently unhook and pull away the L1 display cover from the chassis.

8. Unscrew the two Phillips screws holding the L1 controller display panel to the L1 display protective cover, as shown in Figure 3-36C.

9. Gently disconnect the L1 controller cable from the connector on the L1 controller display, as shown in Figure 3-36D.
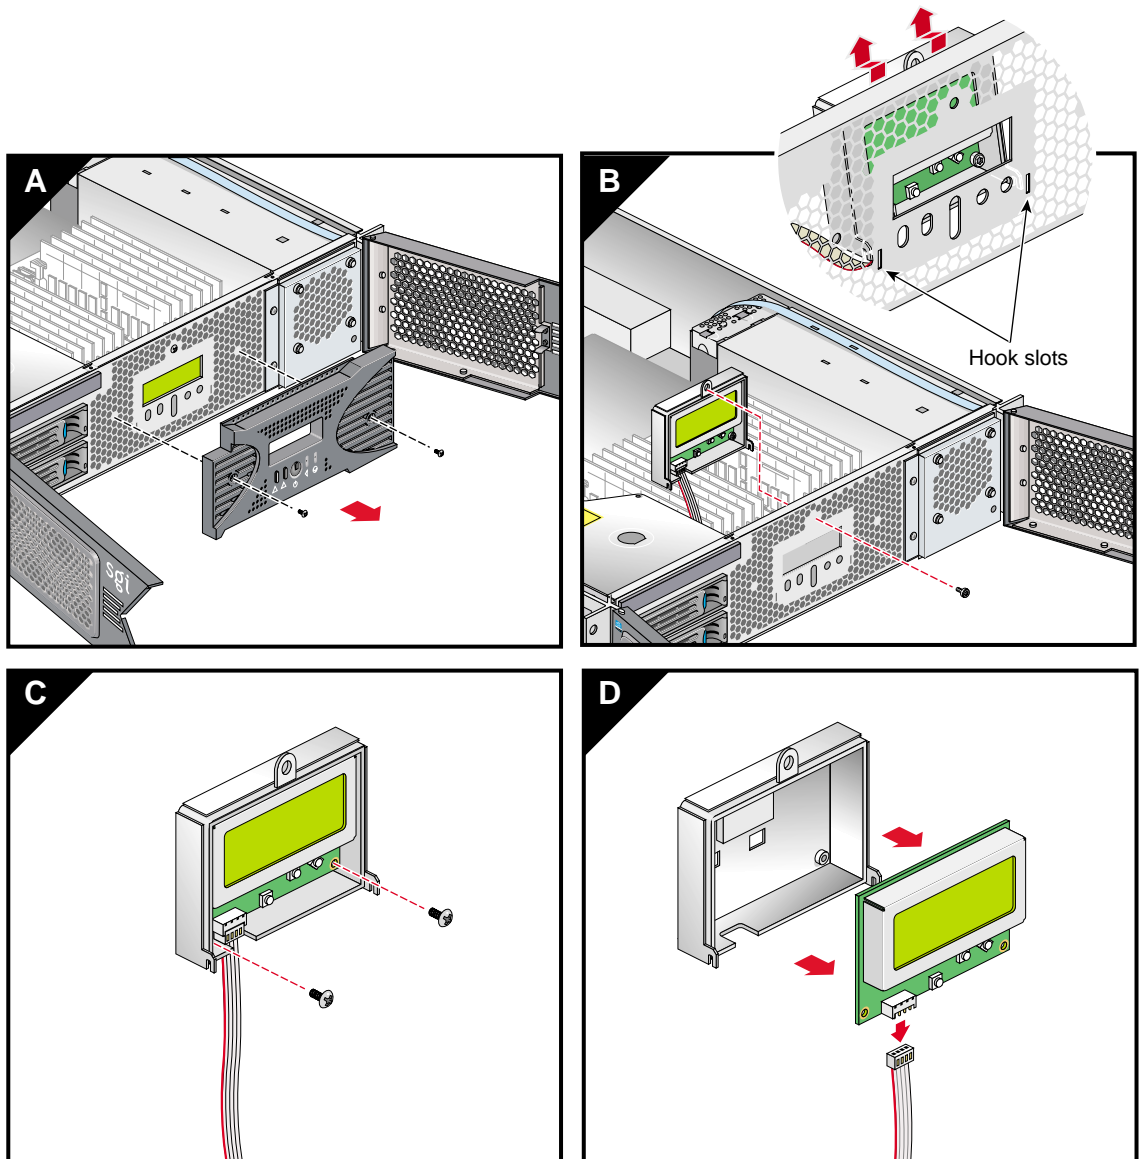
Figure 3-36   Removing the L1 Controller Display Panel

10. Connect the L1 controller cable to the connector on the new L1 controller display, making sure that the red stripe is to your left, as shown in Figure 3-37A.

11. Align the two screw holes on the L1 controller display with the holes on the L1 display protective cover, and screw in the two Phillips screws, as shown in Figure 3-37B.

12. Hook in the L1 display protective cover onto the slots on the front chassis and, holding the L1 display cover up against the front chassis, screw in the Phillips screw, as shown in Figure 3-37C.

13. Replace the front bezel onto the front chassis of the system by screwing in the two Phillips screws holding the bezel to the chassis, as shown in Figure 3-37D.
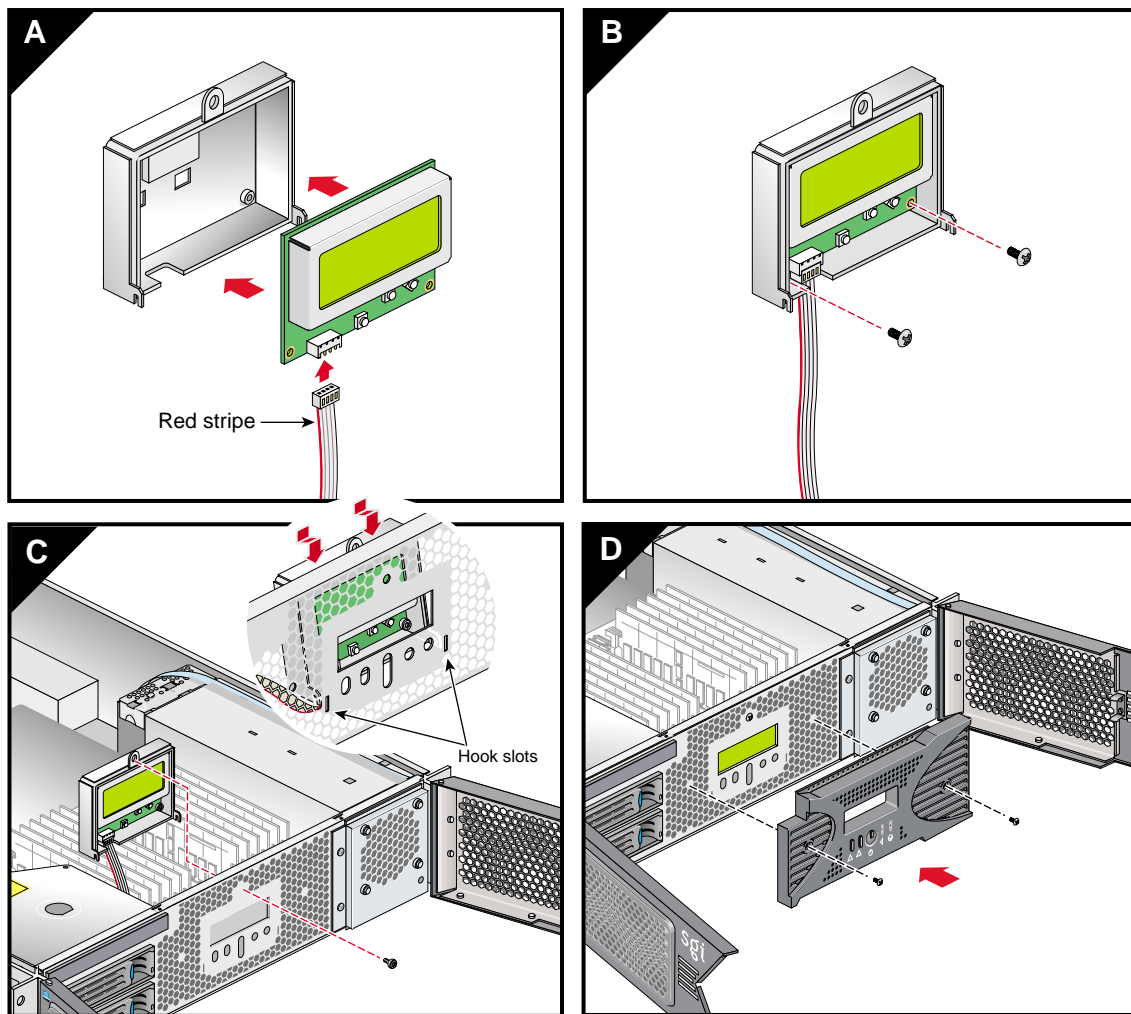
Figure 3-37   Installing an L1 Controller Display Panel

14. Attach the hinged cover and secure it to the module with the ten Phillips screws.

15. If using optional slide rails, press the safety latches on both sides of the module and slide the module into the rack. For shelf-mounted units, use two people to slide the module back into the rack.

16. Install the two screws that secure the module to the front rails of the rack. If shelf rails are used, install the two rear screws.

17. Install all of the cables at the rear of the module.

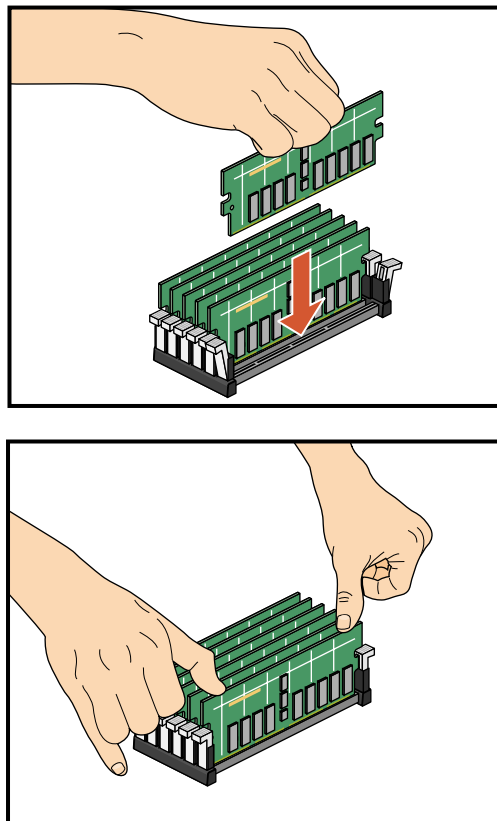18. Power on the Silicon Graphics Prism platform as described in "Power-On the System" on page 16.

Troubleshooting

This chapter provides the following sections to help you troubleshoot your system:

Troubleshooting Chart

Table 4-1 lists recommended actions for problems that can occur on your system. For problems that are not listed in this table, use the SGI Electronic Support system to help solve your problem or contact your SGI system support engineer (SSE). More information about the SGI Electronic Support system is provided later in this chapter.

Table 4-1   Troubleshooting Chart

Problem Description	Recommended Action
The system will not power on.	Ensure that the power cord of the PDU is seated properly in the power receptacle.
	Ensure that the PDU circuit breaker is on.
	If the power cord is plugged in and the circuit breaker is on, contact your SSE.
An individual module will not power on.	View the L1 display; see Table 4-2 if an error message is present.
	If the L1 controller is not running, contact your SSE.
	Check the connection between the module and its power source.
The system will not boot the operating system.	Contact your SSE.
The service-required LED illuminates on a module.	View the L1 display of the failing module; see Table 4-2 for a description of the error message.
The failure LED illuminates on a module.	View the L1 display of the failing module; see Table 4-2 for a description of the error message.
The green or yellow LED of a NUMAlink port is not illuminated.	Ensure both ends of the NUMAlink cable are seated properly and the destination module is powered on.
The PWR LED of a populated PCI slot is not illuminated.	Reseat the PCI card.
The fault LED of a populated PCI slot is illuminated (on).	Reseat the PCI card. If the fault LED remains on, replace the PCI card.
The system status LED of the TP900 is amber.	Contact your SSE.

Table 4-1   Troubleshooting Chart **(continued)**

Problem Description	Recommended Action
The power status LED of the TP900 is amber.	Contact your SSE to replace the power supply module. The power supply module also has an amber LED that indicates a fault.
The cooling status LED of the TP900 is amber.	Contact your SSE to replace the cooling module. The cooling module also has an amber LED that indicates a fault.
The amber LED of a disk drive is on.	Replace the disk drive.

L1 Controller Error Messages

Table 4-2 lists error messages that the L1 controller generates and displays on the L1 display. This display is located on the front of the base compute, compute expansion, XG2N, and CMPX modules. For serial number related errors, check with your service provider for documentation on prevention and solutions.

The serial number error messages listed at the end of Table 4-2 are messages that will come across the L1 console (at the L1 or optional L2 prompt). The log being referenced is the L1 log. Users can get the contents by using the log command from an L1 prompt, or l1 log command from an optional L2 prompt.

Actions that could cause serial number error messages include:

- Moving a base compute, XG2N, expansion, or CMPX module from one system to another.
- Replacing the interface board of a system module.
- L1 NVRAM memory failure.
- Incorrect serial number setting on an optional L2 system controller.

Note: In Table 4-2, a voltage warning occurs when a supplied level of voltage is below or above the nominal (normal) voltage by 10 percent. A voltage fault occurs when a supplied level is below or above the nominal voltage by 20 percent.

Table 4-2   L1 Controller Messages

L1 System Controller Message	Message Meaning and Action Needed
Internal voltage messages:	
ATTN: x.xV high fault limit reached @ x.xxV	30-second power-off sequence for the module.
ATTN: x.xV low fault limit reached @ x.xxV	30-second power-off sequence for the module.
ATTN: x.xV high warning limit reached @ x.xxV	A higher than nominal voltage condition is detected.
ATTN: x.xV low warning limit reached @ x.xxV	A lower than nominal voltage condition is detected.
ATTN: x.xV level stabilized @ x.xV	A monitored voltage level has returned to within acceptable limits.
Fan messages:	
ATTN: FAN # x fault limit reached @ xx RPM	A fan has reached its maximum RPM level. The ambient temperature may be too high. Check to see if a fan has failed.
ATTN: FAN # x warning limit reached @ xx RPM	A fan has increased its RPM level. Check the ambient temperature. Check to see if the fan stabilizes.
ATTN: FAN # x stabilized @ xx RPM	An increased fan RPM level has returned to normal.

Table 4-2 L1 Controller Messages **(continued)**

L1 System Controller Message	Message Meaning and Action Needed
Temperature messages: low alt.	
ATTN: TEMP # advisory temperature reached @ xxC xxF	The ambient temperature at the module's air inlet has exceeded 30 ºC.
ATTN: TEMP # critical temperature reached @ xxC xxF	The ambient temperature at the module's air inlet has exceeded 35 ºC.
ATTN: TEMP # fault temperature reached @ xxC xxF	The ambient temperature at the module's air inlet has exceeded 40 ºC.
Temperature messages: high alt.	
ATTN: TEMP # advisory temperature reached @ xxC xxF	The ambient temperature at the module's air inlet has exceeded 27 ºC.
ATTN: TEMP # critical temperature reached @ xxC xxF	The ambient temperature at the module's air inlet has exceeded 31 ºC.
ATTN: TEMP # fault temperature reached @ xxC xxF	The ambient temperature at the module's air inlet has exceeded 35 ºC.
Temperature stable message:	
ATTN: TEMP # stabilized @ xxC/xxF	The ambient temperature at the module's air inlet has returned to an acceptable level.
Power off messages:	
Auto power down in xx seconds	The L1 controller has registered a fault and is shutting down. The message displays every five seconds until shutdown.
Base module appears to have been powered down	The L1 controller has registered a fault and has shut down.
Serial number messages:	
Brick Serial Number mismatch	See L1 log for details.
System Serial Number mismatch	See L1 log for details.
Invalid System Serial Number format	See L1 log for details.
No assigned System Serial Number	See L1 log for details.

Reading Power Supply Status LEDs

Use the LED located on the front (towards the top) of the power supply to read the condition of the power supply. Table 4-3 shows the LED status and the power supply condition the LED status indicates. See "Removing and Replacing Power Supplies" in Chapter 3 for information on removing and replacing a power supply.

Table 4-3   LED Status and Power Supply Condition

LED Status	Power Supply Condition Indicated
Off	If your system has one power supply, it indicates that the power supply is not receiving AC power. If your system has two power supplies, the LED on both power supplies would be Off, and it would indicate that both power supplies are not receiving AC power. Power supplies will not be receiving AC power because either the module is not plugged into power, or an electrical fuse has blown.
Amber	Indicates a fault condition for one of the following reasons: - In a system where you have two power supplies, one of the power supplies is not receiving AC power. - The voltage limit has been exceeded. - The temperature limit has been exceeded. - The current limit has been exceeded.
Blinking Green	The power supply is receiving AC power, but the main primary DC power has not yet activated.
Green	The power supply is operating properly.

SGI Electronic Support

SGI Electronic Support provides system support and problem-solving services that function automatically, which helps resolve problems before they can affect system availability or develop into actual failures. SGI Electronic Support integrates several services so they work together to monitor your system, notify you if a problem exists, and search for solutions to the problem.

Figure 4-1 shows the sequence of events that occurs if you use all of the SGI Electronic Support capabilities.
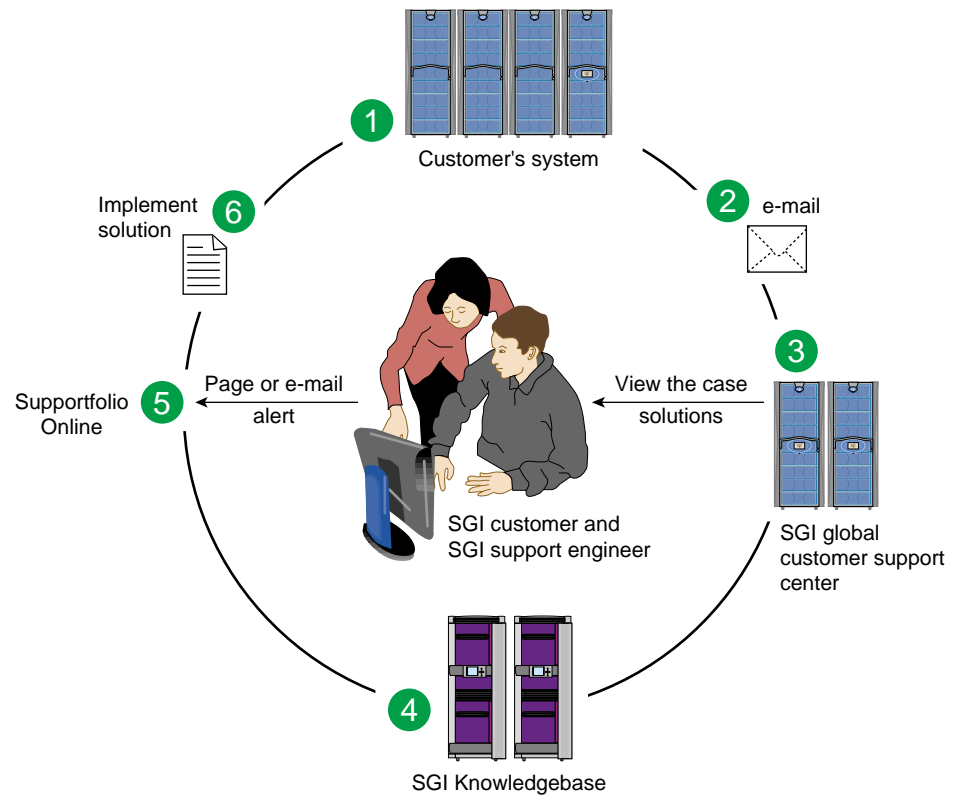
Figure 4-1   Full Support Sequence

The sequence of events can be described as follows:

1. Embedded Support Partner (ESP) monitors your system 24 hours a day.

2. When a specified system event is detected, ESP notifies SGI via e-mail (plain text or encrypted).

3. Applications that are running at SGI analyze the information, determine whether a support case should be opened, and open a case if necessary. You and SGI support engineers are contacted (via pager or e-mail) with the case ID and problem description.

4. SGI Knowledgebase searches thousands of tested solutions for possible fixes to the problem. Solutions that are located in SGI Knowledgebase are attached to the service case.

5. You and the SGI support engineers can view and manage the case by using Supportfolio Online as well as search for additional solutions or schedule maintenance.

6. Implement the solution.

Most of these actions occur automatically, and you may receive solutions to problems before they affect system availability. You also may be able to return your system to service sooner if it is out of service.

In addition to the event monitoring and problem reporting, SGI Electronic Support monitors both system configuration (to help with asset management) and system availability and performance (to help with capacity planning).

The following three components compose the integrated SGI Electronic Support system:

SGI Embedded Support Partner (ESP) is a set of tools and utilities that are embedded in the operating system. ESP can monitor a single system or group of systems for system events, software and hardware failures, availability, performance, and configuration changes, and then perform actions based on those events. ESP can detect system conditions that indicate potential problems, and then alert appropriate personnel by pager, console messages, or e-mail (plain text or encrypted). You also can configure ESP to notify an SGI call center about problems; ESP then sends e-mail to SGI with information about the event.

SGI Knowledgebase is a database of solutions to problems and answers to questions that can be searched by sophisticated knowledge management tools. You can log on to SGI Knowledgebase at any time to describe a problem or ask a question. Knowledgebase searches thousands of possible causes, problem descriptions, fixes, and how-to instructions for the solutions that best match your description or question.

Supportfolio Online is a customer support resource that includes the latest information about patch sets, bug reports, and software releases.

The complete SGI Electronic Support services are available to customers who have a valid SGI Warranty, FullCare, FullExpress, or Mission-Critical support contract. To purchase a support contract that allows you to use the complete SGI Electronic Support services, contact your SGI sales representative. For more information about the various support contracts, see the following website:

http://www.sgi.com/support/customerservice.html

For more information about SGI Electronic Support, see the following website:

http://www.sgi.com/support/es

Regenerating the XF86Config File After Graphics Module Changes

Many details of the graphics system are controlled by the `/etc/X11/XF86Config` file.

After adding or removing an XG2N graphics module to your system, it may be necessary to regenerate the system's `XF86Config` file. If the change was done by your authorized SGI service provider, these processes have most likely been checked and taken care of. If your system graphics hardware was changed by an end user or a local administrator, the information in the following sections may be useful or required for proper operation after the change.

Depending on the settings of two variables, changes to the `XF86Config` file might be done automatically or may need to be done manually. Both methods are described in the following subsections.

Be sure to check your system software release notes for any updates to `XF86Config` usage or modifications.

Automatic Regeneration of the XF86Config File

If the `gen-XF86Config` script detects no hardware changes, it does nothing.

If any graphics hardware not listed in the existing `XF86Config` file is detected, a new file named `XF86Config.N` is created, containing the new hardware. The existing `XF86Config` file is left unchanged. The graphics system will still start this way (though some graphics pipes will not be active). The user can then choose to manually replace the old `XF86Config` file with the new `XF86Config.N` file and reboot.

If some graphics hardware listed in the existing `XF86Config` file is no longer detected, the existing `XF86Config` is renamed `XF86Config.O`, and a new `XF86Config` file is written in its place. This allows the graphics system to start, which it would not do if the `XF86Config` file listed graphics hardware no longer present in the system.

Manually Regenerating the XF86Config File

In order to manually regenerate the `XF86Config` file, you first need to halt the system graphics. This can be done by using `kill -9 X`. Note that this procedure may be complicated if you have set up your system to respawn X at each `init5` occurrence. In this case it would be necessary to remove any scripts that automatically restart X once it is shutdown (probably by modifying the

`inittab` file). Also note that once X is killed, all the subsequent commands have to be entered from a serial system console.

Alternately, you can manually regenerate the `XF86Config` file from the X-term by following these steps:

1. Stop the system graphics from the X-terminal:

    # `<CTRL> <ALT> <BKSPC>`

2. Make a backup copy of the XF86Config file:

    # `mv /etc/X11/XF86Config /etc/X11/XF86Config.backup`

3. Regenerate your XF86Config file using the `fglrxconfig.pl` utility:

    # `/usr/bin/X11/fglrxconfig.pl`

4. Restart the system graphics:

    # `/usr/bin/X11/startx`

If there are no graphics hardware changes, this procedure will do nothing.

If any graphics hardware not listed in the existing `XF86Config` file is detected, a new file named `XF86Config.N` is created, containing the new hardware. The existing `XF86Config` file is left unchanged. The graphics system will still start this way (though some graphics pipes will not be active). The user can then choose to manually replace the old `XF86Config` file with the new `XF86Config.N` file and restart X as in step 4 in the previous sequence.

If some graphics hardware listed in the existing `XF86Config` file is no longer detected, the existing `XF86Config` is renamed `XF86Config.O`, and a new `XF86Config` file is written in its place. This allows the graphics system to start, which it would not do if the `XF86Config` file listed graphics hardware no longer present in the system.

Note: The `XF86Config` file generated by this procedure provides a good starting point. To take advantage of additional features, the file may be customized, as described in the following sections.

Configuring a System for Stereo

This section describes how to configure a system to display stereo images.

Stereo sync is supported only on systems using ImageSync cards.

Note: Simultaneously running stereo and full scene anti-aliasing can require more graphics-card memory than is available, and thus may not always work correctly.

1. Create a copy of the `XF86Config` file to be customized for stereo:

   ```
   # cp /etc/X11/XF86Config /etc/X11/XF86Config.Stereo
   ```

2. Edit the `XF86Config.Stereo` file to include the following line at the end of each "Device" section:

   ```
   Option "Stereo"          "1"
   Option "StereoSyncEnable" "1"
   ```

   (see the "Example "Device" Section for Stereo" on page 111).

3. Edit the `XF86Config.Stereo` file to include the appropriate stereo modes in the "Monitor" section:

   a. Create an appropriate mode (see "Sample Stereo Mode Entries" on page 111).

   b. Add that mode to the "Monitor" section of your `XF86Config.Stereo` file (see the "Example "Monitor" Section for Stereo" on page 111).

   Note: "Mode" and "Modeline" are two alternative formats used to present the same information.

4. Ensure that the monitor supports the high horizontal sync rate setting. Refer to the documentation for the monitor to determine the horizontal sync rate. Modify the HorizSync setting in the "Monitor" section of the `XF86Config.Stereo` file. For example:

   ```
   HorizSync   22-105
   ```

5. Modify the "Screen" section so that you use the appropriate mode setting. For example:

Modes "1280x1024@96" (see the "Example "Screen" Section for Stereo" on page 112).

6. Create a new/etc/X11/xdm/Xservers.Stereo file containing the following line:

   ```
   :0 secure /usr/bin/X11/X :0 -xf86config /etc/X11/XF86Config.Stereo
   ```

7. Edit the /etc/X11/xdm/`xdm-config` file to point to the new Xservers file:

    Replace the line:

    `DisplayManager.servers: /etc/X11/xdm/Xservers`

    with:

    `DisplayManager.servers: /etc/X11/xdm/Xservers.Stereo`

8. Save the file and reboot the system to restart graphics in stereo mode:

    Note that a stereo sync signal will not be present until you run a stereo application. One such application is `ivview`. If your system has `ivview` installed, you can use it to test the stereo configuration by running:
    `ivview /usr/share/data/models/X29.iv`
    and right click to activate the stereo setting on the preferences panel.

Example "Device" Section for Stereo

```
Section "Device"
    Identifier  "SGI SG-0"
    Driver      "fglrx"
    BusId       "PCI:2:0:0"
# === QBS Management ===
    Option "Stereo"         "1"
    Option "StereoSyncEnable" "1"
EndSection
```

Sample Stereo Mode Entries

```
Modeline "1024x768@96" 103.5  1024 1050 1154 1336  768 771 774 807
Modeline "1280x1024@96" 163.28  1280 1300 1460 1600  1024 1027 1033 1063
Modeline "1024x768@100" 113.309 1024 1096 1208 1392 768 769 772 814
Modeline "1024x768@120" 139.054 1024 1104 1216 1408 768 769 772 823 +hsync +vsync
Modeline "1280x1024@100" 190.960 1280 1376 1520 1760 1024 1025 1028 1085 +hsync +vsync
Mode "1280x1024_96s_mirage"
    DotClock     152.928
    HTimings     1280 1330 1390 1500
    VTimings     1024 1026 1030 1062
EndMode
```

Example "Monitor" Section for Stereo

```
Section "Monitor"
    Identifier  "Stereo Monitor"
```

```
    HorizSync   30-96      # multisync
    VertRefresh 50-160     # multisync
    Modeline "1024x768@96" 103.5  1024 1050 1154 1336  768 771 774 807
EndSection
```

Example "Screen" Section for Stereo

```
Section "Screen"
    Identifier          "Screen SG-0"
    Device              "SGI SG-0"
    Monitor             "Stereo Monitor"
    DefaultDepth         24
    SubSection          "Display"
        Depth   24
        Modes   "1280x1024@96"
    EndSubSection
EndSection
```

Configuring a System for Full Scene Anti-Aliasing

This section describes how to configure a system for global or per-window full scene anti-aliasing.

Note: Simultaneously running stereo and full scene anti-aliasing can require more graphics-card memory than is available, and thus may not work correctly.

1. Create a copy of the `XF86Config` file to be customized for full scene anti-aliasing:

   ```
   # cp /etc/X11/XF86Config /etc/X11/XF86Config.AntiAlias
   ```

   Note: Automatically generated `XF86Config` files should contain the customized multi-sample positions shown in on page 114. If these values are not already present, adding them will significantly improve the quality of your output.

2. Edit the `XF86Config.AntiAlias` file to include the following line at the end of each "Device" section:

   ```
   Option "FSAAScale" "X"
   ```

   where *X* is 1, 2, 4, or 6 (see the example "Device" section on page 114).

   Note: Per-window full scene anti-aliasing is accomplished by setting "FSAAScale" to 1. The anti-aliasing level may then be set by the appropriate selection of visuals.
   Global anti-aliasing is accomplished by setting "FSAAScale" to 2, 4, or 6. In this case, the setting will apply to all OpenGL windows, regardless of the visual being displayed.

3. Create a new/etc/X11/xdm/Xservers.AntiAlias file containing the following (all on one line):

   ```
   :0 secure /usr/bin/X11/X :0 –xf86config /etc/X11/XF86Config.AntiAlias
   ```

4. Edit the /etc/X11/xdm/xdm–config file to point to the new Xservers file:

   Replace the line:

   ```
   DisplayManager.servers: /etc/X11/xdm/Xservers
   ```

   with:

   ```
   DisplayManager.servers: /etc/X11/xdm/Xservers.AntiAlias
   ```

5. Stop the system graphics from the X-terminal:

 # <CTRL> <ALT> <BKSPC>

6. Restart graphics:

 # /usr/bin/X11/startx

Example "Device" Section for Full Scene Anti-Aliasing

```
Section "Device"
    Identifier  "SGI SG-0"
    Driver      "fglrx"
    BusId       "PCI:2:0:0"
# === FSAA Management ===
    Option "FSAAScale"              "1"
    Option "FSAADisableGamma"       "no"
    Option "FSAACustomizeMSPos"     "yes"
    Option "FSAAMSPosX0"            "0.250000"
    Option "FSAAMSPosY0"            "0.416666"
    Option "FSAAMSPosX1"            "0.083333"
    Option "FSAAMSPosY1"            "0.083333"
    Option "FSAAMSPosX2"            "0.416666"
    Option "FSAAMSPosY2"            "0.750000"
    Option "FSAAMSPosX3"            "0.750000"
    Option "FSAAMSPosY3"            "0.916666"
    Option "FSAAMSPosX4"            "0.583333"
    Option "FSAAMSPosY4"            "0.250000"
    Option "FSAAMSPosX5"            "0.916666"
    Option "FSAAMSPosY5"            "0.583333"
EndSection
```

Configuring a System for Dual-Channel Operation

To configure a system for dual-channel operation, follow the steps in this section.

Note: If any pipes managed by an X server have their second channel enabled, then every pipe managed by that X server must have its second channel enabled.

Note: Both channels on a pipe must have the same display resolution.

1. Create a copy of the XF86Config file to be customized for dual-channel operation:

    ```
    # cp /etc/X11/XF86Config /etc/X11/XF86Config.DualChannel
    ```

2. Edit the `XF86Config.DualChannel` file to include the following line at the end of each "Device" section:

    ```
    Option "DesktopSetup" mode
    ```

    where *mode* is one of the following:

    ```
    "0x00000100" [this mode clones the managed area]
    "0x00000200" [this mode scales the managed area by 2 horizontally]
    "0x00000300" [this mode scales the managed area by 2 vertically]
    ```

    (see the example "Device" section on page 116).

    Note: All pipes managed by the same X server must be set to the same mode.

3. When using monitors or monitor cables which do not conform to the VESA Display Data Channel (DDC) standard, append the following entry in the "Device" section to enable proper display configuration:

    ```
    Option "NoDDC" "on"
    ```

4. Create a new /etc/X11/xdm/Xservers.DualChannel file containing the following line:

    ```
    :0 secure /usr/bin/X11/X :0 -xf86config
    /etc/X11/XF86Config.DualChannel
    ```

5. Edit the /etc/X11/xdm/xdm-config file to point to the new Xservers file:

    Replace the line:

    ```
    DisplayManager.servers: /etc/X11/xdm/Xservers
    ```

with:

```
DisplayManager.servers: /etc/X11/xdm/Xservers.DualChannel
```

6. Stop the system graphics from the X-terminal:

    # <CTRL> <ALT> <BKSPC>

7. Restart graphics:

    # /usr/bin/X11/startx

Example "Device" Section for Dual Channel

```
Section "Device"
    Identifier  "SGI SG-0"
    Driver      "fglrx"
    BusId       "PCI:2:0:0"
    Option      "DesktopSetup" "0x00000200"
EndSection
```

Configuring a System for SwapReady

To configure a system for SwapReady, it must first be enabled in all pipes connected to a particular ImageSync board, as described in these steps:

1. On the screen 0 desktop, open a winterm.

2. Within that winterm, run `xsetmon`, setting the target to the pipe you are configuring and setting the display to screen 0. For example:

   ```
   # xsetmon -target :0.1 -display :0.0
   ```

   would start the GUI on screen 0 to configure pipe 1.

3. Click the two radio buttons on the GUI main menu:

   ```
   Swap Buffers on Vertical Blank: On
   GLX Swap Barrier Extension: On
   ```

4. Save the configuration by clicking the **Load** button.

5. Repeat steps 1 through 4 for each additional pipe connected to the same ImageSync board.

6. Log out from the desktop, then log back on.

Configuring the System for External Genlock or Framelock

External genlock and framelock may be used on systems with ImageSync2 cards.

To configure your system to receive an external genlock or framelock signal you must run the `setmon` command with the appropriate options.

Before running setmon, use `printenv DISPLAY` to ensure that the DISPLAY environment variable is set to the local system (for example, :0.0). If it is not, use `setenv DISPLAY :0.0` to change it (substituting other numbers for :0.0 if appropriate).

To set the system for genlock, execute the following command:

`# setmon -p`*pipenumber* `-g` *graphicsformat*

where *pipenumber* is the pipe to which this setting should be applied, and *graphicsformat* is one of the timings (modes) listed in the "Monitor" section of the `/etc/X11/XF86Config` file.

To set the system for framelock, execute the following command:

`# setmon -p`*pipenumber* `-L`*videoformat*   *graphicsformat*

where *pipenumber* is the pipe to which this setting should be applied, *videoformat* is the input video format to be used as a framelock source, and *graphicsformat* is one of the framelock-certified timings (modes) listed in the "Monitor" section of the `/etc/X11/XF86Config` file that is compatible with the chosen input video format (Table 4-4 on page 119 provides a list of compatible formats).

Note: The default behavior of `setmon` is to load the new format for the current session only and to prompt for input to determine if the format should be saved as the default. To save the new format as the default you must be logged in as root.

For more information about the `setmon` command, see the `setmon` manpage (`man setmon`).

Note: Framelock-certified timings will include an "f" appended to their name (i.e., "1280x1024_5994f" is certified for NTSC (525 line) video timing).

Table 4-4   Input Video Formats (Framelock)

Input Video Format (Framelock Source)	Format Name	Compatible Graphics Formats
525 line at 59.94Hz (NTSC)	525 (or use the alias NTSC)	1280x1024_5994f 1920x1154_5994f
625 line at 50Hz (PAL)	625 (or use the alias PAL)	1280x1024_50f 1920x1154_50f
720-line progressive-scan at 59.94Hz	720p_5994	1920x1154_5994f
720-line progressive-scan at 60Hz	720p_60	1280x1024_60f 1920x1154_60f 1920x1200_60f
1080-line progressive-scan at 25Hz	1080p_25	1280x1024_50f 1920x1154_50f
1080-line interlaced at 25Hz	1080i_25	1280x1024_50f 1920x1154_50f
1080-line progressive-scan at 29.97Hz	1080p_2997	1920x1154_5994f
1080-line interlaced at 29.97Hz	1080i_2997	1920x1154_5994f
1080-line progressive-scan at 30Hz	1080p_30	1280x1024_60f 1920x1154_60f 1920x1200_60f
1080-line interlaced at 30Hz	1080i_30	1280x1024_60f 1920x1154_60f 1920x1200_60f

Configuring Monitor Positions

When an X-Server is managing multiple monitors, it needs to know their relative positions in order to properly handle cursor cross-over locations.

The monitor positions are specified in the "ServerLayout" section of the `/etc/X11/XF86Config` file as follows:

Each screen is listed, followed by a list of the screens above, below, to the left, and to the right of it (in that order). Figure 4-2 and and the following subsection show an example of four monitors arranged in a line.

Programs started by clicking on an icon appear on the screen from which you invoked them. Note that once a program has been launched, it is not possible to move it from one screen to another.

Figure 4-2   Four Monitors in a Line

Example "ServerLayout" Section for Four Monitors in a Line

```
Section "ServerLayout"
    Identifier  "Four-in-a-Line"
    Screen "Screen SG-0"        ""      ""      ""       "Screen SG-1"
    Screen "Screen SG-1"        ""      ""      "Screen SG-0"  "Screen SG-2"
    Screen "Screen SG-2"        ""      ""      "Screen SG-1"  "Screen SG-3"
    Screen "Screen SG-3"        ""      ""      "Screen SG-2"  ""
    InputDevice "Mouse1" "CorePointer"
    InputDevice "Keyboard1" "CoreKeyboard"
EndSection
```

Figure 4-3 and the subsection following it show an example of four monitors arranged in a square.
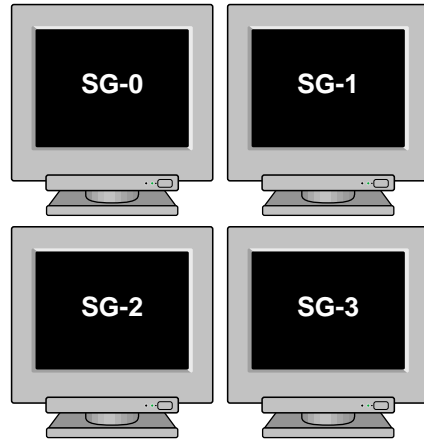

Figure 4-3   Four Monitors in a Square

Example "ServerLayout" Section for Four Monitors in a Square

```
Section "ServerLayout"
    Identifier  "Four-in-a-Square"
    Screen "Screen SG-0"         ""       "Screen SG-2"      ""        "Screen SG-1"
    Screen "Screen SG-1"         ""       "Screen SG-3"    "Screen SG-0"    ""
    Screen "Screen SG-2"       "Screen SG-0"     ""      ""   "Screen SG-3"
    Screen "Screen SG-3"       "Screen SG-1"     ""    "Screen SG-2"    ""
    InputDevice "Mouse1" "CorePointer"
    InputDevice "Keyboard1" "CoreKeyboard"
EndSection
```

Configuring Monitor Types

The system graphics cards support both analog and digital monitors. The type of monitor connected to each graphics card is specified in the "Device" sections of the `/etc/X11/XF86Config` file.

Table 4-5 lists the allowable options for the MonitorLayout line. If the line is not present, both channels default to AUTO.

Table 4-5   Options for Monitor Layout

Monitor Type	Meaning
AUTO	Automatically select monitor type (default)
TMDS	Digital monitor
CRT	Analog monitor
NONE	No monitor

The format is:

```
Option      "MonitorLayout" "channel1type, channel2type"
```

where *channel1type* is the type (AUTO, TMDS, CRT or NONE) of monitor attached to channel 1 (the left DVI-I connector) for this pipe, and
channel2type is the type (AUTO, TMDS, CRT or NONE) of monitor attached to channel 2 (the right DVI-I connector) for this pipe.

Example "Device" Section for Use With Two Analog Monitors

```
Section "Device"
    Identifier  "SGI SG-0"
    Driver      "fglrx"
    BusId       "PCI:2:0:0"
    Option      "MonitorLayout" "CRT, CRT"
EndSection
```

Technical Specifications

This appendix contains technical specification information about your system, as follows:

- "Environmental System Specifications" on page 124
- "Compute Module Specifications" on page 125
- "XG2N Module Connectors" on page 131
- "CMPX Module Specifications" on page 127
- "Rack Specifications" on page 128
- "SGI TP900 Storage Module Specifications" on page 129
- "Non-proprietary I/O Port Specifications" on page 129

Environmental System Specifications

Table A-1 lists the environmental specifications of the Silicon Graphics Prism.

Table A-1   Environmental Specifications

Characteristic	Specification
Temperature, operating	+5 ºC (+41 ºF) to +35 ºC (+95 ºF) (up to 1500 m [5,000 ft]) +5 ºC (+41 ºF) to +30 ºC (+86 ºF) (1500 m to 3000 m [5,000 ft to 10,000 ft])
Temperature, non-operating	-40 ºC (-40 ºF) to +60 ºC (+140 ºF)
Humidity	10% to 95% RH, noncondensing
Altitude	Sea level to 40,000 ft (nonoperating) Sea level to 10,000 ft (3000 m) (operating)

Compute Module Specifications

Table A-2 lists the bandwidth characteristics of the compute module.

Table A-2   Bandwidth Characteristics of the Compute Module

Characteristic	Peak Bandwidth	Sustainable Bandwidth
NUMAlink channel	6.4 GB/s full duplex 3.2 GB/s each direction	~2.8 GB/s each direction
Main memory	Up to 10.8 GB/s	Up to 10.8 GB/s
Front Side Bus (FSB)	6.4 GB/s	~6.4 GB/s

Table A-3 summarizes the general features of the compute module.

Note: The expansion compute module does not include an optional IO10 PCI card.

Table A-3   General Features of the Compute Module

Feature	Base Compute Module	Expansion Compute Module
NUMAlink ports	2 (3.2 GB/s each direction)	2 (3.2 GB/s each direction)
Serial console port	1	1
L1 port	1	1
RT interrupt input port	1 (not functional under Linux)	
RT interrupt output port	1 (not functional under Linux)	
Ethernet port	1 10BaseT/100BaseT/1000BaseT	
Serial RS232/422 (4 ports)	4 serial ports (external)	
serial ATA port (internal)	2 serial ATA	
3.5-in. drive bay	1 or 2	
64-bit Intel processor	2	1 or 2

Table A-3   General Features of the Compute Module **(continued)**

Feature	Base Compute Module	Expansion Compute Module
Memory	up to 24 GB	up to 24 GB
Expansion slot	2 PCI-X	4 PCI-X

Note: Power consumption in the XG2N module is slightly higher (approximately 563 watts) than in the base compute module. The XG2N always requires use of two power supplies at all times. Otherwise, the basic physical and electrical specifications are the same.

Table A-4 lists the specifications for the compute module.

Table A-4   Compute Module Specifications

Characteristic	Specification
Height	3.44 in. (8.74 cm)
Width	17.06 in. (43.33 cm)
Depth	27 in. (68.58 cm) (with bezel)
Weight	50 lb (22.7 kg) maximum configuration[a]
Noise	6.0 Bells sound power, up to 30 °C
Heat dissipation	1315 Btu/hr maximum
Input power	120 - 240 VAC

a. Weight will vary depending on whether the system has one or two power supplies, on the amount of DIMMs installed, and on whether you have one or two disk drives in your system.

CMPX Module Specifications

Table A-5 lists the bandwidth characteristics of the CMPX module.

Table A-5   Bandwidth Characteristics of the CMPX Module

Characteristic	Peak Bandwidth	Sustainable Bandwidth
NUMAlink channel	6.4 GB/s full duplex 3.2 GB/s each direction	~2.8 GB/s each direction
Main memory	Up to 10.8 GB/s	Up to 10.8 GB/s
Front Side Bus (FSB)	6.4 GB/s	~6.4 GB/s

Table A-6 lists the specifications for the CMPX module.

Table A-6   CMPX Module Specifications

Characteristic	Specifications
Height	3.44 in. (8.80 cm)
Width	17.06 in. (43.36 cm)
Depth	27 in. (68.58 cm) (with bezel)
Weight	50 lb (22.7 kg) maximum configuration[a]
Input voltage	120 - 240 VAC

a. Weight will vary slightly depending on whether the module has processors, one or two power supplies, on the number of DIMMs installed, and on PCI cards in the unit.

Rack Specifications

The Silicon Graphics Prism modules can be housed in optional short (17U) or tall (39U) racks.

Note: One U is 1.75 in. (4.45 cm).

Table A-7 lists the specifications of the short rack.

Table A-7   Short Rack Specifications (with Skins)

Characteristic	Specification
Height	36.06 in. (916 mm)
Width	25.38 in. (645 mm)
Depth	40.63 in. (1032 mm)
Weight (maximum)	610 lb (277 kg)
Shipping weight (maximum)	685 lb (312 kg)

Table A-8 lists the specifications of the tall rack.

Table A-8   Tall Rack Specifications

Characteristic	Specification
Height	75.82 in. (1925.83 mm)
Width	23.62 in. (599.95 mm)
Depth	41.25 in. (1048 mm)
Weight (maximum)	1,366 lb (620 kg)
Shipping weight (maximum)	1,547 lb (702 kg)

SGI TP900 Storage Module Specifications

Table A-9 lists the specifications of the SGI TP900 storage module.

Table A-9   TP900 Storage Module Specifications

Characteristic	Specification
Height	3.37 in. (85.7 mm)
Width	17.6 in. (447 mm)
Depth	21.46 in. (545 mm)
Input power	100 - 254 VAC (~175 W)
Weight:	
Maximum configuration	48.5 lb (22 kg)
Empty enclosure	14.3 lb (6.5 kg)

Non-proprietary I/O Port Specifications

This section provides pin assignment information for the non-proprietary connectors on the following components:

- Compute module (base and expansion)
- XG2N module
- CMPX module
- Optional L2 controller
- TP900 storage module

Compute Module

Table A-10 lists the non-proprietary connectors that are located on the rear panel of the compute module (see Figure A-1). The third column of the table indicates where you can find the pin assignments for these connectors.

Table A-10   Compute Module Connectors

Port	Connector	Pin Assignments
Optional PCI serial ports [a]	DB-9	See Figure A-6 on page 136
Console port	DB-9	See Figure A-6 on page 136
Ethernet port[b]	RJ-45	See Figure A-7 on page 137 and Table A-14 on page 137
External serial port	IEEE 1284 36-pin MDR multi-port serial adapter	See Figure A-8 on page 138 and Table A-15 on page 138
RT interrupt input and output ports[c]	Stereo jack	Not supported under Linux
L1 port	USB type B	See Figure A-10 on page 141 and Table A-17 on page 141

a. Optional serial ports are available on a PCI card.

b. These connectors are available only when the compute module has an IO10 PCI card.

c. These connectors are not functionally supported under SGI Linux + ProPack.
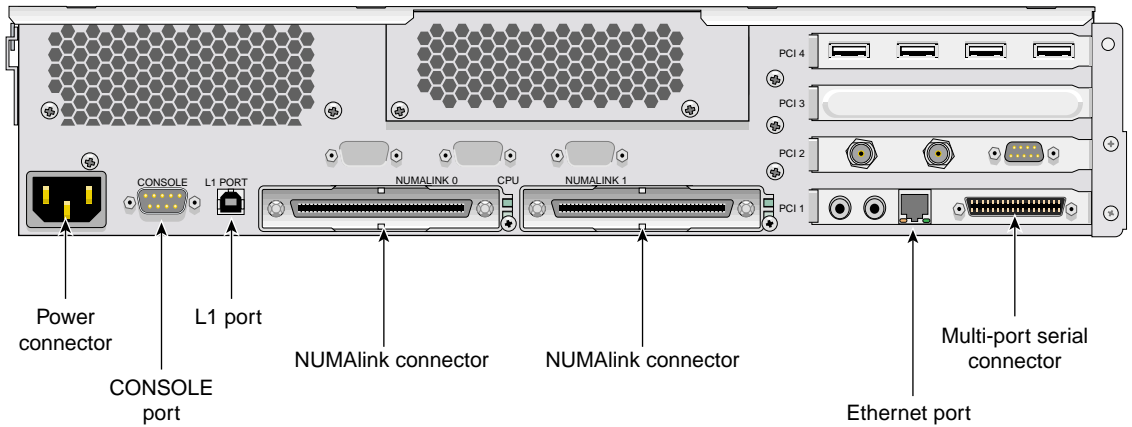
Figure A-1   Rear Panel of Compute Module

Note: The RT interrupt input and RT interrupt output functionality of the IO10 PCI card is not supported under SGI Linux + ProPack.

XG2N Module Connectors

The XG2N module contains two graphics pipes, each capable of supporting two display devices. The graphics connector and pinouts are shown in Figure A-2 and Table A-11.
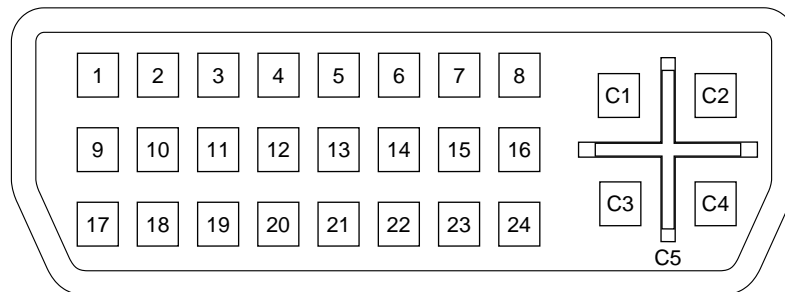

Figure A-2   DVI-I Connector Pinout

The console and L1 connectors on the XG2N module are the same as the compute module.

Table A-11   DVI-I Connector Signals and Pins

Pin	Function	Pin	Function
1	T.M.D.S.[a] Data 2-	16	Hot Plug Detect
2	T.M.D.S. Data 2+	17	T.M.D.S. Data 0-
3	T.M.D.S. Data 2/4 Shield	18	T.M.D.S. Data 0+
4	T.M.D.S. Data 4-	19	T.M.D.S. Data 0/5 Shield
5	T.M.D.S. Data 4+	20	T.M.D.S. Data 5-
6	DDC Clock	21	T.M.D.S. Data 5+
7	DDC Data	22	T.M.D.S. Clock Shield
8	Analog Vertical Sync	23	T.M.D.S. Clock+
9	T.M.D.S. Data 1-	24	T.M.D.S. Clock-
10	T.M.D.S. Data 1+		
11	T.M.D.S Data 1/3 Shield	C1	Analog Red Video Out
12	T.M.D.S. Data 3-	C2	Analog Green Video Out
13	T.M.D.S. Data 3+	C3	Analog Blue Video Out
14	+5V Power	C4	Analog Horizontal Sync
15	Ground (for +5V)	C5	Analog Common Ground

a. Transition Minimized Differential Signaling

CMPX Module

Table A-12 lists the non-proprietary connectors that are located on the rear panel of the optional CMPX module (see Figure A-3). The third column of the table indicates where you can find the pin assignments for these connectors.

Table A-12   CMPX Module Connectors

Port	Connector	Pin Assignments
Console port	DB-9	See Figure A-6 on page 136
L1 port	USB type B	See Figure A-10 on page 141 and Table A-17 on page 141
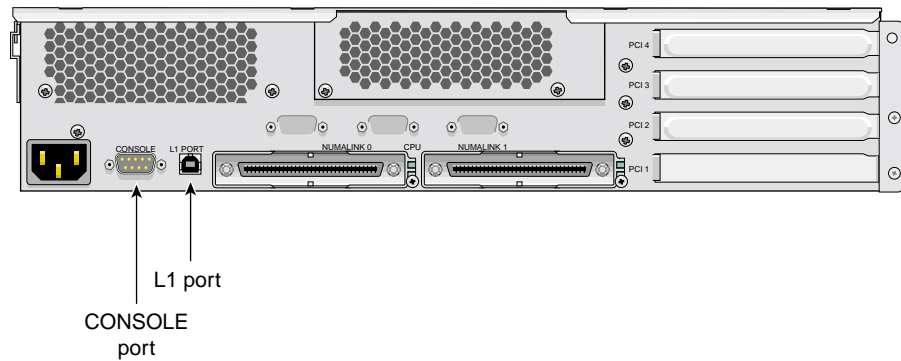
Figure A-3   Non-Proprietary Connectors on Rear Panel of CMPX Module

L2 Controller

Table A-13 lists the non-proprietary connectors that are located on the rear panel of the optional L2 controller (see Figure A-4). The third column of the table indicates where you can find the pin assignments for these connectors.

Table A-13   L2 Controller Connectors

Port	Connector	Pin Assignments
Console (serial port)	DB-9	See Figure A-6 on page 136
Modem (serial port)	DB-9	See Figure A-6 on page 136
Ethernet port	RJ-45	See Figure A-7 on page 137 and Table A-14 on page 137
L1 ports (four ports)	USB type A	See Figure A-9 on page 140 and Table A-16 on page 140
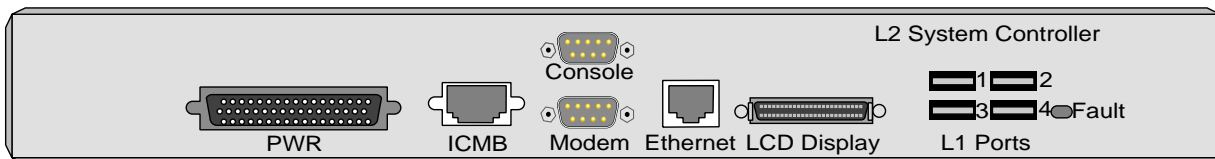
Figure A-4   Non-Proprietary Connectors on Rear of Optional L2 Controller

SGI TP900 Storage Module

Figure A-5 shows the two SCSI port connectors on the rear panel of the TP900 storage module.

Note that your system will need an optional SCSI PCI card installed to interface with the TP900 storage option module.

SCSI port connectors
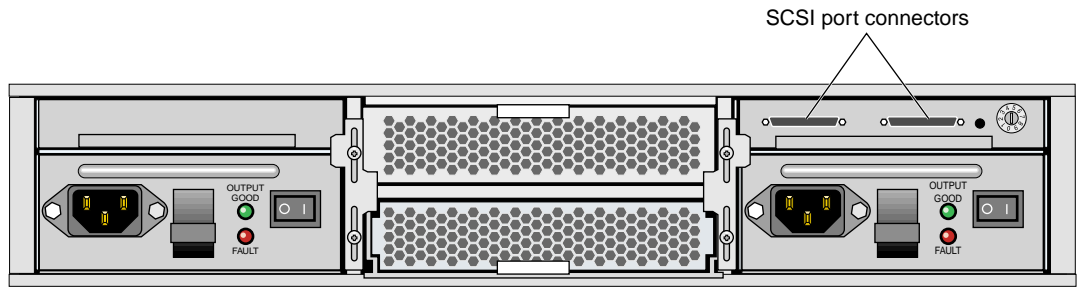
Figure A-5   Non-proprietary Connectors on Rear Panel of TP900 Module

DB-9 Connector

Figure A-6 shows the DB-9 connector pin assignments. This connector is used for the console port of the compute module, XG2N module, and CMPX module, and the console and modem ports of the optional L2 controller.

Pin 2
Receive
data
(RD)

Pin 4
Data terminal
ready (DTR)
(not used)

Pin 3
Transmit
data (TD)

Pin 1
Data carrier
detect (DCD)
(not used)

Pin 5
Ground

Pin 6
Data set
ready (DSR)
(not used)

Pin 9
Ringing indicator (RI)
(not used)

Pin 7
Request to
send (RTS)

Pin 8
Clear to
send (CTS)
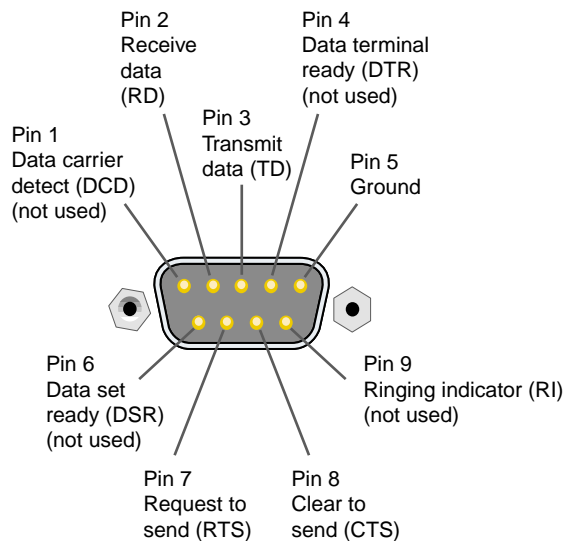
Figure A-6   DB-9 Connector Pin Assignments

RJ-45 Connector

Figure A-7 shows the pin locations for the RJ-45 connector on the IO10 PCI card and the L2 controller. Table A-14 shows the pin assignments for the RJ-45 connector.
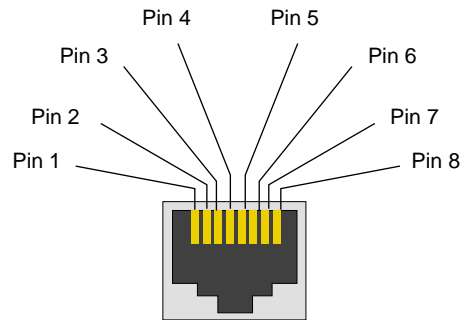
Figure A-7   RJ-45 Connector Pin Assignments

Table A-14   Ethernet Connector Pin Assignments

10/100BaseT Ethernet Pinouts Pin	Assignment	1000BaseT Ethernet Pinouts Pin	Assignment
1	Transmit +	1	Transmit/Receive 0+
2	Transmit –	2	Transmit/Receive 0–
3	Receive +	3	Transmit/Receive 1+
4	Not used	4	Transmit/Receive 2+
5	Not used	5	Transmit/Receive 2–
6	Receive –	6	Transmit/Receive 1–
7	Not used	7	Transmit/Receive 3+
8	Not used	8	Transmit/Receive 3–

External Multi-port Serial Adapter Connector

Figure A-8 shows the connector pin locations for the multi-port serial adapter connector located on the IO10 PCI card.
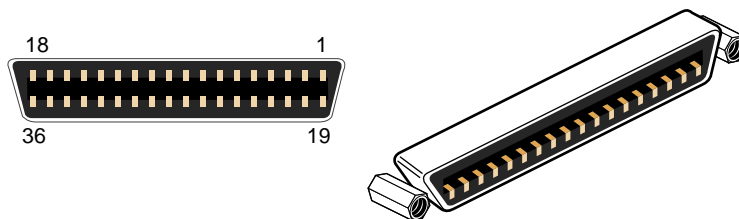
Figure A-8   Pin Number Locations for 36-pin MDR Connector

Table A-15 lists the pin assignments for the multi-port serial adapter connector.

Table A-15   Multi-port Serial Adapter Pinouts

Pin	Assignment	Pin	Assignment
1	GND	19	S0 DSR
2	S0 DCD	20	S0 RTS
3	S0 RXD	21	S0 CTS
4	S0 TXD	22	S0 RI
5	S0 DTR	23	GND
6	S1 DSR	24	S1 DCD
7	S1 RTS	25	S1 RXD
8	S1 CTS	26	S1 TXD
9	S1 RI	27	S1 DTR
10	GND	28	S2 DSR
11	S2 DCD	29	S2 RTS
12	S2 RXD	30	S2 CTS
13	S2 TXD	31	S2 RI

Table A-15   Multi-port Serial Adapter Pinouts **(continued)**

Pin	Assignment	Pin	Assignment
14	S2 DTR	32	GND
15	S3 DSR	33	S3 DCD
16	S3 RTS	34	S3 RXD
17	S3 CTS	35	S3 TXD
18	S3 RI	36	S3 DTR

USB Type A Connector

Figure A-9 shows the USB type A connector that is used for USB ports 1 through 4 of the optional L2 controller that connects to the compute and/or CMPX modules.
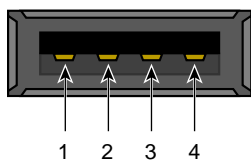
Figure A-9   Pin Number Locations for USB Type A Connector

Table A-16 lists the pin assignments.

Table A-16   Pin Assignments for USB (Type A) Connector

Signal	Color	Pin Number
VCC	Red	1
-Data	White	2
+Data	Green	3
Ground	Black	4

USB Type B Connector

Figure A-10 shows the USB type B connector that is used for the USB L1 port of the compute module, XG2N module and the CMPX module. Table A-17 lists the pin assignments.
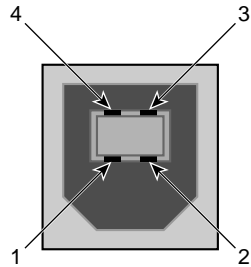
Figure A-10   Pin Number Locations for USB Type B Connector

Table A-17   USB (Type B) Connector Pinouts

Signal	Color	Pin Number
VCC	Red	1
-Data	White	2
+Data	Green	3
Ground	Black	4

Installing Silicon Graphics Prism Modules in a Rack

Depending on how you purchased your Silicon Graphics Prism visualization system, you may need to install or move one or more modules within a rack. This Appendix describes:

- "Choosing Rack Locations" on page 143.

- "Using Silicon Graphics Prism Modules with a Rack" on page 145.

- "Rackmounting with Optional Slide Rails" on page 145.

- "Removing a Rail-mounted Module from a Rack" on page 159

Choosing Rack Locations

The rack locations for your XG2N modules will be partially determined by cable lengths and routings.

Since the Silicon Graphics Prism is a highly-configurable system, it is not practical to list every possible configuration. For valid configurations not shown in Chapter 1 or in this Appendix, consult your SGI sales or support representative.

For power connection examples, see "Positioning and Power for Your Silicon Graphics Prism" in Chapter 1.

Programming L1 Rack and Slot Numbers

Each module in a system contains a Level 1 controller (L1). This controller should be programmed with the location of the brick in which it is installed. Programming the L1 with the correct rack and slot numbers allows for easier system maintenance and diagnostics, and is necessary for predictable operation of the system.

When a module is connected to AC power, the L1 display will indicate its position as follows:

There will be a string of the form XXXyZZ (for example, "001c12").

Decode this string as follows:

"XXX" is the rack number (in this example, rack 001)

"y" is the module type (in this case, c indicates a compute module)

"ZZ" is the "U" number within that rack, counting from the bottom (in this case, slot 12).

If the position indicated on an L1 display is not correct, you should correct it as follows:

1. Connect a serial terminal to the console port on the brick in question.

2. Display the current location setting in the L1:

    ```
    001c12-L1> brick
    rack: 001 slot: 12 partition: 0 type: C source: EEPROM
    ```

    Enter the new rack number:

    ```
    001c12-L1> brick rack 1
    brick rack set to 001
    ```

3. Enter the new slot number:

    ```
    001c12-L1> brick slot 7
    brick slot set to 07
    ```

4. Verify the newly entered information:

    ```
    001c12-L1> brick
    rack: 001 slot: 07 partition: 0 type: C source: EEPROM
    ```

If other bricks need to be changed, repeat steps 1 through 4 for each additional brick.

Using Silicon Graphics Prism Modules with a Rack

This section describes how to install or remove a module in a rack using slide rails.

This rackmounting information is found in the following sections:

- "Rackmounting with Optional Slide Rails" on page 145
- "Removing a Rail-mounted Module from a Rack" on page 159

Rackmounting with Optional Slide Rails

This section describes how to rackmount modules with slide rail assemblies. This section includes the following topics:

- "Determining Module Space Requirements" on page 145
- "Slide Rail Mounting Hardware" on page 146
- "Preparing the Optional Slide Rail Assemblies" on page 147
- "Preparing the Module" on page 150
- "Determining Where to Attach the Slide Rail in the Rack" on page 151
- "Attaching the Slide Rail to the Rack" on page 152
- "Installing Clip Nuts in Rack Rails" on page 154
- "Installing the Module in the Rack" on page 156
- "Adjusting the Position of the Rackmounted Module" on page 158

Determining Module Space Requirements

Table B-1 specifies the space requirements when rackmounting a module in a 19-inch rack.

Table B-1   Rackmounting Space Requirements for Modules

Height	3.44 inches (8.74 cm)
Width	17.06 inches (43.33 cm)
Depth	27 inches (68.58 cm) (with bezel)

Table B-1   Rackmounting Space Requirements for Modules **(continued)**

Weight	50 lb. (22.7 kg) maximum configuration[a]
Required front clearance for module	8.25 inches (20.96 cm)
Required rear clearance for module	10 inches (25.40 cm)
Required side clearance for module	6 inches (15.24 cm) (right) None required for left side.
Required front clearance for rack	36 inches (91 cm)
Required rear clearance for rack	36 inches (91 cm)

a. Weight will vary depending on configuration of memory, PCI cards, and disks.

Slide Rail Mounting Hardware

Table B-2 lists the hardware that you will use to mount each module in a 19-inch rack.

Table B-2   Slide Rail Rackmounting Hardware

Hardware Type	Qty	Usage
Slide rail assembly (includes chassis rail)	2	Allows the module to slide in and out of rack. (The left and right slides are identical.)
2-inch rear mounting bracket	2	Mounts the slide rails to the rear rack rails. (The left and right brackets are identical.)
10-24 x 1/4-inch Phillips screw	10	Secures the chassis rails to the module.
10-32 x 1/2-inch Phillips screw Shoulder washer Bar nut	8 8 4	Secures the slide rails to the rack rails.
10-32 x 1/2-inch Phillips screw Bar nut	4 2	Secures the slide rails to their mounting brackets.
10-32 clip nut	2	Provides a threaded hole for fastening the module front panel to the rack rails.
10-32 x 1/2-inch Phillips screw	2	Fastens the module front panel to the clip nut.

Preparing the Optional Slide Rail Assemblies

The slide rail assembly consists of a chassis rail and a slide rail. You need to remove the chassis rail from the slide rail so that you can install a mounting bracket to the slide rail and attach the chassis rail to the module (see "Preparing the Module" on page 150). To remove the chassis rail from the slide rail, follow these steps:

1. Remove the two slide rail assemblies and the rear mounting brackets from the shipping container.

2. Extend each slide rail assembly until the safety latch snaps into place.

3. Press the safety latch and remove the chassis rail from the slide rail, as shown in Figure B-1.
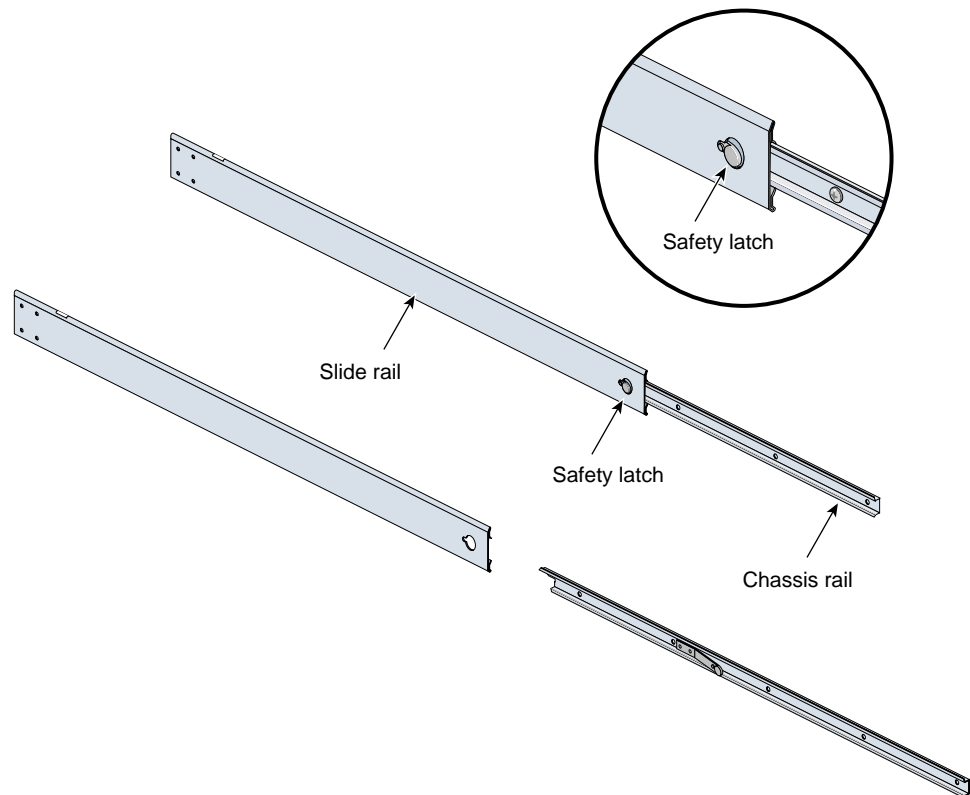
Figure B-1   Removing the Chassis Rail from the Slide Rail

4. Place one of the mounting brackets on the back of the slide rail as shown in Figure B-2. Adjust the position of the mounting bracket on the slide rail according to the depth of the rack.

5. Place a bar nut next to the mounting bracket. Secure the mounting bracket to the slide rail by inserting two 10-32 x 1/2-inch screws through the assembly and into the bar nut as shown in Figure B-2.

6. Repeat steps 4 and 5 to attach a mounting bracket to the other slide rail.
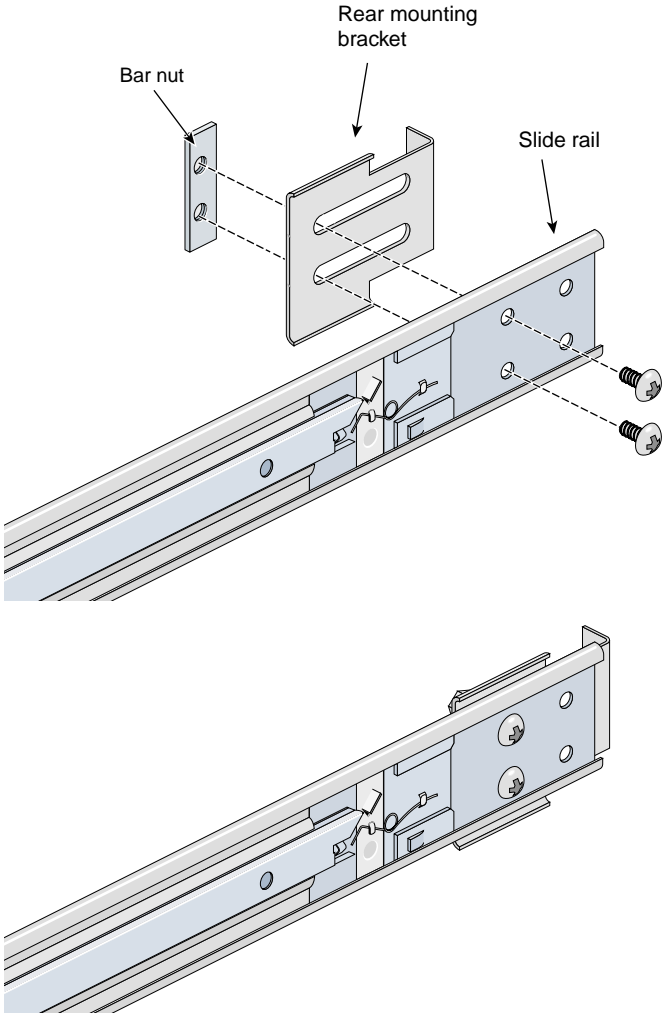
Figure B-2   Attaching the Rear Mounting Bracket to the Slide Rail

Preparing the Module

To attach the chassis rails to the module, follow these steps:

1. Place the module on a flat, stable surface.

2. Using four 10-24 x 1/4-inch screws, attach one of the chassis rails to the right side of the module chassis. Ensure that the rail is installed in the correct direction (see Figure B-3).

⚠️ **Caution:** Use only the 1/4-inch (0.64 cm) length screws. Longer screws damage internal components in the module.

3. Using five 10-24 x 1/4-inch screws, attach the second rail to the left side of the module chassis. Again, ensure that the rail is installed in the correct direction.
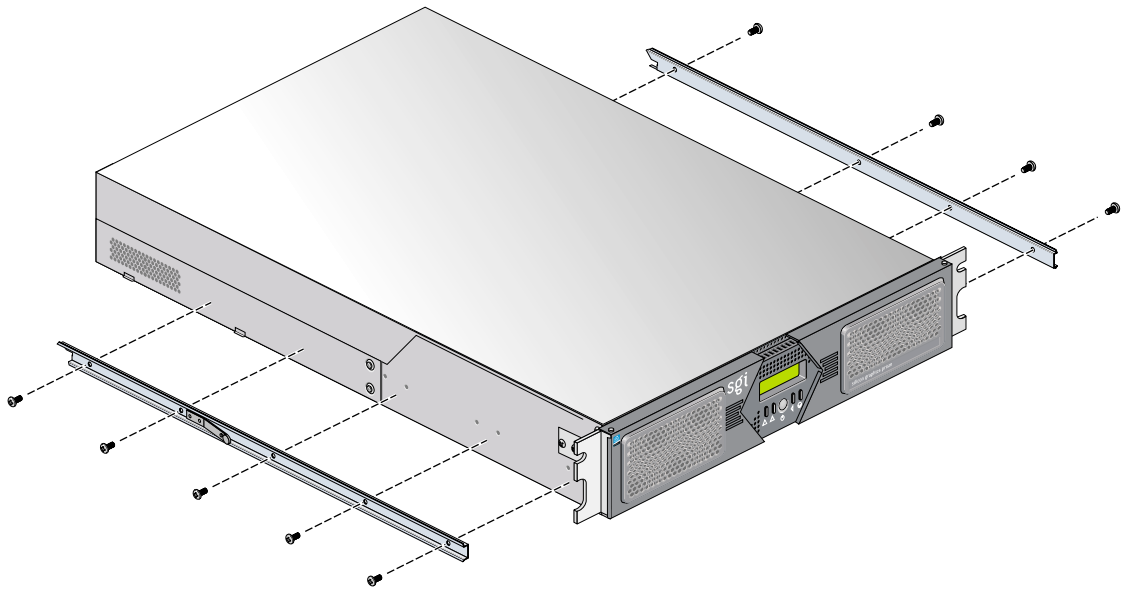
Figure B-3   Attaching Optional Chassis Rails to the Module Chassis

Determining Where to Attach the Slide Rail in the Rack

The module requires two units (2U) of space within the rack (one unit is equivalent to 1.75 inches [44.5 cm]). To determine where you should install the slide rails in the rack, you must count mounting holes. Each U contains three mounting holes; therefore, in the 2U of space that the module occupies, there are six mounting holes. The bottom hole of the 2U space is hole 1. The top mounting hole in the 2U space is hole 6. See Figure B-4.

Note: A module in the rack is identified by the lowest U number that it occupies. For example, in Figure B-4 the module resides in U5 (the fifth unit within the rack).
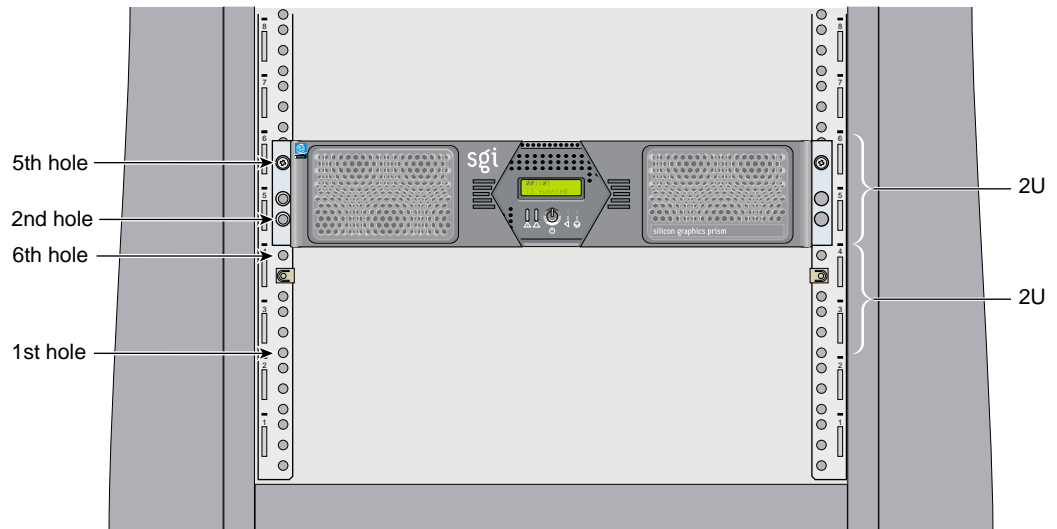
Figure B-4   Mounting-hole Pattern of Rack Vertical Rails

To determine how many mounting holes you must count, use the following formula:

3 x (the lowest U number that the module will occupy) - 2.

For example, when you want to install the module in locations U9 and U10, count 25 mounting holes (3 x 9 - 2) starting from the bottom of the rack. The 25th hole is the first mounting hole of U9.

Attaching the Slide Rail to the Rack

To attach the slide rail to the rack, follow these steps:

Tip: The slide rails must be level in the rack. To ensure that you install the slide rails correctly, carefully count the mounting holes on all of the rack rails (the two front rails and the two rear rails).

1. Locate eight 10-32 x 1/2-inch Phillips screws, eight shoulder washers, and four bar nuts.

2. Place one of the bar nuts inside the rack and align it with the second and third holes of the selected 2U of space (see Figure B-5).

    Note: The holes in the bar nuts are not centered. The bar nuts need to be placed in such a way that the holes are closest to the inside edge of the rack rails. See Figure B-5.

3. Insert two screws with shoulder washers through the rack rail to hold the bar nut in place. The screws should not be tightened at this point.

4. Repeat steps 2 and 3 to install the remaining three bar nuts on the other three rack rails (front and rear of rack).
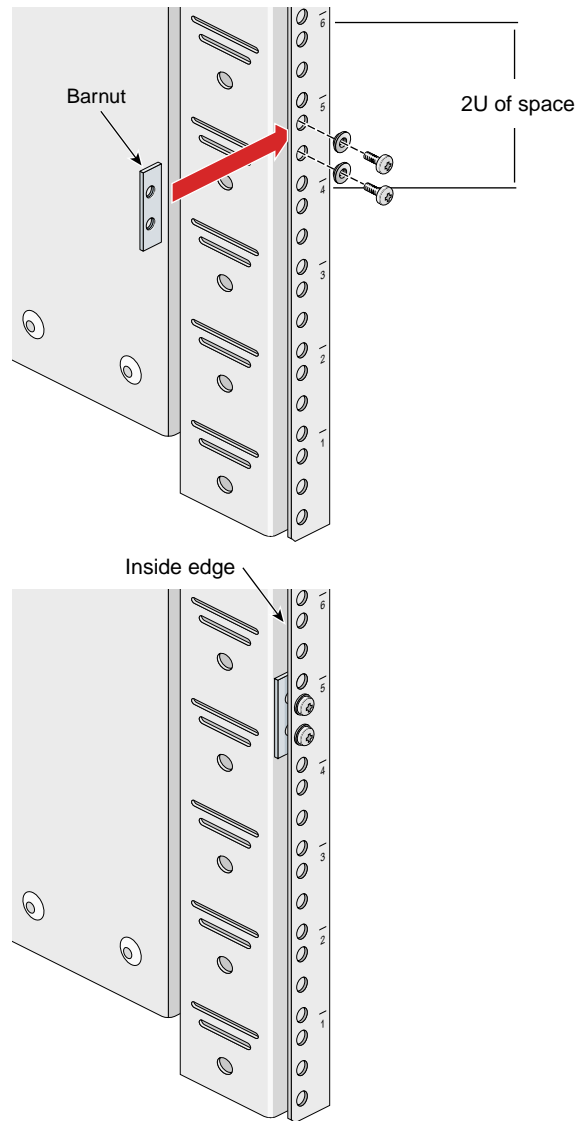
Figure B-5   Placing the Bar Nuts on the Rack Rails

5. Insert the front and rear brackets of one of the slide rails between the rack rails and the bar nuts, as shown in Figure B-6.

6. Tighten the screws on the front- and rear-end of the rails. Do not tighten firmly at this point, because all screws will be firmly tightened once the module is installed in the rack.

7. Repeat steps 5 and 6 to attach the second slide rail to the other side of the rack.
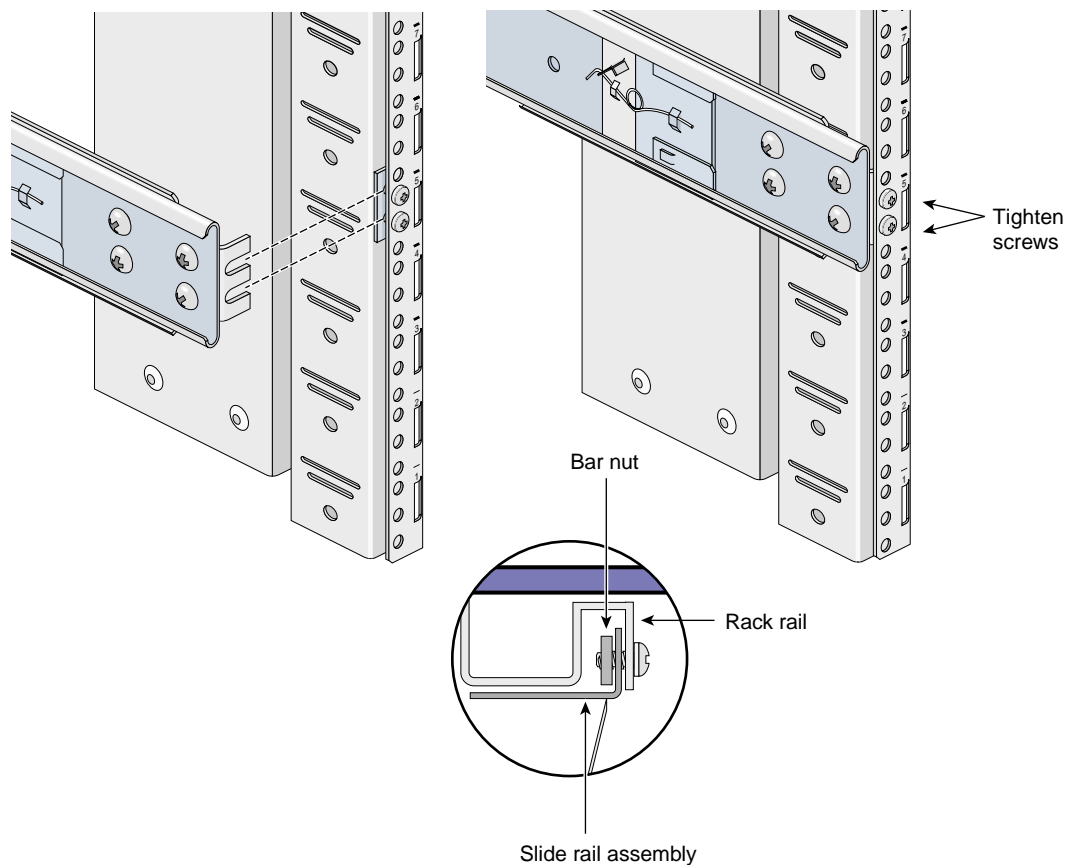
Figure B-6   Attaching the Slide Rail to the Rack

Installing Clip Nuts in Rack Rails

Clip nuts secure the modules to the rack. To install the clip nuts, slide the clip nuts over the fifth hole of the selected 2U of space on each of the front rails. See Figure B-7 for details.
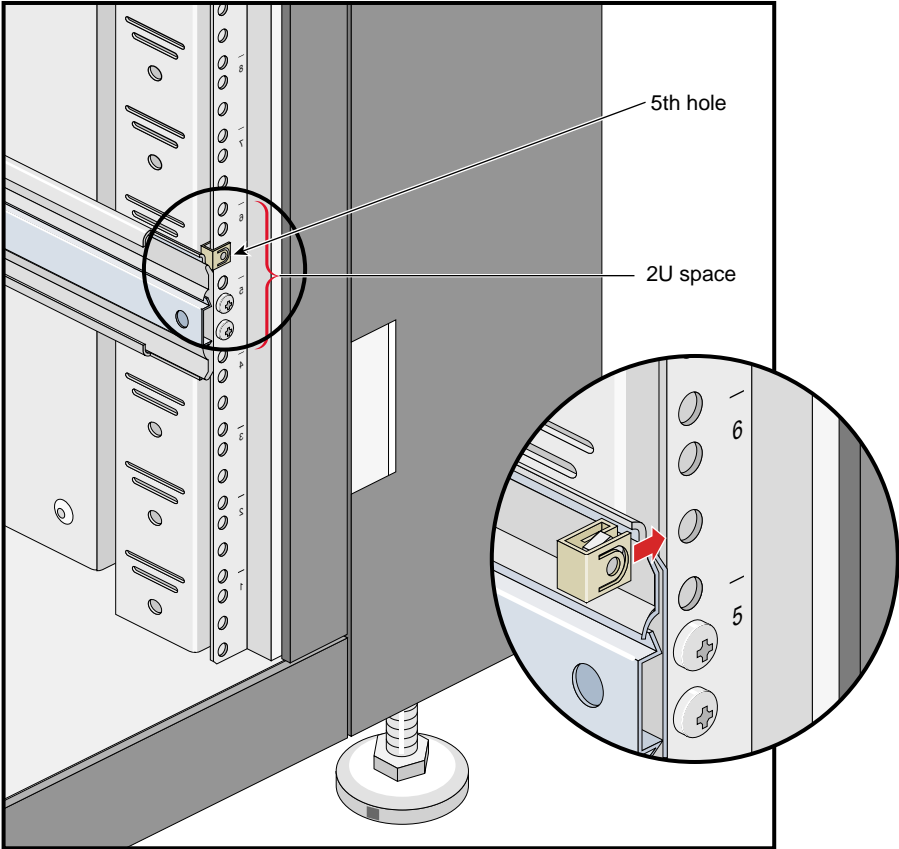
Figure B-7   Installing Clip Nuts in Rack Rails

Installing the Module in the Rack

To install the module in the rack, follow these steps:

Note: Step 2 requires two people.

---

1. Fully extend the left and right slide rails from the rack until they lock into place.

2. With one person holding each side of the module, align the chassis rails of the module with the slide rails of the rack.

3. Slide the chassis rails into the slide rails until the chassis rails are stopped by the safety latches.

4. Press the safety latches on both sides of the module to fully seat the chassis rails into the slide rails (see Figure B-8).

5. Firmly tighten all screws (the eight screws that secure the slide rails to the rack rails).
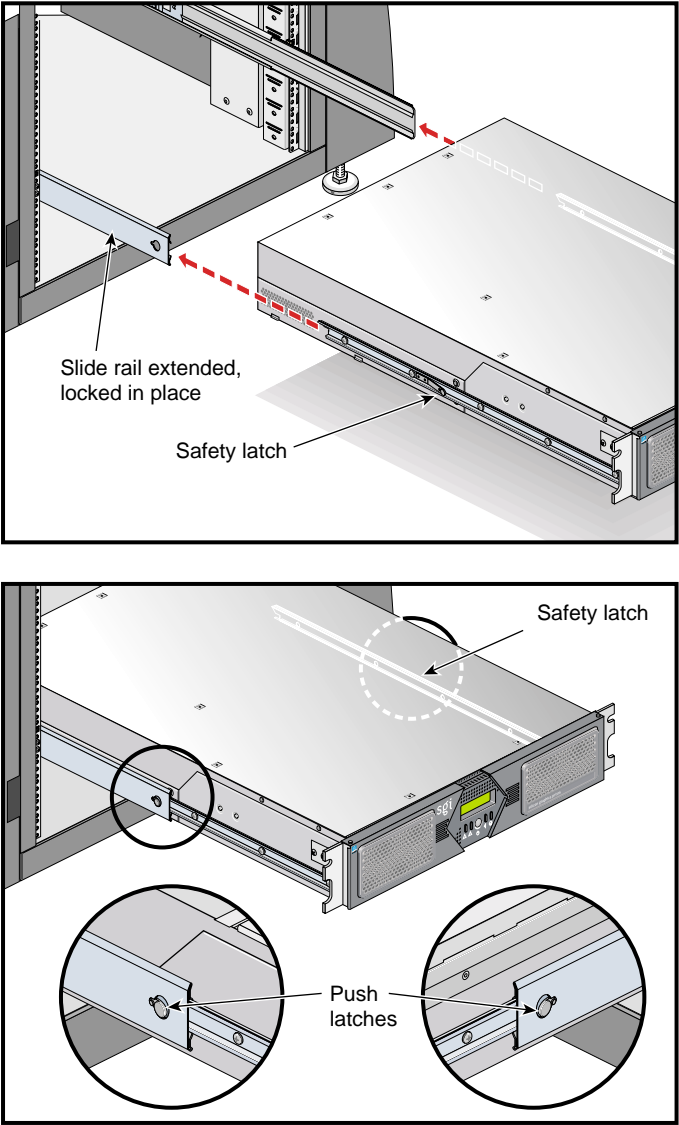
Figure B-8   Pressing the Safety Latches

6. Secure the module to the rack by inserting a 10-32 x 1/2-inch Phillips screw in the top hole of each chassis ear (see Figure B-9).
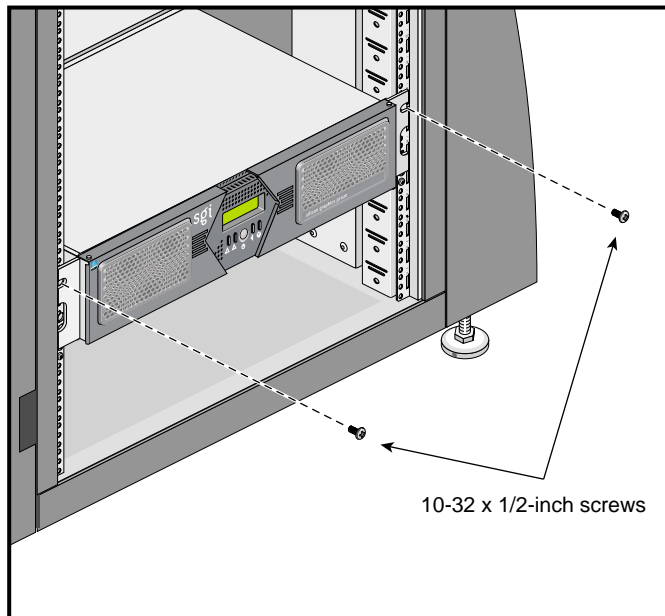
Figure B-9   Securing the Module to the Rack

Adjusting the Position of the Rackmounted Module

Once the module is installed in the rack, you can adjust the position of the module in the rack (up and down, side-to-side). To adjust the position of the module, loosen the front mounting screws, adjust the module to the desired position, then tighten the mounting screws.

⚠️ **Caution:** Do not lift the module by its bezel; it is not designed to handle the weight of the module. Instead, use the chassis ears to move the module (see Figure B-9).

007-4701-001

Removing a Rail-mounted Module from a Rack

To remove a module from a rack, follow these steps:

1. Power off the module.

2. Disconnect all of the cables at the rear of the module.

Warning: **Components may be hot. To avoid injury, allow the components to cool for approximately five minutes before you proceed with these instructions.**

3. Remove the two screws that secure the module to the front rails of the rack.

4. Carefully pull the module from the rack until it is stopped by the safety latches.

5. With one person holding each side, release the safety latches on both sides of the module and pull the module out of the slide rail (see Figure B-10).

6. Place the module on a flat, stable surface.
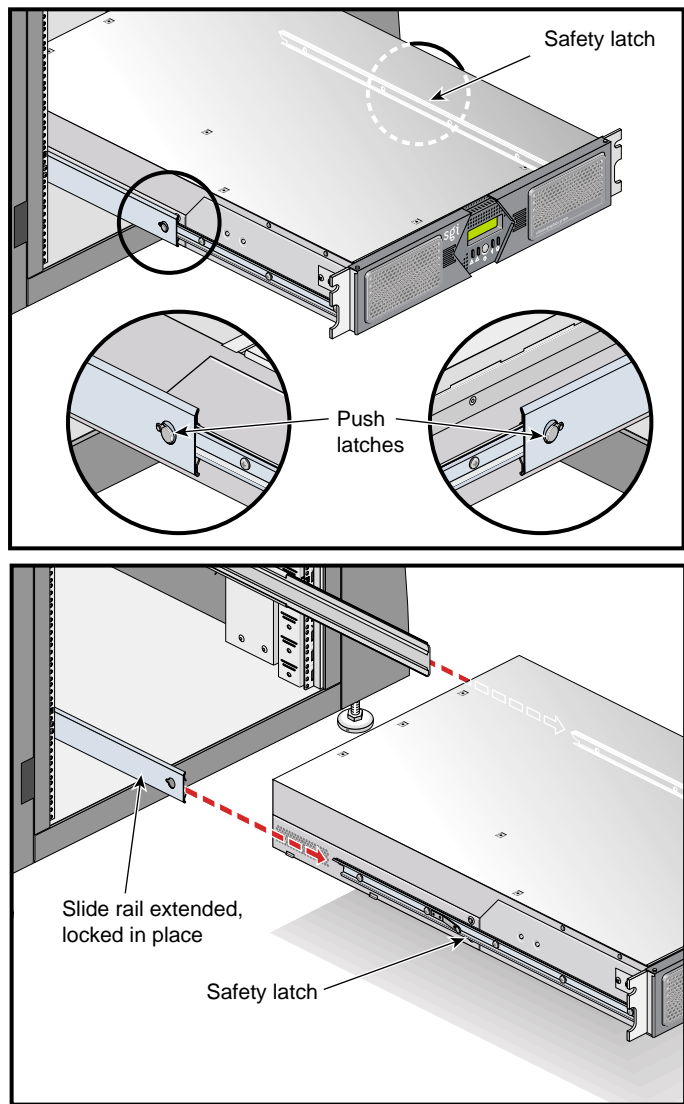
Figure B-10   Releasing the Safety Latches

7. To slide the slide rails back into the rack, push down on the slide latches as shown in Figure B-11.

Note: Before you reinstall a module into the rack, fully extend the slide rails from the rack until they lock into place.
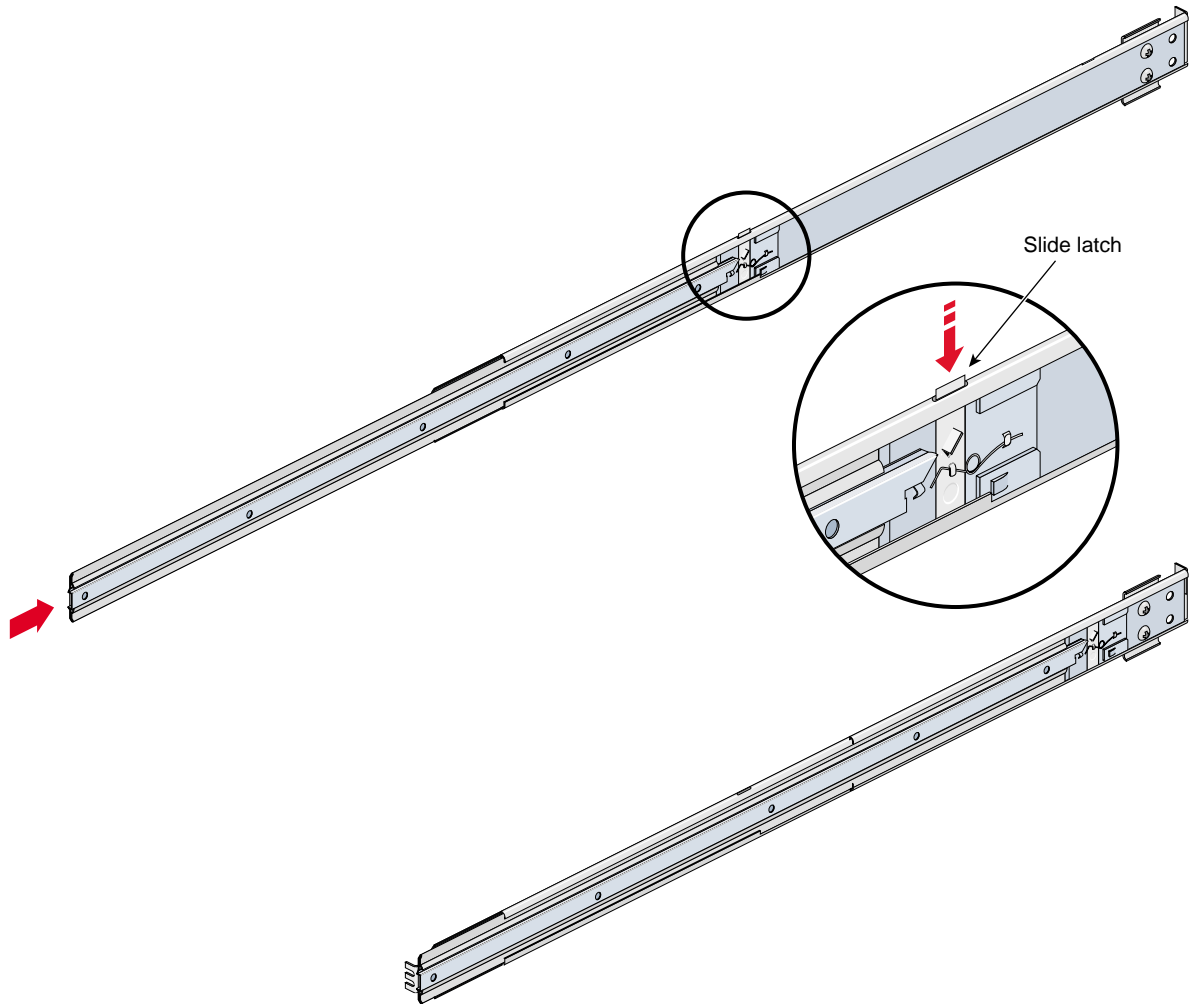
Slide latch

Figure B-11   Releasing the Slide Latches

Regulatory Specifications and Safety Information

This appendix presents regulatory information that may be important to the operation of your Silicon Graphics Prism.

Manufacturer's Regulatory Declarations

The Silicon Graphics Prism products conform to several national and international specifications and European Directives listed on the "Manufacturer's Declaration of Conformity." The CE insignia displayed on each device is an indication of conformity to the European requirements.

> ⚠️ **Caution:** Each SGI system has several governmental and third-party approvals, licenses, and permits. Do not modify this product in any way that is not expressly approved by SGI. If you do, you may lose these approvals and your governmental agency authority to operate this device.

System Model Number

The CMN (model) number for each system is printed on the system label on the unit.

CE Notice and Manufacturer's Declaration of Conformity

The "CE" symbol indicates compliance of the device to directives of the European Community. A "Declaration of Conformity" in accordance with the standards has been made and is available from SGI upon request.

Electromagnetic Emissions

This section provides the contents of electromagnetic emissions notices for various countries.

FCC Notice (USA Only)

This equipment complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions:

• This device may not cause harmful interference.

• This device must accept any interference received, including interference that may cause undesired operation.

Note: This equipment has been tested and found to comply 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 you will be required to correct the interference at your own expense.

If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, you are encouraged to try to correct the interference by using one or more of the following methods:

• Reorient or relocate the receiving antenna.

• Increase the separation between the equipment and receiver.

• Connect the equipment to an outlet on a circuit different from that to which the receiver is connected.

Consult the dealer or an experienced radio/TV technician for help.

Caution: Changes or modifications to the equipment not expressly approved by the party responsible for compliance could void your authority to operate the equipment.

Industry Canada Notice (Canada Only)

> This Class A digital apparatus meets all requirements of the Canadian Interference-Causing Equipment Regulations.
>
> Cet appareil numérique német pas de perturbations radioélectriques dépassant les normes applicables aux appareils numériques de Classe A préscrites dans le Règlement sur les interferences radioélectriques établi par le Ministère des Communications du Canada.

VCCI Notice (Japan Only)

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

Chinese Class A Regulatory Notice

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

Korean Class A Regulatory Notice

> 이 기기는 업무용으로 전자파적합등록을 한 기기이오니 판매자 또는 사용자는 이 점을 주의하시기 바라며 만약 잘못 판매 또는 구입하였을 때에는 가정용으로 교환하시기 바랍니다.

Shielded Cables

The Silicon Graphics VSL product is FCC compliant under test conditions that include the use of shielded cables between the system and its peripherals. Your system and any peripherals that you purchase from SGI have shielded cables. Shielded cables reduce the possibility of interference with radio, television, and other devices. If you use any cables that are not from SGI, ensure that they are shielded. Telephone cables do not require shielding.

Optional monitor cables supplied with your system use additional filtering molded into the cable jacket to reduce radio frequency interference. Always use the cable that is supplied with your system. If your monitor cable becomes damaged, obtain a replacement cable from SGI.

Electrostatic Discharge

SGI designs and tests its products to be resistant to the effects of electrostatic discharge (ESD). ESD is a source of electromagnetic interference and can cause problems ranging from data errors and lockups to permanent component damage.

It is important that you keep all the covers and doors, including the plastics, in place while you are operating the system. The shielded cables that came with the system and its peripherals should be installed correctly, with all thumbscrews fastened securely.

An ESD wrist strap may be included with some products, such as memory or PCI upgrades. Use the wrist strap when you install these upgrades to prevent the flow of static electricity; it is designed to protect your system from ESD damage.

Laser Compliance Statements

The DVD-ROM drive in this computer is a Class 1 laser product. The DVD-ROM drive-classification label is located on the drive.

Warning: **Invisible laser radiation when open. Avoid exposure to beam.**

⚡ **Warning:** Attention: Radiation du faisceau laser invisible en cas d'ouverture. Evitter toute exposition aux rayons.

⚡ **Warning:** Vorsicht: Unsichtbare Laserstrahlung, Wenn Abdeckung geöffnet, nicht dem Strahl aussetzen.

⚡ **Warning:** Advertencia: Radiación láser invisible al ser abierto. Evite exponerse a los rayos.

⚡ **Warning:** Advarsel: Laserstråling vedåbning se ikke ind i strålen

⚡ **Warning:** Varo! Lavattaessa Olet Alttina Lasersåteilylle

⚡ **Warning:** Varning: Laserstrålning når denna del år öppnad ålå tuijota såteeseenstirra ej in i strålen.

⚡ **Warning:** Varning: Laserstrålning nar denna del år öppnadstirra ej in i strålen.

⚡ **Warning:** Advarsel: Laserstråling nar deksel åpnesstirr ikke inn i strålen.

Lithium Battery Statement

Warning: Only qualified service personnel should replace the soldered lithium battery (or batteries) in the SGI system.

Warning: Advarsel!: Lithiumbatteri - Eksplosionsfare ved fejlagtig håndtering. Udskiftning må kun ske med batteri af samme fabrikat og type. Léver det brugte batteri tilbage til leverandøren.

Warning: Advarsel: Eksplosjonsfare ved feilaktig skifte av batteri. Benytt samme batteritype eller en tilsvarende type anbefalt av apparatfabrikanten. Brukte batterier kasseres i henhold til fabrikantens instruksjoner.

Warning: Varning: Explosionsfara vid felaktigt batteribyte. Använd samma batterityp eller en ekvivalent typ som rekommenderas av apparattillverkaren. Kassera använt batteri enligt fabrikantens instruktion.

Warning: Varoitus: Paristo voi räjähtää, jos se on virheellisesti asennettu. Vaihda paristo ainoastaan laitevalmistajan suosittelemaan tyyppiin. Hävitä käytetty paristo valmistajan ohjeiden mukaisesti.

Warning: Vorsicht!: Explosionsgefahr bei unsachgemäßen Austausch der Batterie. Ersatz nur durch denselben oder einen vom Hersteller empfohlenem ähnlichen Typ. Entsorgung gebrauchter Batterien nach Angaben des Herstellers.

Index