

Silicon Graphics® Prism™ Visualization System User's Guide

007-4701-002

CONTRIBUTORS

Written by Mark Schwenden and Eric Zamost

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.

COPYRIGHT

© 2005, Silicon Graphics, Inc. All rights reserved; provided portions may be copyright in third parties, as indicated elsewhere herein. No permission is granted to copy, distribute, or create derivative works from the contents of this electronic documentation in any manner, in whole or in part, without the prior written permission of Silicon Graphics, Inc.

LIMITED RIGHTS LEGEND

The electronic (software) version of this document was developed at private expense; if acquired under an agreement with the USA government or any contractor thereto, it is acquired as "commercial computer software" subject to the provisions of its applicable license agreement, as specified in (a) 48 CFR 12.212 of the FAR; or, if acquired for Department of Defense units, (b) 48 CFR 227-7202 of the DoD FAR Supplement; or sections succeeding thereto. Contractor/manufacturer is Silicon Graphics, Inc., 1500 Crittenden Lane, Mountain View, CA 94043.

TRADEMARKS AND ATTRIBUTIONS

Silicon Graphics, SGI, the SGI logo are registered trademarks and NUMALink and Silicon Graphics Prism are trademarks of Silicon Graphics, Inc., in the United States and/or other countries worldwide.

Intel and Itanium are trademarks or registered trademarks of Intel Corporation and its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds.

All other trademarks mentioned herein are the property of their respective owners.

Record of Revision

Version	Description
001	October 2004 First Release
002	February 2005 Added Prism Team configurations (routed NUMALink) Added Multi-Xserver configuration instructions Added additional rack-handling instructions Added description of optional IO9 card

Contents

Figures	xi
Tables	xv
Important Information	xviii
Chapter Descriptions	xix
Related Publications	xx
Conventions	xxi
Product Support	xxi
Reader Comments	xxii
1. Getting Started with the Silicon Graphics Prism.	1
XG2N and Compute Modules	1
Front Panel Controls	3
Compute Module Rear Panel Items.	5
Internal Components.	6
Safety Precautions	6
Hazard Statements	7
ESD Precautions	7
Safety Measures	8
Quick Start Information	9
Positioning and Power for Your Silicon Graphics Prism	9
NUMAlink Cabling for a Silicon Graphics Prism.	12
Optional SGI ImageSync Cabling	13
Connecting a Monitor	15
Connecting a Keyboard and Mouse	16
Connecting the Multi-Port Serial Cable	17

	Connecting an L1 System Console 18
	Power-On the System 19
	Verifying System Connections 20
	Powering Off Manually 22
2.	System Overview and Options 23
	Physical Features 24
	Functional Architecture 26
	System Components 29
	Base Compute Module 29
	XG2N Graphics Module 31
	CPU Expansion Module 32
	CMPX Module 34
	Router Module 36
	Router Module Front Panel Components 38
	Router Module Rear Panel Components 39
	Storage Expansion 40
	SGI TP900 Storage Module 41
	2Gb SGI TP9100 Optional Storage System 42
	Power Components 44
	SGI Racks for Silicon Graphics Prism 45
	Other Optional Devices 48
	Optional IO9 PCI Card 49
	Optional Tape Devices 49
	Optional DM8 Audio Board 50
	Board Installation 50
	Cables and Connectors 51
	Functional Features 51
	Troubleshooting Tips 52
	Optional L2 Controller 53
	Optional Keyboard Extender 54
3.	Configurations and Cabling 57
	System Configurations 58

NUMALink Cabling in Ring Topologies	60
NUMALink Cabling in Routed Systems.	63
Minimum Routed Configuration	64
Maximum Routed Configuration	66
4. Installing and Removing Customer-Replaceable Units.	69
Safety Instructions.	70
Accessing the System’s Internal Components	72
PCI and PCI-X Cards	74
Rules for Card Installation and Removal	75
Installing a PCI or PCI-X Card	76
Removing a PCI or PCI-X Card	84
Replacing an XG2N Graphics Card	93
Disk Drives.	100
Installing a Disk Drive	101
Removing a Disk Drive	103
Removing and Replacing Power Supplies	105
Memory DIMM Group Placement and Guidelines	109
Installing a DIMM Group	111
Removing a DIMM	114
L1 Controller Display	117
5. Troubleshooting	123
Troubleshooting Chart	124
L1 Controller Error Messages	126
Reading Power Supply Status LEDs.	129
SGI Electronic Support	130
Customizing the XF86Config File	133
Configuring a System for Stereo	133
Example “Device” Section for Stereo	134
Sample Stereo Mode Entries	134
Example “Monitor” Section for Stereo.	135
Example “Screen” Section for Stereo	135

Configuring a System for Full Scene Anti-Aliasing	136
Example “Device” Section for Full Scene Anti-Aliasing	137
Configuring a System for Dual-Channel Operation	138
Example “Device” Section for Dual Channel	139
Configuring a System for SwapReady	140
Configuring a System for External Genlock or Framelock	141
Configuring Monitor Positions	143
Example “ServerLayout” Section for Four Monitors in a Line	143
Example “ServerLayout” Section for Four Monitors in a Square	144
Configuring Monitor Types	145
Example “Device” Section for Use With Two Analog Monitors.	145
Configuring a System for Multiple Xservers	146
Identifying Event Devices	147
Identifying Event Devices	147
Creating a Multi-Seat XF86Config File	148
Example “ServerLayout” Sections for Three Xservers	152
Pointing X to the New XF86Config-Nserver File	153
Example /etc/X11/xdm/gdm.conf Servers Section for Three Xservers	154
A. Technical Specifications	155
Environmental Specifications	156
Compute Module Specifications	157
CMPX Module Specifications	159
Router Module Technical Specifications	160
Rack Specifications	161
SGI TP900 Storage Module Specifications	162
Non-proprietary I/O Port Specifications	163
Compute Module	163
XG2N Module Connectors	164
CMPX Module	166
L2 Controller	167
SGI TP900 Storage Module	168
DB-9 Serial Connector.	169
RJ-45 Connector	171

External Multi-port Serial Adapter Connector172
USB Type A Connector174
USB Type B Connector175
B. Installing Silicon Graphics Prism Modules Into a Rack177
Choosing Locations Within a Rack177
Programming L1 Rack and Slot Numbers178
Using Silicon Graphics Prism Modules with a Rack179
Rackmounting with Optional Slide Rails179
Determining Module Space Requirements.179
Slide Rail Mounting Hardware180
Preparing the Optional Slide Rail Assemblies.181
Preparing the Module184
Determining Where to Attach the Slide Rail in the Rack185
Attaching the Slide Rail to the Rack186
Installing Clip Nuts in Rack Rails189
Installing the Module in the Rack190
Adjusting the Position of the Rackmounted Module192
Removing a Rail-Mounted Module from a Rack193
C. Installing Rack Systems197
Pre-Installation Activities.198
Site Plan Verification198
Tools Required198
Power Receptacle Verification199
Unloading and Moving System Equipment202
Unloading the Equipment from the Truck.202
Inspecting the Shipping Crate204
Transporting the Shipping Crate204
Removing a Short Rack from the Shipping Crate206
Removing a Tall Rack from the Shipping Crate208
Positioning and Leveling a Single-Rack System211
Positioning and Leveling Multiple Racks213

D. Regulatory Specifications and Safety Information	217
Manufacturer’s Regulatory Declarations	217
System Model Number	217
CE Notice and Manufacturer's Declaration of Conformity	217
Electromagnetic Emissions	218
FCC Notice (USA Only)	218
Industry Canada Notice (Canada Only)	219
VCCI Notice (Japan Only)	219
Chinese Class A Regulatory Notice.	219
Korean Class A Regulatory Notice	219
Shielded Cables	220
Electrostatic Discharge	220
Laser Compliance Statements.	220
Lithium Battery Statement	222
Index.	223

Figures

Figure 1-1	Rear View of XG2N Graphics Module (Showing First Module Pipe Numbering)	2
Figure 1-2	Pipe Numbering of Additional XG2N Graphics Modules	3
Figure 1-3	Silicon Graphics Prism XG2N and Compute Module Front Panel Controls	3
Figure 1-4	Silicon Graphics Prism Compute Module Rear Panel With ImageSync	5
Figure 1-5	Silicon Graphics Prism XG2N Module Internal View	6
Figure 1-6	Example Connection of a Non-Racked System to Power	10
Figure 1-7	Connection Example of a Racked System to Power	11
Figure 1-8	NUMAlink Cabling in a Two-module System	12
Figure 1-9	Cabling Example of SGI ImageSync Card to Graphics	14
Figure 1-10	Connecting the USB Keyboard and Mouse to the System Compute Module	16
Figure 1-11	Multi-port Adapter Serial Cable	17
Figure 1-12	Connecting a Serial Terminal to the Base Compute Console Port	18
Figure 2-1	Example of Silicon Graphics Prism Rack Systems	25
Figure 2-2	Front and Rear Views of Base Compute Module with IO10 PCI Card.	30
Figure 2-3	Front and Rear Views of XG2N Module	31
Figure 2-4	Front and Rear Views of CPU Expansion Module	33
Figure 2-5	Front and Rear Views of CMPX Option Module	35
Figure 2-6	Router Functional Block Diagram	37
Figure 2-7	Front View of the Router Module	38
Figure 2-8	Rear View of Router Module	39
Figure 2-9	SGI TP900 Storage Module	41
Figure 2-10	2Gb SGI TP9100 Storage System	43

Figure 2-11	Unit Numbering Within Racks 45
Figure 2-12	Optional Slide Rails 46
Figure 2-13	Shelf Rails 46
Figure 2-14	Front Views of Short and Tall Racks 47
Figure 2-15	DM8 Interface Panel and Connectors 50
Figure 2-16	Optional L2 Controller Unit 53
Figure 2-17	Keyboard and Mouse Connected via Optional USB Extender 55
Figure 3-1	Two-Module Ring Topology Example 61
Figure 3-2	Three-Module Ring Topology Example 61
Figure 3-3	Four-Module Ring Topology Example 62
Figure 3-4	Minimum Routed Configuration Conceptual Diagram. 64
Figure 3-5	Module Ordering in Minimum Routed Configuration 65
Figure 3-6	Maximum Routed Configuration Conceptual Diagram 67
Figure 3-7	Module Ordering in Maximum Routed Configuration 68
Figure 4-1	Opening the System Cover. 72
Figure 4-2	Closing the System Cover 73
Figure 4-3	PCI and PCI-X Card Slots in the Base Compute Module 74
Figure 4-4	Removing an Optional Chassis Rail 77
Figure 4-5	Removing the Lower PCI/PCI-X Support Bracket 78
Figure 4-6	Removing the Blanking Plate 79
Figure 4-7	Installing the Card and Retaining Screw 80
Figure 4-8	Replacing the Lower PCI/PCI-X Support Bracket 81
Figure 4-9	Replacing the Chassis Rail 82
Figure 4-10	Opening the Cover to Remove the Card 85
Figure 4-11	Removing the Chassis Rail. 86
Figure 4-12	Removing the Lower PCI/PCI-X Support Bracket 87
Figure 4-13	Removing the Retaining Screw and Extracting the Card 88
Figure 4-14	Replacing a Blanking Plate. 89
Figure 4-15	Replacing the Lower PCI/PCI-X Support Bracket 90
Figure 4-16	Replacing the Optional Chassis Rail 91
Figure 4-17	Close and Screw Down Cover 92
Figure 4-18	Removing the XG2N Graphics Card Bracket 94
Figure 4-19	Remove XG2N Graphics Board Assembly 95

Figure 4-20	Synchronization Cable Connection on Underside of Graphics Board	96
Figure 4-21	Installing the Graphics Card and Retaining Screws	97
Figure 4-22	Replacing the PCI/PCI-X Support Bracket	98
Figure 4-23	Disk Drive Locations	100
Figure 4-24	Installing a Disk Drive	102
Figure 4-25	Removing a Disk Drive	104
Figure 4-26	Power Supplies Location	105
Figure 4-27	Removing a Power Supply	107
Figure 4-28	Installing the Power Supply	108
Figure 4-29	Layout of Slots and DIMM Memory Groups	110
Figure 4-30	Opening Module Cover to Install DIMMs	112
Figure 4-31	Inserting a DIMM	113
Figure 4-32	Opening the Module to Remove a DIMM	115
Figure 4-33	Removing a DIMM.	116
Figure 4-34	L1 Controller Display on Base Compute Module	117
Figure 4-35	Opening the Module to Access the L1 Controller Display.	118
Figure 4-36	Removing the L1 Controller Display Panel	119
Figure 4-37	Installing an L1 Controller Display Panel	121
Figure 5-1	Full Support Sequence	130
Figure 5-2	Four Monitors in a Line	143
Figure 5-3	Four Monitors in a Square.	144
Figure A-1	Rear Panel of Compute Module	164
Figure A-2	DVI-I Connector Pinout	164
Figure A-3	Non-Proprietary Connectors on Rear Panel of CMPX Module	166
Figure A-4	Non-Proprietary Connectors on Rear of Optional L2 Controller	167
Figure A-5	Non-proprietary Connectors on Rear Panel of TP900 Module	168
Figure A-6	DB-9 Serial Connector Pin Assignments	169
Figure A-7	RJ-45 Connector Pin Assignments	171
Figure A-8	Pin Number Locations for 36-pin MDR Connector.	172
Figure A-9	Pin Number Locations for USB Type A Connector.	174
Figure A-10	Pin Number Locations for USB Type B Connector	175
Figure B-1	Removing the Chassis Rail from the Slide Rail	181
Figure B-2	Attaching the Rear Mounting Bracket to the Slide Rail	183

Figure B-3	Attaching Optional Chassis Rails to the Module Chassis	184
Figure B-4	Mounting-hole Pattern of Rack Vertical Rails	185
Figure B-5	Placing the Bar Nuts on the Rack Rails	187
Figure B-6	Attaching the Slide Rail to the Rack	188
Figure B-7	Installing Clip Nuts in Rack Rails	189
Figure B-8	Pressing the Safety Latches.	191
Figure B-9	Securing the Module to the Rack	192
Figure B-10	Releasing the Safety Latches	194
Figure B-11	Releasing the Slide Latches.	195
Figure C-1	30-amp Single-phase Power Receptacle for North American Sites	200
Figure C-2	32-amp Single-phase Power Plug for International Sites	201
Figure C-3	Dimensions of Tall Rack Shipping Crate	203
Figure C-4	Dimensions of Short Rack Shipping Crate	204
Figure C-5	Removing a Short Rack from the Shipping Crate	207
Figure C-6	Removing a Tall Rack from the Shipping Crate	210
Figure C-7	Leveling Bolts	211
Figure C-8	Seismic Tie-down Attachment Points.	212
Figure C-9	Leveling Bolts	213
Figure C-10	Joining Locations	215

Tables

Table 3-1	Silicon Graphics Prism Power System Configuration Ranges . . .	58
Table 3-2	Silicon Graphics Prism Team System Configuration Ranges . . .	59
Table 3-3	Bisection Bandwidth of Ring Topology	60
Table 3-4	Minimum Routed Configuration Cable Chart	64
Table 3-5	Maximum Routed Configuration Cable Chart	66
Table 5-1	Troubleshooting Chart.	124
Table 5-2	L1 Controller Messages	127
Table 5-3	LED Status and Power Supply Condition	129
Table 5-4	Input Video Formats (Framelock)	142
Table 5-5	Options for Monitor Layout	145
Table A-1	Environmental Specifications	156
Table A-2	Bandwidth Characteristics of the Compute Module	157
Table A-3	General Features of the Compute Module	157
Table A-4	Compute Module Physical Specifications	158
Table A-5	Bandwidth Characteristics of the CMPX Module	159
Table A-6	CMPX Module Specifications	159
Table A-7	Router Port Specifications	160
Table A-8	Router Technical Specifications	160
Table A-9	Short Rack Specifications (with Skins)	161
Table A-10	Tall Rack Specifications	161
Table A-11	TP900 Storage Module Specifications	162
Table A-12	Compute Module Connectors.	163
Table A-13	DVI-I Connector Signals and Pins	165
Table A-14	CMPX Module Connectors	166
Table A-15	L2 Controller Connectors	167
Table A-16	DB-9 Serial Connector Pin Assignments.	170

Table A-17	Ethernet Connector Pin Assignments	171
Table A-18	Multi-port Serial Adapter Pinouts.	172
Table A-19	Pin Assignments for USB Type A Connector	174
Table A-20	Pin Assignments for USB Type B Connector.	175
Table B-1	Rackmounting Space Requirements for Modules	179
Table B-2	Slide Rail Rackmounting Hardware	180
Table C-1	Installation Tools	198

About This Guide

This guide provides an overview of the Silicon Graphics Prism Visualization System components and 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 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.
- A Router module provides eight NUMAlink-4 ports and functions as a high-speed switch to route network packets between base compute, XG2N, CPU expansion, and CMPX modules.
- 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* (007-4428-00x). Note that 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, “Configurations and Cabling,”
- Chapter 4, “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 5, “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 Into a Rack,” provides instructions for rack mounting modules with optional slide rails.
- Appendix C, “Installing Rack Systems,” provides instructions for installing a pre-assembled rack system.
- Appendix D, “Regulatory Specifications and Safety Information,” contains regulatory specifications and safety information for 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 optional 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 vX.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.
<i>variable</i>	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(<i>x</i>)	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:

- 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 more than one XG2N graphics module the pipes are numbered as shown in Figure 1-2.




Figure 1-1 Rear View of XG2N Graphics Module (Showing First Module Pipe Numbering)




Figure 1-2 Pipe Numbering of Additional XG2N Graphics Modules

Front Panel Controls

This section describes the front panel controls and indicators of the XG2N and compute modules as shown in Figure 1-3.




Figure 1-3 Silicon Graphics Prism XG2N and Compute 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 green LED illuminates when the internal components are on.
 - **Service-required LED.** This amber LED illuminates to indicate that an item is not functioning properly (for example, a fan is off), but the system is still operating.
 - **Failure LED.** This red LED illuminates 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 are used to connect base compute, XG2N, CPU expansion, or CMPX modules to each other or to optional router modules. These connections are made with a NUMALink cable at 3.2 GB/s in each direction.
- **NUMALink LEDs.** Each NUMALink connector has a yellow LED and a green LED (both located to the right of the NUMALink connector). These LEDs provide information about this specific NUMALink connection:
 - The yellow LED illuminates when this module and the module to which this port is connected are both powered on.

- The green LED illuminates when a link has been established to another module through this NUMAlink connector.

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 7
- “ESD Precautions” on page 7
- “Safety Measures” on page 8

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 10, or placed in a rack. For instructions on mounting modules in a rack see Appendix B, “Installing Silicon Graphics Prism Modules Into a Rack,”.

If your system was shipped pre-mounted in a rack, see Appendix C, “Installing Rack Systems,” for important information regarding unloading, unpacking, and installing a rack system.



Warning: Follow the guidelines in “Safety Measures” on page 8 and Appendix C, “Installing Rack Systems,” to avoid damage to equipment, injury, or death.

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




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 13 to continue.

This section describes the NUMAlink cabling for the basic configuration. For additional NUMAlink cabling information, see “System Configurations” on page 58.

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 15, 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.



Caution: Although the image sync subsystem uses DB9 connectors, these connectors, whether on the ImageSync card or on the XG2N modules, are not serial ports. Connecting a serial device to these connectors may cause damage to both the ImageSync devices and the serial devices.




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 XG2N graphics modules, 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 54 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 17 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.




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 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 5, "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 5, "Troubleshooting," for additional trouble shooting tips and L1 messaging.

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

1. Enter **/sbin/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 5, "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 24
- “Functional Architecture” on page 26
- “System Components” on page 29
- “Power Components” on page 44
- “SGI Racks for Silicon Graphics Prism” on page 45
- “Other Optional Devices” on page 48

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 16 processors, 192 GB local DIMM memory, and 28 PCI/PCI-X slots. 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 Super-Hub (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 192 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 16 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.

A Silicon Graphics Prism system contains a combination of the following modules:

- **Base compute module.** All Silicon Graphics Prism systems include at least one base compute module that contains the following components:
 - A node board with two Intel Itanium 2 processors (each processor has integrated L1, L2, and L3 caches), between 2 GB and 24 GB of local memory, and a SHUB ASIC (the crossbar between the processors, local memory, the network interface, and the I/O interface).
 - 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 a minimum of 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 (no more than one card per module).

The IO10 card has real-time interrupt input and output ports, 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.

- A bootable hard drive
- A DVD-ROM drive

- **XG2N graphics module.** All systems contain at least one graphics module that has the following components:
 - A node board with zero or two Intel Itanium 2 processors (each processor has integrated L1, L2, and L3 caches), between 0 GB and 24 GB of local memory, and a SHUB ASIC (the crossbar between the processors, local memory, the network interface, and the I/O interface).
 - Two graphics pipes
- **CPU Expansion module.** Like the base compute module, the CPU module contains a node board and local memory. This module is different from the base compute module, however, in the following ways:
 - It has zero, one, or two CPUs (the base compute module must have two CPUs)
 - It does not contain PCI/PCI-X slots, hard drives, or a removable media device.
- **CMPX module.** Like the base compute module, the CMPX module contains a node board and local memory. This module is different from the base compute module, however, in the following ways:
 - It has zero, one, or two CPUs (the base compute module must have two CPUs)
 - It does not contain an IO10 PCI card, hard drives, or a removable media device.
- **Router module.** ~Like the base compute module, the CMPX module contains a node board and local memory. This module is different from the base compute module, however, in the following ways:
 - It has zero, one, or two CPUs (the base compute module must have two CPUs)
 - It does not contain an IO10 PCI card, hard drives, or a removable media device.

When the system consists of a base compute module and XG2N graphics module only, it may contain a maximum of 4 processors and a maximum of 48 GB of memory. 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 or through router modules.

System Components

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

- “Base Compute Module” on page 29
- “XG2N Graphics Module” on page 31
- “CPU Expansion Module” on page 32
- “CMPX Module” on page 34
- “Router Module” on page 36

Base Compute Module

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

- Two Intel Itanium 2 processors (each with integrated L1, L2, and L3 caches)
- Between 2 GB and 24 GB of local DIMM memory (4, 8, or 12 DIMMs)
- 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 card.
- 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 2-2 shows the front and rear views of a base compute module.




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

XG2N Graphics Module

The XG2N graphics module (see Figure 2-3) is a 2U AC-powered device that consists of the following:

- Two graphics pipes
- Zero or two Intel Itanium 2 processors (each with integrated L1, L2, and L3 caches)
- Between 0 GB and 24 GB of 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




Figure 2-3 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, one, or two Intel Itanium 2 processors (each with integrated L1, L2, and L3 caches)
- Between 2 GB and 24 GB of 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

Figure 2-4 shows the front and rear views of a CPU expansion module.




Figure 2-4 Front and Rear Views of CPU Expansion Module

CMPX Module

The CMPX module (see Figure 2-5) is a 2U AC-powered module that offers:

- Zero, one, or two Intel Itanium 2 processors (each with integrated L1, L2, and L3 caches)
- Between 2 and 24 GB of local DIMM memory (4, 8, or 12 DIMMs)
- Four PCI/PCI-X slots
- 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 2-5 Front and Rear Views of CMPX Option Module

Router Module

The optional router module (sometimes called an R-brick) is an eight-port 2U-high module that functions as a high-speed switch to route network packets between base compute, XG2N, CPU expansion, and CMPX modules. This creates a NUMAlink-4 interconnect fabric (as opposed to the ring topology normally used in smaller system configurations). The optional router is generally only used when there are four and eight modules interconnected as a single system image.

The section “NUMAlink Cabling in Routed Systems” on page 63 details the cable connection points used in a system with a router.

The key component within the module is the router chip, an SGI custom-designed ASIC. The router chip is an eight-port crossbar that connects any input-link channel to any of the seven possible output-link channels (ports).

The router has the following features:

- Eight NUMAlink-4 channels
- One USB port for system controller support
- One L1 controller and LCD display
- One 9-pin console connector
- Two hot-pluggable cooling fans (not customer replaceable)

Figure 2-6 shows a block diagram of the router.

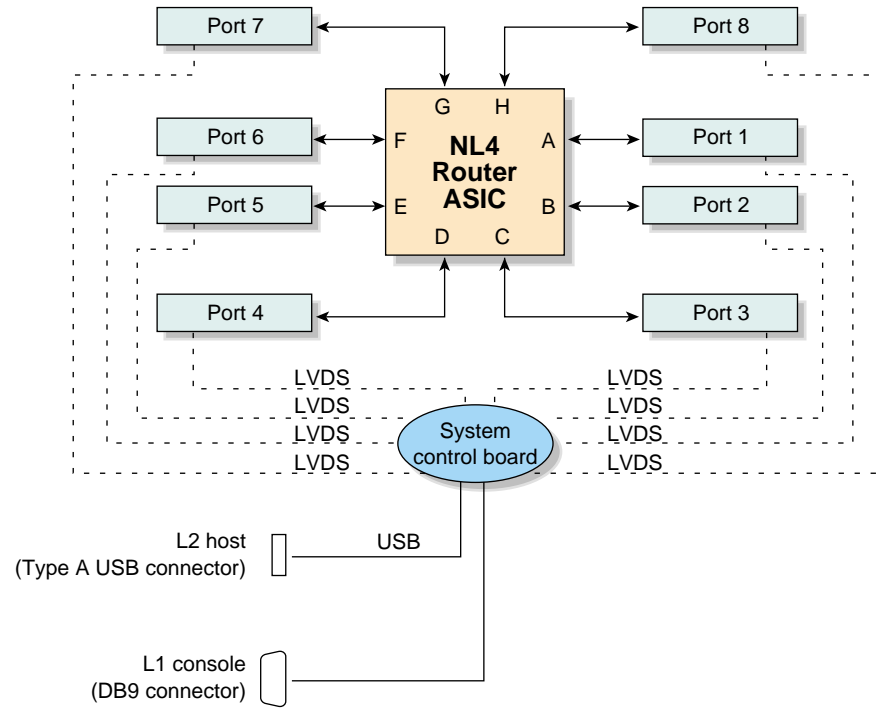


Figure 2-6 Router Functional Block Diagram

Router Module Front Panel Components

The router module contains the following front panel items (as shown in Figure 2-7):

- **L1 display.** The L1 display is a 55.7 mm X 32 mm backlit liquid crystal display (LCD) that displays system messages. It displays two lines with a maximum of 12 characters on each line.
- **On/Off switch with LED.** Press this button to turn on the router internal components. You can also turn on the router internal components at a system console.
- **Three LEDs:**
 - **Power-button LED.** This green LED illuminates when the router internal components are on and turns off when they are off.
 - **Service-required LED.** This amber LED illuminates to indicate that an item is broken or not operating properly (for example, a fan is off), but the router is still operating.
 - **Failure LED.** This red LED illuminates to indicate that a system failure has occurred and the router is down.
- **Fans.** Two hot-pluggable fans provide N+1 redundant cooling.




Figure 2-7 Front View of the Router Module

Router Module Rear Panel Components

The router module has the following rear panel items (see Figure 2-8):

- **PWR (power) connector.** This connects the router to the power outlet (120V or 220V, autosensing).
- **NUMAlink connectors.** These NUMAlink connectors connect the router to the Prism compute or graphics modules to form an interconnect fabric.
- **L1 port connector.** This connects the internal USB hub of the router to an optional console/controller. The internal USB hub can receive the USB signals from the controller via this port and distribute these signals to the L1 of the router.
- **Console connector.** This serial port provides optional connection to a terminal for system control purposes.
- **Link connector LEDs.** Each NUMAlink connector has two LEDs, as follows:
 - The yellow LED illuminates to indicate that both the router and the module to which this NUMAlink port is connected are powered on.
 - The green LED illuminates when a link has been established between the router and the module to which it is connected through this NUMAlink port.




Figure 2-8 Rear View of Router Module

Storage Expansion

A base compute module contains an IO10 base I/O 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. With the addition of an optional SCSI PCI card, the TP900 can be connected to base compute modules or CMPX expansion modules.
- 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 2-9, 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 (007-4428-00x)*.




Figure 2-9 SGI TP900 Storage Module

2Gb SGI TP9100 Optional Storage System

The 2Gb SGI TP9100, shown in Figure 2-10, is an affordable, entry-level RAID storage array that is easily expandable and comes in either a desktide 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* (007-4522-00x).




Figure 2-10 2Gb SGI TP9100 Storage System

Power Components

The Silicon Graphics Prism platform can contain 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, router modules, and TP900 storage modules.

See Figure 1-7 on page 11 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 rack units (U); one U is equal to 1.75 in. (4.45 cm). The short rack is a 17U rack (see Figure 2-11). The tall rack is a 39U rack (see Figure 2-14).

Each component within the rack is identified by the lowest U number that it occupies. For example, the top (XG2N) module shown in Figure 2-11 is described as being located in U12.




Figure 2-11 Unit Numbering Within Racks

Both rack types are industry-standard 19-inch racks, and they support two types of mounting rails (shelf rails or optional slide 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 2-12). The optional TP900 storage modules always use shelf rails, which are two parallel L-shaped mounting rails within the rack (see Figure 2-13).




Figure 2-12 Optional Slide Rails

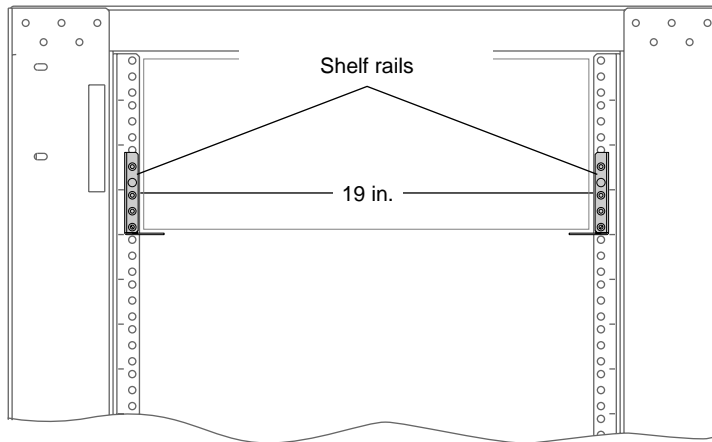


Figure 2-13 Shelf Rails

Both short and tall racks, as shown in Figure 2-14, 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.



Warning: Follow the guidelines in “Safety Measures” on page 8 and Appendix C, “Installing Rack Systems,” to avoid damage to equipment, injury, or death.

The base of each rack has seismic tie-down attachment points. The base of the tall rack also has leveling pads.




Figure 2-14 Front Views of Short and Tall Racks

Other Optional Devices

This section describes some of the optional devices available for the Silicon Graphics Prism system, in the following subsections:

- “Optional IO9 PCI Card” on page 49
- “Optional Tape Devices” on page 49
- “Optional DM8 Audio Board” on page 50
- “Optional L2 Controller” on page 53
- “Optional Keyboard Extender” on page 54

Optional IO9 PCI Card

If you do not wish to use the IO10 base I/O card in your system, an optional IO9 PCI card is required for base I/O functionality. Each system must contain at least one base I/O PCI card for base I/O functionality within the system. This PCI card must reside in bus 1, slot 1 (the bottom slot) of the lowest base compute module in the system. The IO9 PCI card has the following connectors:

- External VHDCI 68-pin SCSI connector
- 10/100/1000BaseT Ethernet connector
- Real-time interrupt output (RTO) and real-time interrupt input (RTI) connectors

Note: The RT interrupt input and RT interrupt output functionality of the IO9 PCI card is not supported under SGI Linux + ProPack.

The optional IO9 card also contains an IOC-4 ASIC that supports the following features:

- One IDE channel for the DVD-ROM
- NVRAM and time-of-day clock

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 2-15.

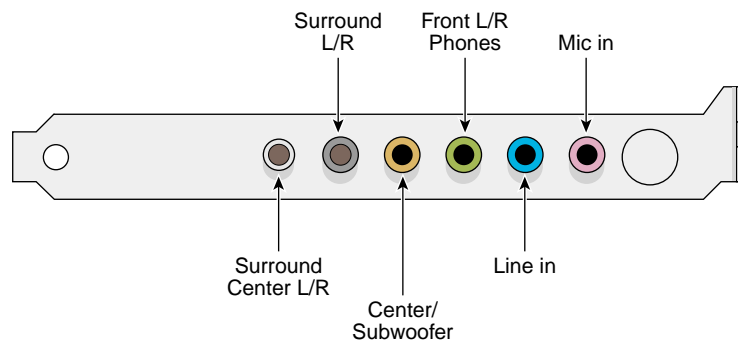


Figure 2-15 DM8 Interface Panel and Connectors

Board Installation

Your DM8 audio board is tested for operation in specific SGI visualization systems. Select SGI Linux based systems support use of a single DM8 audio board (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 board. 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 its rear panel as shown from left to right in Figure 2-15 on page 50:

- 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 board 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

Line in:

- Dynamic range of 105 dB (a-weighted)

Mic in:

- Dynamic range of 90 dB (a-weighted)
- Maximum input level of 1.45 Vrms (3.2 dBV)

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 board:

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 optional L2 controller (shown in Figure 2-16) 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. For example, base compute, XG2N, CPU expansion, CMPX, and router 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 2-16 Optional L2 Controller Unit

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 16), 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 2-17 on page 55.

Note: The 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 2-17 Keyboard and Mouse Connected via Optional USB Extender

Configurations and Cabling

This chapter describes the configuration ranges of the Silicon Graphics Prism Power and Team Systems and provides NUMALink cabling instructions.

System Configurations

The minimum and maximum ranges of the configurable items for the Silicon Graphics Prism Power Systems (ring topology) are listed in Table 3-1, and for the Silicon Graphics Prism Team Systems (router topology) are listed in Table 3-2.

Table 3-1 Silicon Graphics Prism Power System Configuration Ranges

	Configuration Minimum	Configuration Maximum ^a
Compute/expansion modules	1	3
Graphics modules	1	2
Processors	4	8
Memory capacity	2 GB	96 GB
Internal disk storage	1 serial ATA disk	6 serial ATA disks
Internal PCI/PCI-X slots ^b	4	16
Storage device	None	Customer configurable
17U short rack	None	1
39U tall 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 ImageSync card, and one for USB. Therefore, the number of available slots in the base compute module is reduced.

Table 3-2 Silicon Graphics Prism Team System Configuration Ranges

	Configuration Minimum	Configuration Maximum^a
Compute/expansion modules	1	5
Graphics modules	2	4
NUMALink modules (routers)	2	2
Processors	8	16
Memory capacity	8 GB	192 GB
Internal disk storage	1 serial ATA disk	10 serial ATA disks
Internal PCI/PCI-X slots ^b	4	28
Storage device	None	Customer configurable
17U short rack	None	1
39U tall rack	None	2
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 ImageSync card, and one for USB. Therefore, the number of available slots in the base compute module is reduced.

NUMALink Cabling in Ring Topologies

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 if a connection fails between two modules.

The Silicon Graphics Prism system may also be configured in routed configurations, which are described in “NUMALink Cabling in Routed Systems” on page 63.

The bisection bandwidth per node depends on the number of modules on the ring, as well as what type of modules they are; the bisection bandwidth per node is greatest when there are only two modules on the ring. See Table 3-3.

Table 3-3 Bisection Bandwidth of Ring Topology

Module Count	Bisection Bandwidth
2	3.2 GB/s/node
3 (with 2 pipes)	2.13 GB/s/node
3 (with 4 pipes)	1.83 GB/s/node
4 (with 2 pipes)	2.13 GB/s/node
4 (with 4 pipes)	1.6 GB/s/node

Figure 3-1 provides an example of the two-module ring topology.




Figure 3-1 Two-Module Ring Topology Example

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




Figure 3-2 Three-Module Ring Topology Example

Figure 3-3 shows an example of a four-module ring topology.

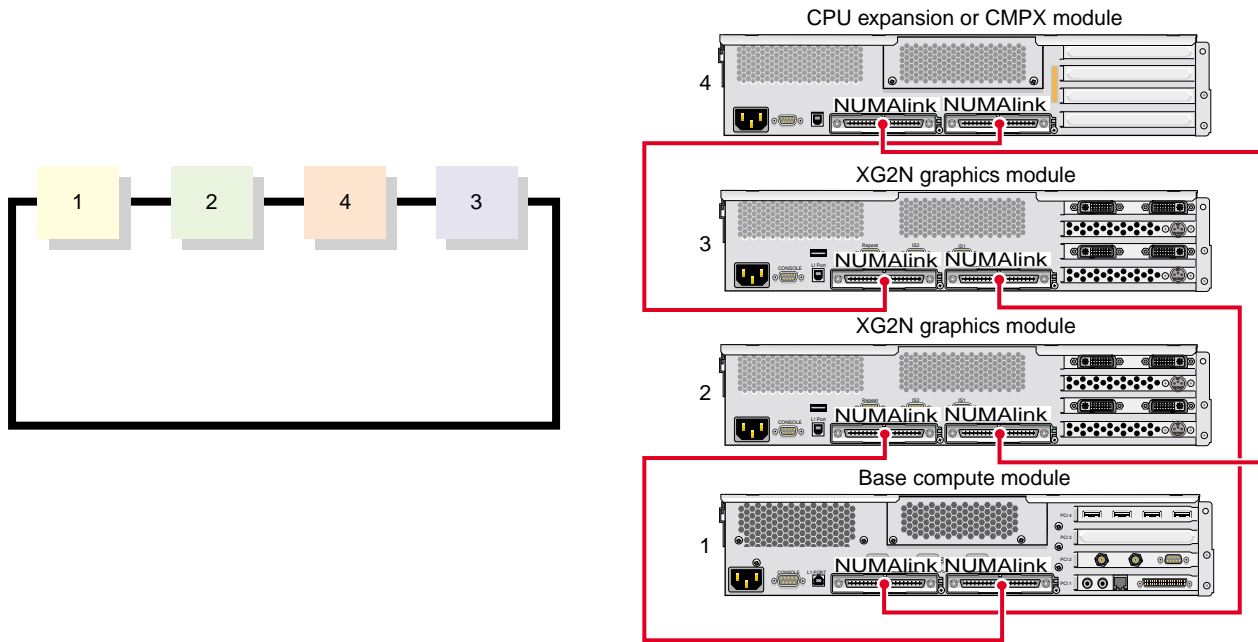


Figure 3-3 Four-Module Ring Topology Example

NUMALink Cabling in Routed Systems

This section describes the NUMALink cabling for systems containing routers.

Prism systems containing routers will have between four and eight base compute, XG2N, CPU, or CMPX modules, as well as two routers. This section provides cabling details for the smallest configuration (four modules, two routers) and largest configuration (eight modules, two routers). Other configurations are similar to the ones shown.

Minimum Routed Configuration

This section describes a minimum routed configuration, having four base compute, XG2N, CPU, or CMPX modules. Table 3-4 details the NUMAlink cabling for this configuration, Figure 3-4 provides a conceptual diagram, and Figure 3-5 shows a representative module ordering.

Table 3-4 Minimum Routed Configuration Cable Chart

Router (a) Cables	Router (b) Cables
Router (a), port A to module 1, port NL-0	Router (b), port A to module 1, port NL-1
Router (a), port B to module 2, port NL-0	Router (b), port B to module 2, port NL-1
Router (a), port C to module 3, port NL-0	Router (b), port C to module 3, port NL-1
Router (a), port D to module 4, port NL-0	Router (b), port D to module 4, port NL-1




Figure 3-4 Minimum Routed Configuration Conceptual Diagram

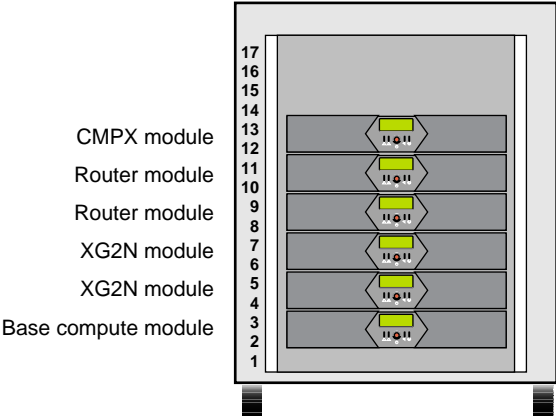


Figure 3-5 Module Ordering in Minimum Routed Configuration

Maximum Routed Configuration

This section describes a maximum routed configuration, having eight base compute, XG2N, CPU, or CMPX modules. Table 3-5 details the NUMAlink cabling for this configuration, Figure 3-6 provides a conceptual diagram, and Figure 3-7 shows a representative the module ordering.

Table 3-5 Maximum Routed Configuration Cable Chart

Router (a) Cables	Router (b) Cables
Router (a), port A to module 1, port NL-0	Router (b), port A to module 1, port NL-1
Router (a), port B to module 2, port NL-0	Router (b), port B to module 2, port NL-1
Router (a), port C to module 3, port NL-0	Router (b), port C to module 3, port NL-1
Router (a), port D to module 4, port NL-0	Router (b), port D to module 4, port NL-1
Router (a), port E to module 1, port NL-0	Router (b), port E to module 1, port NL-1
Router (a), port F to module 2, port NL-0	Router (b), port F to module 2, port NL-1
Router (a), port G to module 3, port NL-0	Router (b), port G to module 3, port NL-1
Router (a), port H to module 4, port NL-0	Router (b), port H to module 4, port NL-1

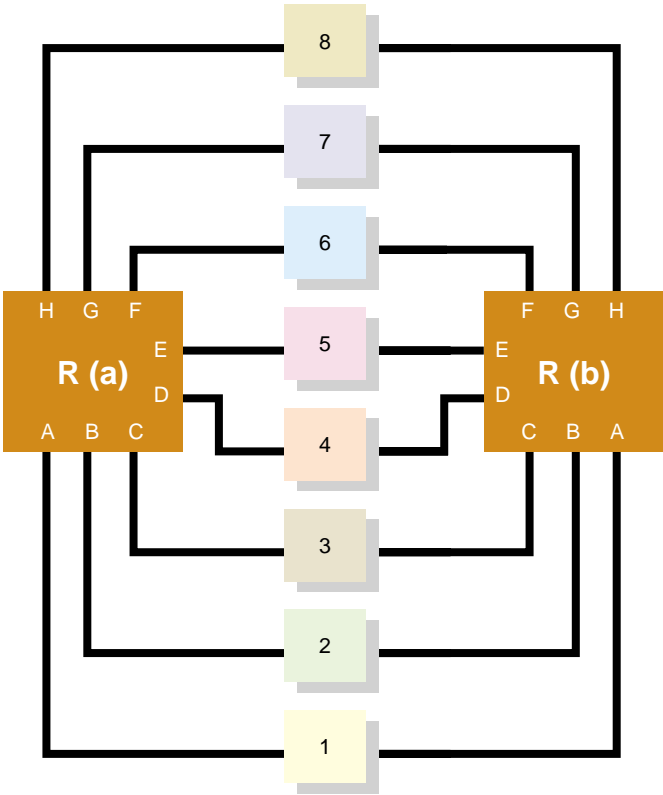


Figure 3-6 Maximum Routed Configuration Conceptual Diagram

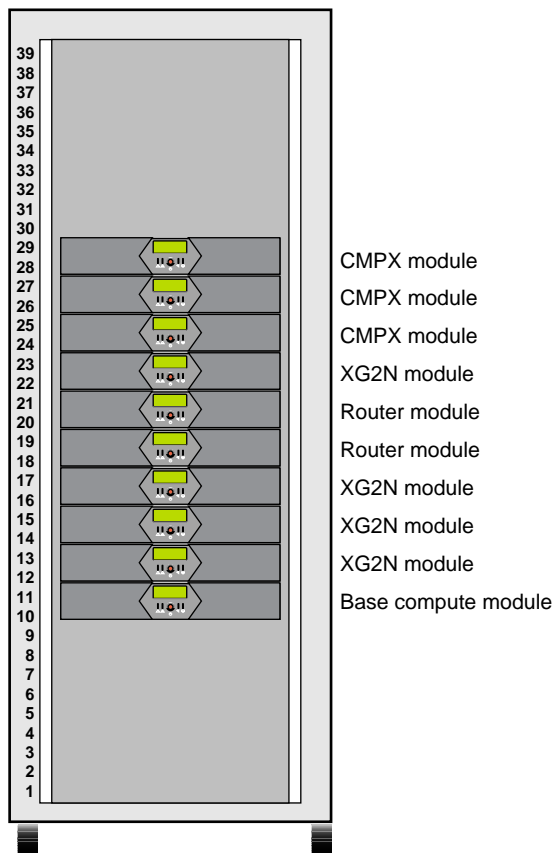


Figure 3-7 Module Ordering in Maximum Routed Configuration

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 70
- “Accessing the System’s Internal Components” on page 72
- “PCI and PCI-X Cards” on page 74
- “Replacing an XG2N Graphics Card” on page 93
- “Disk Drives” on page 100
- “Removing and Replacing Power Supplies” on page 105
- “Memory DIMM Group Placement and Guidelines” on page 109
- “L1 Controller Display” on page 117

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 D, “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 4-1.
2. Lift the hinged cover into the fully open position.




Figure 4-1 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 4-2.




Figure 4-2 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, CPU expansion, or CMPX module. For specific information on removing or replacing XG2N graphics cards, see “Replacing an XG2N Graphics Card” on page 93.

Figure 4-3 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 75
- “Installing a PCI or PCI-X Card” on page 76
- “Removing a PCI or PCI-X Card” on page 84




Figure 4-3 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” on page 22.
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 72.
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 193.

7. Make sure that you read “Safety Instructions” on page 70, and “Rules for Card Installation and Removal” on page 75 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 4-4. (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 4-4 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 4-5, by removing the four Phillips screws.




Figure 4-5 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 4-6, and the blanking plate.




Figure 4-6 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 4-7. If you have installed the card into one of the upper two slots, proceed to step 13.




Figure 4-7 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 4-8.




Figure 4-8 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 4-9.




Figure 4-9 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 "Power-On the System" on page 19.

Removing a PCI or PCI-X 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 “Powering Off Manually” on page 22.
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 193.

7. Make sure that you read “Safety Instructions” on page 70, and “Rules for Card Installation and Removal” on page 75, before installing the card.

8. To access the card, remove the ten Phillips screws, as shown in Figure 4-10, and lift the hinged cover.




Figure 4-10 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 4-11. (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 4-11 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 4-12.




Figure 4-12 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 4-13. Place the card on an ESD-safe surface.




Figure 4-13 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 76. 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 4-14. 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 4-14 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 4-15.




Figure 4-15 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 4-16.




Figure 4-16 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 4-17.




Figure 4-17 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 19.

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” on page 22.
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 72.
7. Make sure that you read “Safety Instructions” on page 70, 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 4-11).

9. Remove the lower PCI/PCI-X support bracket by unscrewing the four Phillips screws, as shown in Figure 4-18.




Figure 4-18 Removing the XG2N Graphics Card Bracket

10. Undo and remove the two screws that hold the card to the back panel, as shown in Figure 4-19.



Caution: Each graphics card has 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 back panel 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 4-20 on page 96.
13. Place the graphics card on an ESD-safe surface or in a protective bag.




Figure 4-19 Remove XG2N Graphics Board Assembly

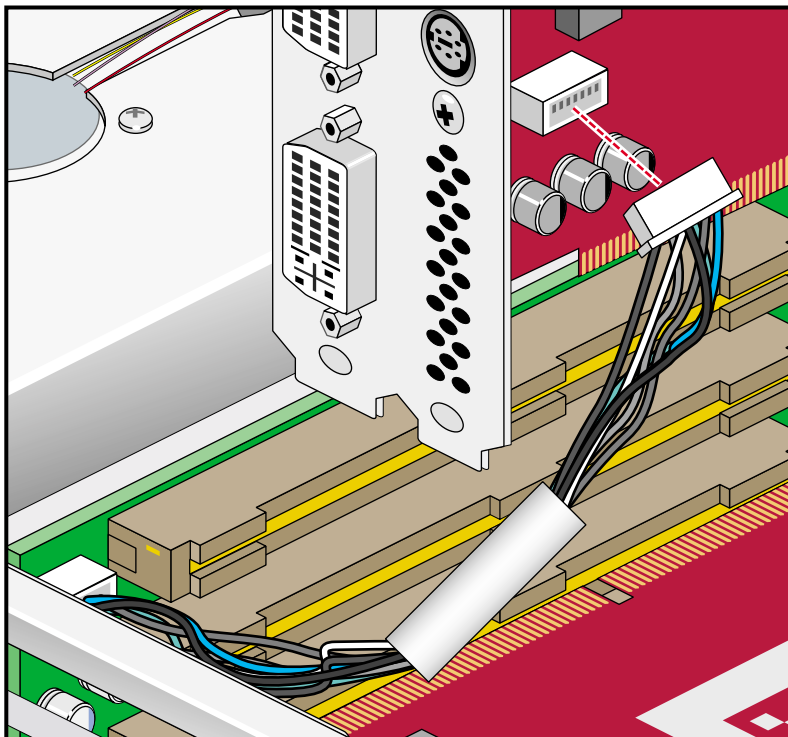


Figure 4-20 Synchronization Cable Connection on Underside of Graphics Board

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 4-20). 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 back panel retaining screws, as shown in Figure 4-21.

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.




Figure 4-21 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 4-22.




Figure 4-22 Replacing the PCI/PCI-X Support Bracket

5. Replace the optional chassis rail (if applicable) by screwing in the five Phillips screws. Reference Figure 4-16 if necessary.
6. Close the hinged cover on the graphics module and screw in the ten Phillips screws to secure the cover, (reference Figure 4-17 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 19.

Disk Drives

Each base compute module contains one or two sled-mounted serial ATA disk drives (see Figure 4-23).




Figure 4-23 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 4-24A.

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

2. Swing the locking handle towards the chassis until the locking handle engages the latch, as shown in Figure 4-24B and Figure 4-24C.
3. Close the bezel door, as shown in Figure 4-24D.




Figure 4-24 Installing a Disk Drive

Removing a Disk Drive

To remove a disk drive, follow these steps:

1. 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 22.

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

2. Open the bezel door located on the left side of the front panel of the module, as shown in Figure 4-25A. (Make sure that you open the door as far as it will open.)
3. Remove the drive by depressing the locking handle with your forefinger (Figure 4-25B). Then swing open the locking handle away from the chassis until the handle disengages the drive connector from the backplane connector (see Figure 4-25C).

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

4. Carefully slide the drive out of the bay (see Figure 4-25D) and gently place it on a flat ESD-safe surface. (Do not use the handle to pull the drive out of the bay.)
5. If you are replacing the disk drive, proceed to “Installing a Disk Drive” on page 101. (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.
6. Close the bezel door.
7. Power on the Silicon Graphics Prism platform as described in “Power-On the System” on page 19.




Figure 4-25 Removing a Disk Drive

Removing and Replacing Power Supplies

Each module in a Silicon Graphics Prism system can contain one or two sled-mounted power supplies (see Figure 4-26). For information on determining the operational status of a power supply, see “Reading Power Supply Status LEDs” on page 129.

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




Figure 4-26 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 22. 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 4-27A.

- b. Swing open the screen cover as shown in Figure 4-27B.
- 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 4-27C.
- d. Gently pull out the power supply from the chassis until it clears the power supply bay as shown in Figure 4-27D. 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 4-28A.

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 4-28B.
- c. After you have installed the power supply, swing the screen cover until it closes as shown in Figure 4-28C.

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 4-28D, and close the bezel door.
4. Power on the Silicon Graphics Prism platform, as described in "Power-On the System" on page 19.




Figure 4-27 Removing a Power Supply




Figure 4-28 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 base compute, XG2N, CPU expansion, or CMPX 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 4-29 on page 110.

DIMMs are installed one per DIMM slot, and must be installed in groups 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 compute 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.

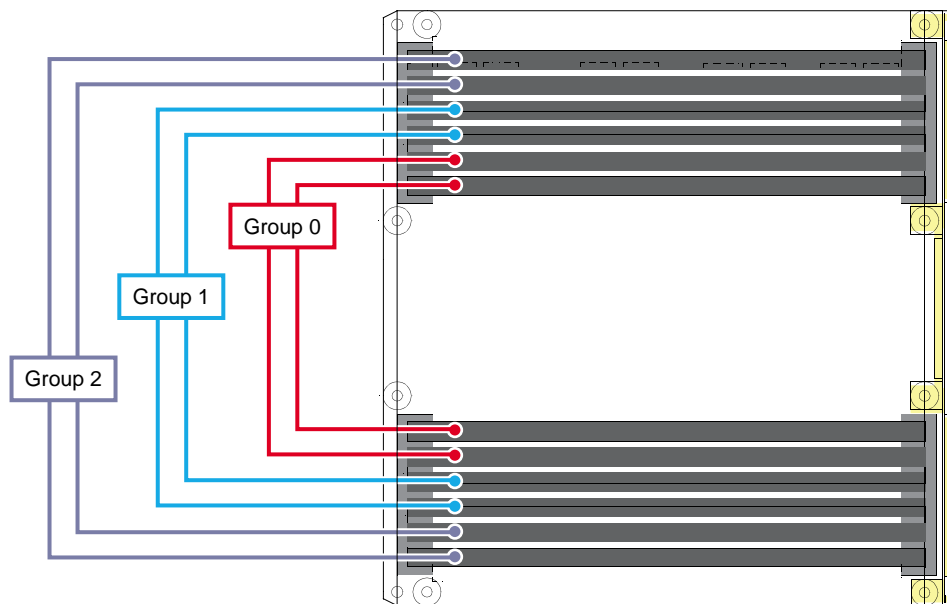


Figure 4-29 Layout of Slots and DIMM Memory Groups



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 22.
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 4-30 and lift and open the hinged cover.




Figure 4-30 Opening Module Cover to Install DIMMs

6. Install the DIMMs, as follows (see Figure 4-31):

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 109.

- 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 4-31.
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 19.




Figure 4-31 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 22.
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 4-32 and open the hinged cover.




Figure 4-32 Opening the Module to Remove a DIMM

6. Remove the DIMM, as follows (see Figure 4-33):

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 109 and see Figure 4-29 on page 110.

- 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 111. 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 19.




Figure 4-33 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 4-34. 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, CPU expansion, or CMPX module.)




Figure 4-34 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 22.
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 4-35, and lift and open the hinged cover.




Figure 4-35 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 4-36A.
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 4-36B. 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 4-36C.
9. Gently disconnect the L1 controller cable from the connector on the L1 controller display, as shown in Figure 4-36D.




Figure 4-36 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 4-37A.
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 4-37B.
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 4-37C.
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 4-37D.




Figure 4-37 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 19.

Troubleshooting

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

- “Troubleshooting Chart” on page 124
- “L1 Controller Error Messages” on page 126
- “Reading Power Supply Status LEDs” on page 129
- “SGI Electronic Support” on page 130
- “Customizing the XF86Config File” on page 133

Troubleshooting Chart

Table 5-1 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 5-1 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 5-2 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 amber service-required LED illuminates on a module.	View the L1 display of the failing module; see Table 5-2 for a description of the error message.
The failure LED illuminates on a module.	View the L1 display of the failing module; see Table 5-2 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 5-1 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 5-2 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 5-2 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 `ll 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 5-2, 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 5-2 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.
Temperature messages: low alt.	
ATTN: TEMP # advisory temperature reached @ xxC xxF	The ambient temperature at the module's air inlet has exceeded 30 °C.
ATTN: TEMP # critical temperature reached @ xxC xxF	The ambient temperature at the module's air inlet has exceeded 35 °C.
ATTN: TEMP # fault temperature reached @ xxC xxF	The ambient temperature at the module's air inlet has exceeded 40 °C.
Temperature messages: high alt.	
ATTN: TEMP # advisory temperature reached @ xxC xxF	The ambient temperature at the module's air inlet has exceeded 27 °C.

Table 5-2 L1 Controller Messages **(continued)**

L1 System Controller Message	Message Meaning and Action Needed
ATTN: TEMP # critical temperature reached @ xxC xxF	The ambient temperature at the module's air inlet has exceeded 31 °C.
ATTN: TEMP # fault temperature reached @ xxC xxF	The ambient temperature at the module's air inlet has exceeded 35 °C.
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 5-3 shows the LED status and the power supply condition that LED status indicates. See “Removing and Replacing Power Supplies” on page 105 for information on removing and replacing a power supply.

Table 5-3 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: <ul style="list-style-type: none"> - 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 5-1 shows the sequence of events that occurs if you use all of the SGI Electronic Support capabilities.

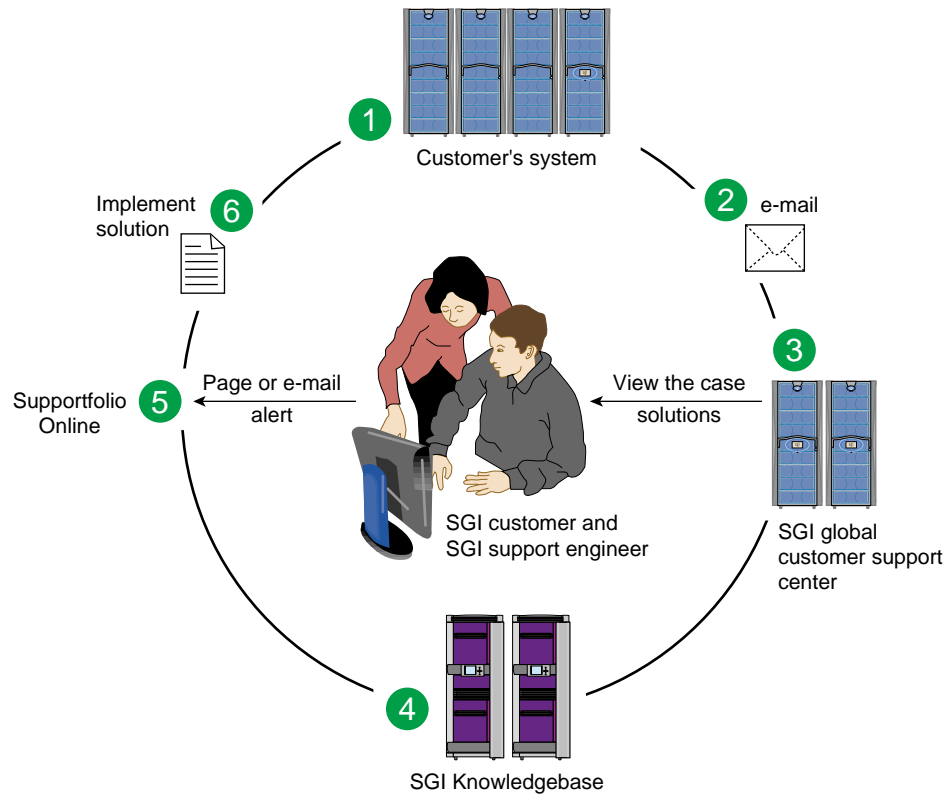


Figure 5-1 Full Support Sequence

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>

Customizing the XF86Config File

The following sections provide information about customizing the XF86Config file for various special configurations.

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 boards.

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 134).

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 134).
 - b. Add that mode to the "Monitor" section of your XF86Config.Stereo file (see the "Example "Monitor" Section for Stereo" on page 135).

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 135).

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 137. 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 137).

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 139).

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 a System for External Genlock or Framelock

External genlock and framelock may be used on systems with at least one optional ImageSync board.

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 -ppipenumber -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 -ppipenumber -LvideofORMAT 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 5-4 on page 142 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` man page (`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 5-4 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 5-2 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 5-2 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 5-3 and the subsection following it show an example of four monitors arranged in a square.

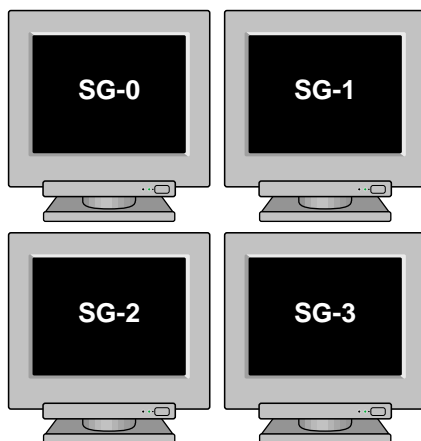


Figure 5-3 Four Monitors in a Square

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

Table 5-5 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
```

Configuring a System for Multiple Xservers

Multiple Xservers allows specific subsets of the keyboards, pointing devices, and monitors attached to a Silicon Graphics Prism system to each be managed by a different Xserver.

Note: The use of multiple Xservers requires ProPack™ 3, Service Pack 4 or a later release of the software.

This section describes a relatively simple configuration. Much more complex configurations are possible, however, and may be accomplished by extending the instructions provided here.

Note: When configuring multiple seats, the best method is to first attach all devices (keyboards, pointing devices, and monitors) and configure the system with a single Xserver. Once this is done, the configuration may be modified to assign individual subsets of these devices to be managed by separate Xservers.

Configuring a system for multi-seat operation involves the following steps, each described in a separate subsection below:

1. Identify the correct event devices (that is, keyboards and pointing devices) for each seat.
2. Create and edit an XF86Config-Nserver file for the desired configuration.
3. Point X to the newly-created XF86Config-Nserver file.

Identifying Event Devices

An “event device” is a keyboard or pointing device. All event devices connected to the system are listed at boot time on lines beginning with the string “input.” These boot messages may be displayed at a Linux command prompt using the `dmesg` command. The output from the `dmesg` command can be quite long, and therefore is usually filtered with a `grep` command. For example:

```
# dmesg | grep ^input
input0: USB HID v1.10 Keyboard [NOVATEK Generic USB Keyboard] on usb1:4.0
input1: USB HID v1.00 Mouse [Logitech N43] on usb1:5.0
input2: USB HID v1.00 Mouse [Logitech N43] on usb1:6.0
input3: USB HID v1.10 Keyboard [NOVATEK Generic USB Keyboard] on usb1:7.0
input4: USB HID v1.00 Keyboard [SILITEK USB Keyboard and Mouse] on usb1:9.0
input5: USB HID v1.00 Mouse [SILITEK USB Keyboard and Mouse] on usb1:9.1
input6: USB HID v1.00 Mouse [Logitech N43] on usb1:10.0
```

All input devices detected during boot-up will have device nodes created for them in the `/dev/input` directory as follows:

- Each keyboard will have an associated event* device node.
- Each pointing device will have both an associated event* device node and an associated mouse* device node.

The mapping of devices to nodes is by number (that is, `input0` maps to `event0`, `input1` maps to `event1`, and so on). The first input that is a pointing device gets mapped to `mouse0`, the next input that is a pointing device gets mapped to `mouse1`, and so on.

The `dmesg` output shown above would therefore create the following mapping:

```
input0: event0
input1: event1, mouse0
input2: event2, mouse1
input3: event3
input4: event4
input5: event5, mouse2
input6: event6, mouse3
```

This mapping can then be used to edit the `XF86Config-Nserver` file, as described in the next subsection, “Creating a Multi-Seat XF86Config File” on page 148.

Creating a Multi-Seat XF86Config File

A multiple-Xserver configuration requires a customized XF86Config file containing a separate ServerLayout section for each Xserver you will be running.

Note: The original ServerLayout section (always identified as “Main Layout”) is typically left unchanged, allowing the system to easily be reconfigured as a single-Xserver system.

Creating a New XF86Config File

Start out by creating a new XF86Config file. The easiest way to do this is to simply make a copy of the system’s regular XF86Config file, as follows:

```
# cp /etc/X11/XF86Config /etc/X11/XF86Config-Nservers
```

(where *N* is the number of servers you will be configuring).

Configuring the Input Devices

Next, configure the input devices as follows:

1. Copy the section beginning:

```
Section "InputDevice"
    Identifier "Keyboard1"
```

and insert a duplicate copy (or copies) below the existing section, until there is one copy for each keyboard (including the original copy in this count).

2. Edit all the keyboard InputDevice sections, in each one changing the driver from "keyboard" to "evdev" and adding an Option line identifying the appropriate event device (in this example, "/dev/input/event0"). The resulting InputDevice sections will look something like the following:

```
Section "InputDevice"
    Identifier "Keyboard1"
    Driver "evdev"
    Option "Device" "/dev/input/event0"
    # ...
EndSection
```

Note: See "Identifying Event Devices" on page 147 for instructions on how to determine the appropriate event device for each section.

Note: You may assign any number of keyboards to a single Xserver, but no keyboard may be assigned to more than one Xserver.

3. Copy the section beginning:

```
Section "InputDevice"
    Identifier "Mouse1"
```

and insert a duplicate copy (or copies) below the existing section, until there is one copy for each pointing device (including the original copy in this count).

4. Edit all the mouse InputDevice sections, changing the Option "Device" line from the default "/dev/input/mice" to one identifying the appropriate event device (in this example, "/dev/input/mouse0"). The resulting InputDevice sections will look something like the following:

```
Section "InputDevice"
    Identifier "Mouse1"
```

```
Driver "mouse"  
Option "Device" "/dev/input/mouse0"  
# ...  
EndSection
```

Note: See “Identifying Event Devices” on page 147 for instructions on how to determine the appropriate event device.

Note: You may assign any number of pointing devices to a single Xserver, but no pointing device may be assigned to more than one Xserver.

Configuring the New ServerLayout Sections

In this new `XF86Config-Nservers` file, perform the following steps:

1. Copy the section beginning:

```
Section "ServerLayout"
    Identifier "Main Layout"
```

and insert a duplicate copy (or copies) below the existing section, until there is one copy for each Xserver you will have (do NOT include the original "Main Layout" copy in this count).

2. While leaving the original ServerLayout section identified as "MainLayout," give each new ServerLayout section a new name. For example, the first server might be named "Layout0":

```
Identifier "Layout0"
```

the next "Layout1," and so on.

3. Within each new Server Layout section, disable (delete or comment out) every screen that will not be used in that layout:

```
Screen "Screen SG-0"      ""      ""      ""      "Screen SG-1"
#   Screen "Screen SG-1"  ""      ""      "Screen SG-0"  ""
```

Note: You may assign any number of screens to a single Xserver, but no screen may be assigned to more than one Xserver.

4. Edit each Server Layout section to make sure than no remaining uncommented screen lists as adjacent another screen that will be managed by a different Xserver:

```
Screen "Screen SG-0"      ""      ""      ""      ""
#   Screen "Screen SG-1"  ""      ""      "Screen SG-0"  ""
```

5. Within each Server Layout section, change the input devices to the correct ones for that Xserver. For example, the first Xserver might use:

```
InputDevice "Mouse1"  "CorePointer"
InputDevice "Keyboard1" "CoreKeyboard"
```

6. Save the `XF86Config-Nservers` file.

For an example ServerLayout section from an `XF86Config-3server` file, see "Example "ServerLayout" Sections for Three Xservers" on page 152. In this example, the first two Xservers manage one screen each, while the third Xserver manages two screens.

Example “ServerLayout” Sections for Three Xservers

```
# *****
# ServerLayout sections.
# *****

Section "ServerLayout"
    Identifier "Main Layout"
    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

Section "ServerLayout"
    Identifier "Layout0"
    Screen "Screen SG-0"      ""      ""      ""      ""
    InputDevice "Mouse1" "CorePointer"
    InputDevice "Keyboard1" "CoreKeyboard"
EndSection

Section "ServerLayout"
    Identifier "Layout1"
    Screen "Screen SG-1"      ""      ""      ""      ""
    InputDevice "Mouse2" "CorePointer"
    InputDevice "Keyboard2" "CoreKeyboard"
EndSection

Section "ServerLayout"
    Identifier "Layout2"
    Screen "Screen SG-2"      ""      ""      ""      "Screen SG-3"
    Screen "Screen SG-3"      ""      ""      "Screen SG-2"  ""
    InputDevice "Mouse3" "CorePointer"
    InputDevice "Keyboard3" "CoreKeyboard"
EndSection
```

Pointing X to the New XF86Config-Nserver File

Once you have created the new `XF86Config-Nserver` file, the last step is to tell X to use the new layouts it contains, rather than the default server layout. To do so, perform the following steps:

1. Make a backup copy of the default single-server `/etc/X11/xdm/gdm.conf` file:


```
# cp /etc/X11/gdm/gdm.conf /etc/X11/gdm/gdm.conf-old
```
2. Edit the `/etc/X11/gdm/gdm.conf` file to use the new server layouts you defined in the `XF86Config` file:

In the `[servers]` section, comment out the standard server, then add the new server layouts you will be using:

```
#0=Standard
0=Layout0
1=Layout1
2=Layout2
```

3. Define each new server layout. For example:

```
[server-Layout0]
name=Layout0 server
command=/usr/X11R6/bin/X :0 -xf86config /etc/X11/XF86Config.3server -layout Layout0
flexible=true
```

For an example of a multi-Xserver `[servers]` section, see “Example `/etc/X11/xdm/gdm.conf` Servers Section for Three Xservers” on page 154.

4. Save the file and reboot the system.

Example /etc/X11/xdm/gdm.conf Servers Section for Three Xservers

```
[servers]

#0=Standard
0=Layout0
1=Layout1
2=Layout2

[server-Standard]
name=Standard server
command=/usr/X11R6/bin/X
flexible=true

[server-Layout0]
name=Layout0 server
command=/usr/X11R6/bin/X :0 -xf86config /etc/X11/XF86Config.3server -layout Layout0
flexible=true

[server-Layout1]
name=Layout1 server
command=/usr/X11R6/bin/X :1 -xf86config /etc/X11/XF86Config.3server -layout Layout1
flexible=true

[server-Layout2]
name=Layout2 server
command=/usr/X11R6/bin/X :2 -xf86config /etc/X11/XF86Config.3server -layout Layout2
flexible=true
```


Technical Specifications

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

- “Environmental Specifications” on page 156
- “Compute Module Specifications” on page 157
- “CMPX Module Specifications” on page 159
- “Router Module Technical Specifications” on page 160
- “SGI TP900 Storage Module Specifications” on page 162
- “Non-proprietary I/O Port Specifications” on page 163
- “XG2N Module Connectors” on page 164

Environmental Specifications

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

Table A-1 Environmental Specifications

Characteristic	Specification
Temperature, operating	+5 °C (+41 °F) to +35 °C (+95 °F) (up to 1500 m [5,000 ft]) +5 °C (+41 °F) to +30 °C (+86 °F) (1500 m to 3000 m [5,000 ft to 10,000 ft])
Temperature, non-operating	-40 °C (-40 °F) to +60 °C (+140 °F)
Humidity	10% to 95% RH, noncondensing
Altitude	Sea level to 40,000 ft (12,000 m) (nonoperating) Sea level to 10,000 ft (3,000 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	0, 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 physical specifications for the compute module.

Table A-4 Compute Module Physical 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.

Router Module Technical Specifications

Table A-7 lists the specifications of each of the 8 router ports.

Table A-7 Router Port Specifications

Port	Quantity	Peak Transfer Rate
NUMAlink	8	3.2 GB/s each direction
L1	1	12 Mbits/s

Table A-8 lists the technical specifications of the router.

Table A-8 Router Technical Specifications

Characteristic	Specification
Height	3.3 in. (83.82 mm)
Width	17.38 in. (441.45 mm)
Depth	27.5 in. (698.50 mm)
Weight	20 lb (9.1 kg)
Heat dissipation	205 Btu/hr maximum
Input power	60 W

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-9 lists the specifications of the short rack.

Table A-9 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-10 lists the specifications of the tall rack.

Table A-10 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-11 lists the specifications of the SGI TP900 storage module.

Table A-11 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-12 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-12 Compute Module Connectors

Port	Connector	Pin Assignments
Optional PCI serial ports ^a	DB-9	See Figure A-6 on page 169
Console port	DB-9	See Figure A-6 on page 169
Ethernet port ^b	RJ-45	See Figure A-7 on page 171 and Table A-17 on page 171
External serial port	IEEE 1284 36-pin MDR multi-port serial adapter	See Figure A-8 on page 172 and Table A-18 on page 172
RT interrupt input and output ports ^c	Stereo jack	Not supported under Linux
L1 port	USB type B	See Figure A-10 on page 175 and Table A-20 on page 175

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-13.




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-13 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-14 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-14 CMPX Module Connectors

Port	Connector	Pin Assignments
Console port	DB-9	See Figure A-6 on page 169
L1 port	USB type B	See Figure A-10 on page 175 and Table A-20 on page 175




Figure A-3 Non-Proprietary Connectors on Rear Panel of CMPX Module

L2 Controller

Table A-15 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-15 L2 Controller Connectors

Port	Connector	Pin Assignments
Console (serial port)	DB-9	See Figure A-6 on page 169
Modem (serial port)	DB-9	See Figure A-6 on page 169
Ethernet port	RJ-45	See Figure A-7 on page 171 and Table A-17 on page 171
L1 ports (four ports)	USB type A	See Figure A-9 on page 174 and Table A-19 on page 174




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.




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

DB-9 Serial Connector

Figure A-6 shows and Table A-16 lists the DB-9 serial connector pin assignments. This connector is used for the console port of the base compute, XG2N, CPU expansion, CMPX, and router modules, and the console and modem ports of the optional L2 controller.



Caution: Although the image sync subsystem also uses DB9 connectors, these connectors, whether on the ImageSync card or on the G2 and G2N bricks, are not serial ports. Connecting a serial device to these connectors may cause damage to both the ImageSync devices and the serial devices.

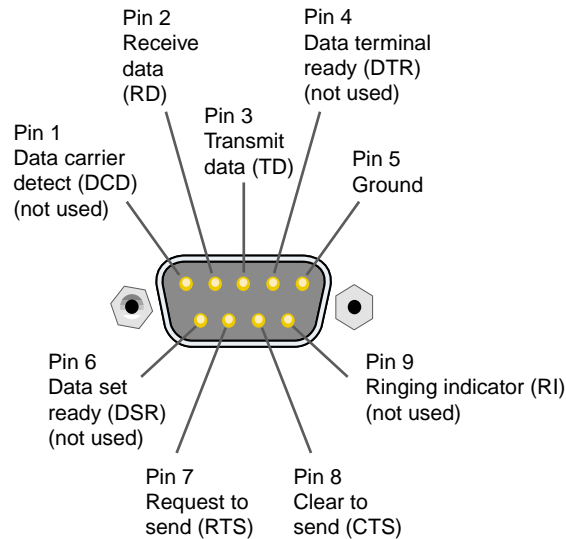


Figure A-6 DB-9 Serial Connector Pin Assignments

Table A-16 DB-9 Serial Connector Pin Assignments

DB-9 Connector Pin	Pin Assignment
1	Data carrier detect (DCD) (not used)
2	Receive data (RD)
3	Transmit data (TD)
4	Data Terminal Ready (DTR) (not used)
5	Ground
6	Data set ready (DSR) (not used)
7	Request to send (RTS)
8	Clear to send (CTS)
9	Ring indicator (RI) (not used)

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-17 shows the pin assignments for the RJ-45 connector.




Figure A-7 RJ-45 Connector Pin Assignments

Table A-17 Ethernet Connector Pin Assignments

10/100BaseT Ethernet Pinouts		1000BaseT Ethernet Pinouts	
Pin	Assignment	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-18 lists the pin assignments for the multi-port serial adapter connector.

Table A-18 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-18 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-19 lists the pin assignments.

Table A-19 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-20 lists the pin assignments.




Figure A-10 Pin Number Locations for USB Type B Connector

Table A-20 Pin Assignments for USB Type B Connector

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

Installing Silicon Graphics Prism Modules Into a Rack

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

- “Choosing Locations Within a Rack” on page 177.
- “Using Silicon Graphics Prism Modules with a Rack” on page 179.
- “Rackmounting with Optional Slide Rails” on page 179.
- “Removing a Rail-Mounted Module from a Rack” on page 193

Choosing Locations Within a Rack

The rack locations for your XG2N and other 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 3, “Configurations and Cabling” consult your SGI sales or support representative.

For power connection examples, see “Positioning and Power for Your Silicon Graphics Prism” on page 9.

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 by displaying 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 179
- “Removing a Rail-Mounted Module from a Rack” on page 193

Rackmounting with Optional Slide Rails

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

- “Determining Module Space Requirements” on page 179
- “Slide Rail Mounting Hardware” on page 180
- “Preparing the Optional Slide Rail Assemblies” on page 181
- “Preparing the Module” on page 184
- “Determining Where to Attach the Slide Rail in the Rack” on page 185
- “Attaching the Slide Rail to the Rack” on page 186
- “Installing Clip Nuts in Rack Rails” on page 189
- “Installing the Module in the Rack” on page 190
- “Adjusting the Position of the Rackmounted Module” on page 192

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	right side: 6 inches (15.24 cm) left side: none required
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	8	Secures the slide rails to the rack rails.
Shoulder washer	8	
Bar nut	4	
10-32 x 1/2-inch Phillips screw	4	Secures the slide rails to their mounting brackets.
Bar nut	2	
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 184). 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:

$$3 \times (\text{the lowest U number that the module will occupy}) - 2.$$

For example, when you want to install the module in locations U9 and U10, count 25 mounting holes ($3 \times 9 - 2$) 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).

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.




Figure B-11 Releasing the Slide Latches

Installing Rack Systems

This Appendix describes how to install a rack that already has your system modules rack-mounted and cabled together. The following information is provided:

- “Pre-Installation Activities” on page 198
- “Unloading and Moving System Equipment” on page 202
- “Removing a Short Rack from the Shipping Crate” on page 206
- “Removing a Tall Rack from the Shipping Crate” on page 208
- “Positioning and Leveling a Single-Rack System” on page 211
- “Positioning and Leveling Multiple Racks” on page 213

After you have completed installing your rack, return to “Positioning and Power for Your Silicon Graphics Prism” on page 9 to continue your installation.

Pre-Installation Activities

Perform the following pre-installation activities:

- Perform site verifications (see “Site Plan Verification” on page 198)
- Gather appropriate tools to complete the installation (see “Tools Required” on page 198)
- Ensure that the correct power receptacle is installed and properly wired (see “Power Receptacle Verification” on page 199)

Note: You can perform the pre-installation activities days or weeks before you receive your system.

Site Plan Verification

Ensure that all site requirements are met before your system arrives. If you have questions about the site requirements or you would like to order full-size floor templates for your site, contact a site planning representative by e-mail (site@sgi.com) or by telephone (+1 715 726 2820).

Tools Required

Table C-1 lists the tools that you need in order to complete the installation.

Table C-1 Installation Tools

Tool	Part Number	Purpose
13-mm wrench	7260744	Adjust the leveling pads.
13-mm socket (3/8-in. drive)	7260726	Remove bracket bolts from tall rack shipping crate.
19-mm socket (3/8-in. drive)	9470618	Remove bolts from short rack shipping crate.
Extension, 6-in. (3/8-in. drive)	7260655	Used with ratchet and sockets.
Ratchet, reversible (3/8-in. drive)	7260755	Used with extension and sockets.
Level, 9-in.	9470556	Level the rack.

Power Receptacle Verification

Ensure that a qualified technician installs the correct power receptacles. Your Silicon Graphics Prism system uses one or two single-phase power receptacles. For North American sites, the single-phase receptacle is a 30-amp, 200- to 240-volt receptacle that has two phase sockets and one ground socket. For international sites, the single-phase receptacle is a 32-amp, 200-volt receptacle that has one phase socket, one neutral socket, and one ground socket. Note that the ground pin is slightly longer than the other two pins on the North American plug.

For North American sites, follow these steps to ensure that a single-phase power receptacle is properly wired:

1. Set the voltmeter to a high AC voltage range.
2. Check the voltage between socket X and socket Y (see Figure C-1). The meter should read between 200 and 240 VAC.
3. Check the voltage between socket X and the ground socket. The meter should read approximately 120 VAC.
4. Check the voltage between socket Y and the ground socket. The meter should read approximately 120 VAC.
5. Check the voltage between the ground socket and an earth-ground location. The meter should read 0 VAC.
6. Change the voltmeter to a low-resistance setting.
7. Measure between the ground socket and an appropriate earth-ground location and ensure that resistance is less than 1 ohm.
8. Repeat steps 1 through 7 for any additional single-phase power receptacles.



Caution: If a voltage reading is incorrect, or if the resistance measured in step 7 is more than 1 ohm, contact a site-approved electrician. Do not proceed with the installation.




Figure C-1 30-amp Single-phase Power Receptacle for North American Sites

For international sites, follow these steps to ensure that a single-phase power receptacle is properly wired:

1. Set the voltmeter to a high AC voltage range.
2. Check the voltage between socket 1 and socket 2 (see Figure C-2). The meter should read between 200 and 240 VAC.
3. Check the voltage between socket 1 (line) and the ground socket. The meter should read between 200 and 240 VAC.

Note: The ground pin is slightly larger in diameter than the other two pins.

4. Check the voltage between socket 2 (neutral) and the ground socket. The meter should read approximately 0 VAC.
5. Check the voltage between the ground socket and an earth-ground location. The meter should read 0 VAC.
6. Change the voltmeter to a low-resistance setting.

7. Measure between the ground socket and an appropriate earth-ground location and ensure that resistance is less than 1 ohm.
8. Repeat steps 1 through 7 for any additional single-phase power receptacles.



Caution: If a voltage reading is incorrect, or if the resistance measured in step 7 is more than 1 ohm, contact a site-approved electrician. Do not proceed with the installation.




Figure C-2 32-amp Single-phase Power Plug for International Sites

Unloading and Moving System Equipment

Your Silicon Graphics Prism rack system arrives at the site in cardboard shipping crates.

For a short rack system, the documentation carton and the accessories carton are packed with the system. The documentation carton contains the system manuals as well as warranty and licensing information. The accessories carton contains the I/O, peripheral, and system cables, and any additional connectors or tools that are required for a specific configuration. Any PC, workstation, or terminal is shipped in a separate carton.

For a tall rack system, the system documentation; accessories; and any PC, workstation, or terminal arrive in separate cartons.

This section describes how to unload and transport the system to its designated location, as follows:

- “Unloading the Equipment from the Truck” on page 202
- “Inspecting the Shipping Crate” on page 204
- “Transporting the Shipping Crate” on page 204

Unloading the Equipment from the Truck

If your loading dock is the same height as the transportation vehicle, use a pallet jack to unload the system from the transportation vehicle. The pallet jack should have 48-in. tines or forks. Follow any instructions that are printed on the packing crates.

If the loading dock is not the same height as the vehicle, you must provide a forklift or another approved method to unload the system. You can use a platform or ramp to obtain the desired level as long as the ramp incline does not exceed a ratio of one unit vertical to six units horizontal. For more information on site requirements, contact site planning by e-mail (site@sgi.com) or by telephone (+1 715 726 2820).



Warning: Use two or more people to prevent computer equipment from rolling off the transportation vehicle. Failure to do so could result in serious damage to the computer equipment.

If your site does not have a loading dock, arrange for a forklift to unload the system from the transportation vehicle. Ensure that two or three people are available to help unload the equipment. Move all crates slowly and carefully.

Figure C-3 shows the lift openings and dimensions of a tall rack shipping crate. This figure also shows where to position the pallet jack.




Figure C-3 Dimensions of Tall Rack Shipping Crate

Figure C-4 shows the lift openings and dimensions of a short rack shipping crate. This figure also shows where to position the pallet jack.




Figure C-4 Dimensions of Short Rack Shipping Crate

Inspecting the Shipping Crate

After the system is unloaded from the truck, follow these steps before you unpack it:

1. Ensure that the crates and cartons arrive unopened.
2. Inspect the shipping crate for signs of external damage such as dents, holes, crushed corners, and water marks.
3. Ensure that the tilt watch has not been tripped.
4. If the crate is damaged, file a damage claim with the carrier immediately. In addition, notify your local Customer Support Center (CSC) for any missing, incorrect, or damaged items. For CSC contact information, see <http://www.sgi.com/support/supportcenters.html>.

Transporting the Shipping Crate

Use a pallet jack with forks that are 48 in. (122 cm) long or longer to transport the shipping crate to the designated location. See Figure C-3 and Figure C-4 for the crate dimensions and location to position the pallet jack. For system weight and dimensions, contact site planning by e-mail (site@sgi.com) or by telephone (+1 715 726 2820).

If the crate does not fit through all access doors, you may need to partially disassemble the crate.



Caution: If the system shipping or storage environment is significantly colder than the environment in which it will be installed [40 °F (22 °C) or greater disparity], leave the rack in its shipping crate for at least 24 hours at room temperature before you start the installation. This acclimation prevents damage to the equipment that could result from thermal shock and condensation.

Removing a Short Rack from the Shipping Crate



Warning: Be careful when you unpack and move the short rack system. Ensure that the rack remains on a level surface and that the rack weight remains evenly distributed across the four casters. If you must lift the casters over an obstacle, such as a door threshold, use proper lifting techniques and employ a minimum of two people.



Caution: Do not subject the rack to any unnecessary shocks or vibration while you unpack and install the system.

See Figure C-5 as you follow these steps; the numbered illustrations correspond to the numbered steps.

1. Ensure that the temperature of the rack is acclimated to the environment in which you are installing it.
2. Remove the crate cover.
3. Lift the ramp out of the crate and set it aside.
4. Remove the documentation carton, accessories carton, and cardboard packing material.
5. Lift the sidewalls of the crate up and over the system.
6. Remove the four bolts that secure the rack to the crate. You must reach underneath the crate and feel for the bolts.
7. Align the holes in the edge of the ramp with the pegs in the base of the crate. Ensure that the ramp is secure.
8. Remove the gate pins from the left and right ends of the gate. Then remove the gate.
9. Pull the rack down the ramp.



Warning: The maximum weight of the short rack is 610 lbs (277 kg). Use caution when you roll the rack down the ramp.

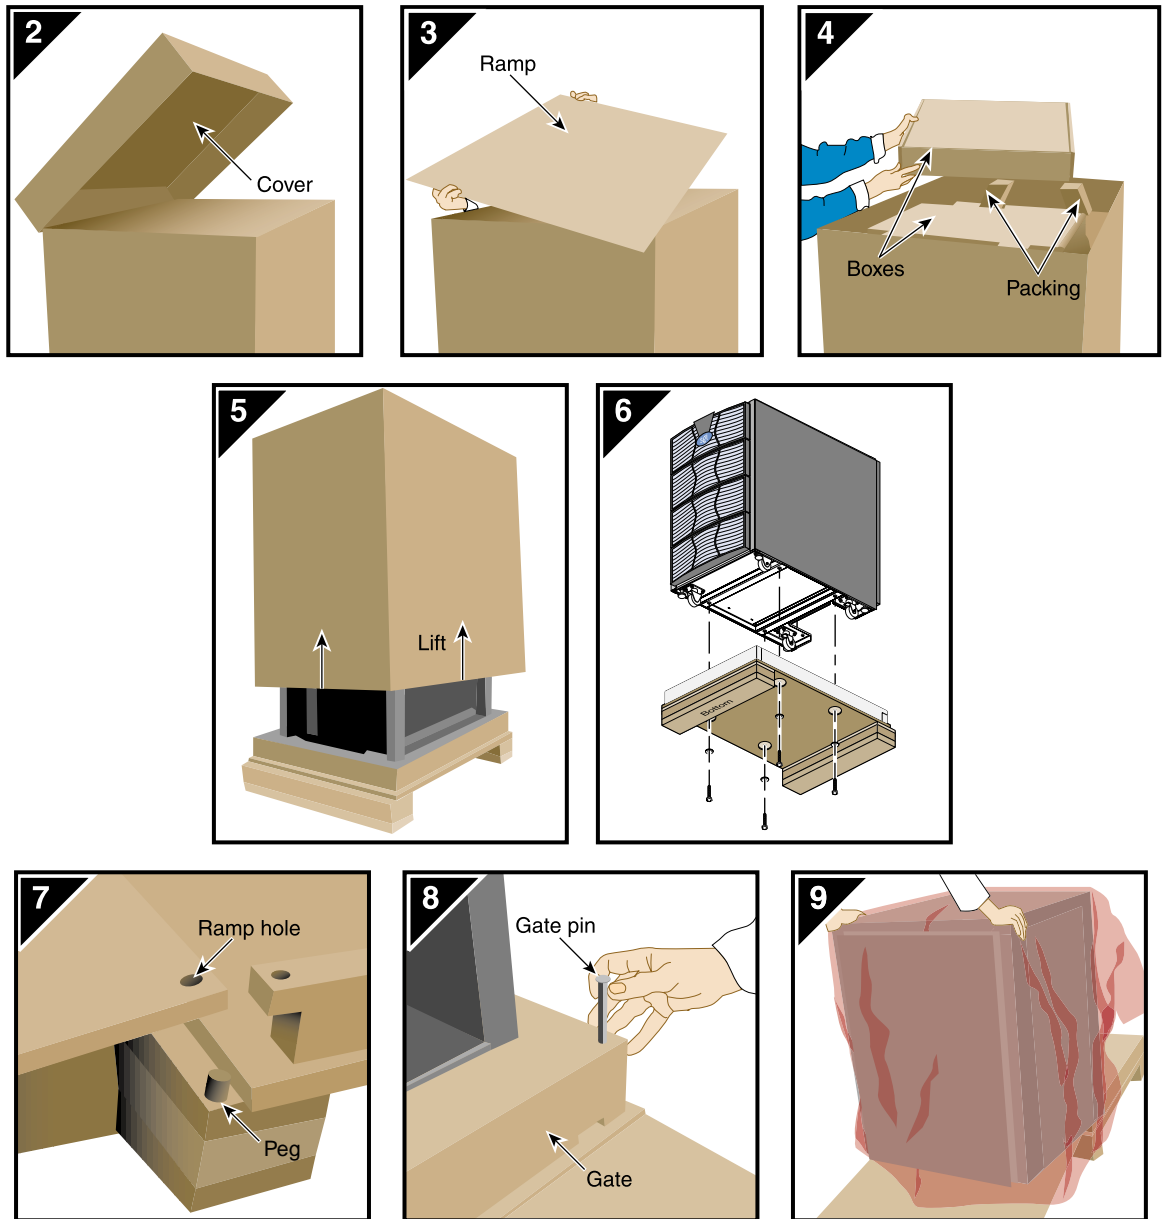


Figure C-5 Removing a Short Rack from the Shipping Crate

Removing a Tall Rack from the Shipping Crate



Warning: In its maximum configuration, a tall rack system weighs approximately 1,366 lb (620 kg). Use caution when you unpack and move this rack. Ensure that the rack remains on a level surface and that the rack weight remains evenly distributed across the four casters.

To unpack a tall rack, you will need the following tools:

- Extension, 6-inch, 3/8-in. drive
- 13-mm standard 3/8-in. drive socket
- Ratchet, reversible, 3/8-in. drive

See Figure C-6 as you follow these steps; the numbered illustrations correspond to the numbered steps.

1. Ensure that the temperature of the rack is acclimated to the environment in which you are installing it and that the system crate is in a stable, upright position.
2. Remove the bands that secure the crate.

Note: Brace the wooden ramp as you remove the horizontal band that surrounds the crate and the wooden ramp. The ramp moves freely after you remove this band.

3. Place the ramp so that the three holes in the edge of the ramp align with the pegs in the base of the pallet deck.
4. Remove the cardboard cover, the two cardboard sidewalls, and the foam cushion.
5. Remove the bolts that secure the rack to the pallet deck, as follows:
 - a. Remove the top four bolts from the rear mounting bracket; do not remove the bottom bolts.
 - b. Remove the four bolts that secure the front mounting bracket and tip tray to the bottom of the pallet deck. Set the tip tray aside.
 - c. Remove the top four bolts from the front mounting bracket. Set the mounting bracket aside.
6. Use two people to roll the rack out of the crate and down the ramp.



Warning: Use extreme caution when you roll the tall rack down the ramp. Personal injury and system damage could result if the rack becomes unbalanced or gains too much momentum when it rolls down the ramp.

7. Bolt the tip tray to the front of the rack before you move the rack to its designated location. This tray prevents the rack from tipping while you move the rack.




Figure C-6 Removing a Tall Rack from the Shipping Crate

Positioning and Leveling a Single-Rack System



Caution: To avoid ESD damage to the electronic components, be sure to position the rack before you remove the ESD bag that covers the rack assembly.

To position and level a single-rack system, follow these steps:

1. Grasp the rear of the rack and roll the rack to its designated location.
2. Remove the ESD bag.
3. If you are installing a tall rack, adjust the leveling bolts, as shown in Figure C-7, until the rack is level.





Figure C-7 Leveling Bolts

4. Ensure that the circuit breaker on the power distribution unit is in the OFF (○) position. Then connect the power cord to a grounded power outlet. Plugging in the power cord grounds the rack.


5. Secure the rack with seismic tie-downs if you are installing the system in an earthquake zone.

Note: Tall and short racks have four threaded holes that are located at the bottom of the rack (see Figure C-8). Use these holes to secure the seismic tie-downs. SGI does not supply the seismic tie-downs.

Top view of short rack



Top view of tall rack



- Seismic tie-down attachment points
- Casters
- Leveling pads

Figure C-8 Seismic Tie-down Attachment Points

Positioning and Leveling Multiple Racks



Caution: To avoid ESD damage to the electronic components, be sure to position the racks before you remove the ESD bags that cover the rack assemblies.

To position and level multiple racks, follow these steps:

1. Grasp the rear of each rack and roll the rack to its designated location.
2. Remove the ESD bags.
3. Adjust the leveling bolts of rack 001, as shown in Figure C-9, until the rack is level.




Figure C-9 Leveling Bolts

4. Using the leveling bolts of rack 002, adjust rack 002 so that the joining holes of rack 002 align with the joining holes of rack 001. See Figure C-10.
5. Using the provided straps, screws, and washers, bolt the racks together in the four designated locations shown in Figure C-10.
6. If your system has additional racks, repeat steps 4 and 5 until all of the racks are bolted together.
7. Ensure that the circuit breakers on the power distribution units are in the OFF (O) position. Then connect the power cords to grounded power outlets. Plugging in the power cords grounds the racks.

8. Secure the racks with seismic tie-downs if you are installing the system in an earthquake zone.

Note: Tall racks have four threaded holes that are located at the bottom of the rack (see Figure C-8 on page 212). Use these holes to secure the seismic tie-downs. SGI does not supply the seismic tie-downs.




Figure C-10 Joining Locations

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

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. Eviter 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 Lasersäteilylle



Warning: Varning: Laserstrålning når denna del är öppnad åla tuijota säteeseenstirra ej in i strålen.



Warning: Varning: Laserstrålning nar denna del är 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: Päristö 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

Numbers

17U rack install, 206-207
39U rack install, 208-210

A

AC power input, 5
adding or replacing
 disk drives, 100
 PCI cards, 71
adding or replacing memory, 109
analog monitors, 145

B

bandwidth
 compute node, 157
base component
 L2 controller, 53
battery statements, 222

C

Chinese Class A regulatory notice, 219
compute module
 connector pin assignments, 163
 general features, 157
 rear panel items, 5

compute node
 bandwidth, 157
connector
 LINK, 5
connector pin assignment
 compute module, 163
 DB9, 169
 L2 controller, 167
 RJ-45, 171
 TP900 storage module, 168
 USB type A, 174
 USB type B, 175
console port, 5
controller
 L2, 53
customer service, xxi

D

DB9
 connector pin assignments, 169
digital monitors, 145
DIMMs
 adding or replacing, 109
 installation, 111-113
 removal, 114-116
disk drives
 adding or replacing, 100
 installation, 101
 removal, 103

disk storage
 TP900, 41
 TP9100, 42
DM8 audio board, 50
dual-channel
 configuring, 138

E

electromagnetic emissions, 218
electrostatic discharge
 regulatory notice, 220
Embedded Support Partner (ESP), 132
ESD precautions, 7
ESP, 132
expansion
 storage
 TP900, 41
 TP9100, 42

F

Failure LED, 4
FCC notice (USA only), 218
Fibre Channel disks, 40
framelock
 configuring, 141
front panel controls & indicators, 3
front panel LEDs, 4
full scene anti-aliasing
 configuring, 136

G

Genlock
 configuring, 141

H

hazard advisory statements, 7

I

Industry Canada Notice (Canada only), 219
install
 disk drive, 101
 memory, 111-113
 PCI/PCI-X card, 76-83
 short rack, 206-207
 tall rack, 208-210
internal view of XG2N, 6
I/O ports, 5
IO10 base I/O card, 40, 49
IO9 base I/O card, 49

K

keyboard extender, 54
Knowledgebase, 132
Korean Class A regulatory notice, 219

L

L1 console port, 5
L1 controller, 4
L1 controller display, 4
L1 port (USB type B), 5
L2 controller, 53
 connector pin assignments, 167
laser compliance statements, 220-221
LEDs, front panel, 4

leveling pads, 47
LINK connector, 5

M

manufacturer's declaration of conformity, 217
manufacturer's regulatory declarations, 217
memory
 adding or replacing, 109
 installation, 111-113
 removal, 114-116
monitor positions, 143
monitor types (digital and analog), 145
MonitorLayout, 145
Multi-port serial adapter, 172

N

NMI button, 4
non-racked chassis, 9
NUMALink
 cables, 12
 connector, 5
 LED, 5

P

PCI/PCI-X card
 installation, 76-83
 removal, 84-92
PDUs, 44
port specifications
 router module, 160
ports, rear panel, 5
Power button, 4
power distribution units, 44

power receptacles, 8, 199-201
Power Systems, 58
Power-button LED, 4
product support, xxi

R

rack
 locations, 177
 numbers, 178
 short install, 206-207
 tall install, 208-210
 types, 47
rack examples
 17U and 39U, 47
rack, short
 specifications, 161
rack, tall
 specifications, 161
rackmounting bricks, 179
R-brick (router module), 36
rear panel of the XG2N, 1
regulatory declarations
 manufacturer, 217
regulatory specifications, 217
 CE notice and Manufacturer's declaration of
 conformity, 217
 Chinese Class A regulatory notice, 219
 electromagnetic emissions, 218
 electrostatic discharge, 220
 FCC notice (USA) only, 218
 Industry Canada notice (Canada only), 219
 Korean Class A regulatory notice, 219
 manufacturer's regulatory declarations, 217
 shielded cables, 220
 VCCI notice (Japan only), 219

- remove
 - disk drive, 103
 - memory, 114-116
 - PCI/PCI-X card, 84-92
- Reset button, 4
- ring topology, 58, 60
- RJ-45
 - connector pin assignments, 171
- router module
 - block diagram, 37
 - features, 36
 - front panel components, 38
 - overview, 36
 - port specifications, 160
 - rear panel components, 39
 - technical specifications, 160
- router topology, 58

S

- safety, 70-71
- safety measures, 8
- SCSI disks, 40
- server model number, 217
- Service-required LED, 4
- setmon command, 141
- SGI Knowledgebase. See Knowledgebase
- shielded cables
 - regulatory notice, 220
- short rack install, 206-207
- site requirements, 198
- slot number (rack), 178
- space requirements, 179
- specifications
 - environmental, 156
 - short rack, 161

- tall rack, 161
 - TP900 storage module, 162
- Status LEDs, 4
- stereo images
 - configuring, 133
- storage expansion
 - TP900, 41
 - TP9100, 42
- Supportfolio Online, 132
- SwapReady
 - configuring, 140
- system rack
 - 17U, 45
 - 39U, 45

T

- tall rack install, 208-210
- tape devices, 49
- tape drives, 49
- Team Systems, 58
- technical specifications
 - router module, 160
- technical support, xxi
- TP900 storage module, 41
 - connector pin assignments, 168
 - specifications, 162
- TP9100, 42
- troubleshooting
 - problems and recommended actions, 124

U

- USB type A
 - connector pin assignments, 174
- USB type B
 - connector pin assignments, 175

V

- VCCI notice (Japan only), 219

X

- XF86Config file
 - configuring for dual-channel, 138
 - configuring for external framelock, 141
 - configuring for external Genlock, 141
 - configuring for full scene anti-aliasing, 136
 - configuring for stereo, 133
 - configuring for SwapReady, 140
 - configuring monitor types, 145
- XG2N module description, 1

