

SGI® Altix® 3700 Bx2 User's Guide

Document Number 007-4377-002

CONTRIBUTORS

Written by Mark Schwenden and Dick Brownell

Illustrated by Chrystie Danzer and Dan Young

Production by Karen Jacobson

Engineering contributions by Michael T Brown, Steve Bowen, David Collins, Hartmut Gottwald, Paul Wiley, Gary Meyer, Galen Flunker, Mark Koneazny, William Kellerman, Tom Hotle, Gregory Thorson and Paul Wiley.

COPYRIGHT

© 2004 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 software described in this document is "commercial computer software" provided with restricted rights (except as to included open/free source) as specified in the FAR 52.227-19 and/or the DFAR 227.7202, or successive sections. Use beyond license provisions is a violation of worldwide intellectual property laws, treaties and conventions. This document is provided with limited rights as defined in 52.227-14.

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, and Altix are registered trademarks, and NUMALink, and SGIconsole 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 or its subsidiaries in the United States and other countries. UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company, Ltd.

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	September 2004, first release
002	November 2004 Updated IO9 information and L1 bezel color

Installing Optional Components 13
Adding or Removing PCI Cards 13
Adding or Removing Disk Drives 14
2. System Control 15
Levels of System Control 17
System Controller Interaction 18
L1 Controller 20
L1 Controller Functions 20
L1 Front Panel Display 22
Optional Ethernet Switch 23
Console Hardware Requirements 24
Operating the L1 25
L1 Mode 26
Viewing System Configuration (from a Brick’s Perspective) 27
Command Targeting 28
Viewing Information, Warnings, and Error Messages 28
Powering On, Powering Off, and Resetting the Brick 29
Console Mode from L1 30
L1 Console Selection 30
Viewing Information, Warnings, and Error Messages 32
L2 Emulator Operation 33
Configuring an Emulated L2’s IP Address 33
Viewing System Configuration 35
Setting Command Targeting 36
Viewing Information, Warnings, and Error Messages 38
Powering On, Powering Off, and Resetting the System From the L2 Emulator 38
Console Mode from the L2 Emulator 39
Console Selection 40
L1 Mode from L2 Emulation 42
Upgrading L1 Firmware 43
Upgrading L1 Firmware 43

3.	System Overview	45
	System Models	47
	System Architecture	49
	System Features	51
	Modularity and Scalability	51
	Distributed Shared Memory (DSM)	51
	Distributed Shared I/O	52
	ccNUMA Architecture	52
	Cache Coherency	52
	Non-uniform Memory Access (NUMA)	53
	Reliability, Availability, and Serviceability (RAS)	53
	System Components	55
	Bay (Unit) Numbering	57
	Rack Numbering	57
	Optional System Components	58
4.	Rack Information	59
	Overview	59
	Tall Rack (40U)	60
	Locking the Front Door	62
	Technical Specifications	63
5.	CR-brick	65
	Overview	65
	External Components	69
	Front Panel Components	69
	Rear Panel Components	71
	Internal Components	72
	Processors	73
	Memory DIMMs	73
	Pinouts	73
	Product Options	73
	Important Notes	74
	Technical Specifications	74

6.	IX-brick	.75
	Product Overview	.76
	External Components	.79
	Front Panel Components	.79
	Rear Panel Components	.81
	PCI and PCI-X Card Configuration Guidelines	.83
	Important Installation Considerations	.84
	Supported PCI Cards	.84
	PCI Card Carrier	.85
	Technical Specifications	.85
7.	PX-brick	.87
	Product Overview	.88
	External Components	.90
	Front Panel Components	.90
	Rear Panel Components	.92
	PCI and PCI-X Card Configuration Guidelines	.94
	Important Installation Considerations	.95
	Supported PCI and PCI-X Cards	.95
	PCI Card Carrier	.96
	Technical Specifications	.96
8.	R-brick	.97
	Overview	.97
	External Components	100
	Front Panel Components	100
	Rear Panel Components	101
	Technical Specifications	103
9.	Power Bay	105
	Overview	105
	External Components	107
	Front Panel Components	107
	Rear Panel Components	108

	Distributed Power Supplies109
	Technical Specifications111
10.	SGI TP900 Storage Module113
	Overview113
	External Components114
	Front Panel Components.114
	Disk Drive Carrier Module115
	Dummy Carrier Module115
	Rear Panel Components116
	Power Supply Module117
	Input/Output (I/O) Module117
	Loopback module117
	Cooling Module118
	Blank Module.118
	Terminator Plug118
	Technical Specifications119
	Product Options119
11.	D-brick2 Storage Module121
	D-brick2 Overview122
	Functional Description123
	D-brick2 Features123
	D-brick2 External Components123
	Drive Carrier Modules and “Dummy” Modules124
	Power Supply/Cooling Modules125
	Operator’s Panel126
	Loop Resiliency Circuit (LRC) Modules127
	Powering the D-brick2 On or Off128
	Powering On the D-brick2128
	Powering Off the D-brick2129
	D-brick2 Technical and Environmental Specifications130
12.	Maintenance and Upgrade Procedures133
	Maintenance Precautions and Procedures133

	Preparing the System for Maintenance or Upgrade	134
	Returning the System to Operation	134
	Installing or Removing Internal Parts	135
	Adding or Replacing a PCI or PCI-X Card	136
	Installing or Replacing a Disk Drive in the IX-brick	144
	Installing a System Disk Drive	144
	Removing a System Disk Drive	147
	Replacing a TP900 Drive Carrier Module	149
	Installing or Replacing a D-brick2 Drive Carrier Module	153
	Installing a Drive Carrier Module	153
	Replacing a Drive Carrier Module	158
13.	Troubleshooting and Diagnostics	161
	Troubleshooting Chart.	162
	L1 Controller Error Messages	164
	SGI Electronic Support.	166
A.	Technical Specifications and Pinouts	169
	System-level Specifications	169
	Physical Specifications	170
	Environmental Specifications	171
	Power Specifications	172
	I/O Port Specifications.	173
	Ethernet Port	174
	Serial Ports	175
B.	Safety Information and Regulatory Specifications	177
	Safety Information	177
	Regulatory Specifications	179
	CMN Number	179
	CE Notice and Manufacturer's Declaration of Conformity	179
	Electromagnetic Emissions	180
	FCC Notice (USA Only)	180
	Industry Canada Notice (Canada Only)	181
	VCCI Notice (Japan Only)	181

Chinese Class A Regulatory Notice181
Korean Class A Regulatory Notice.181
Shielded Cables182
Electrostatic Discharge182
Laser Compliance Statements183
Lithium Battery Statements184
Index185

Figures

Figure 1-1	Power Switch “On” Setting	4
Figure 1-2	D-brick2 PSU/cooling Module Power Switch	5
Figure 1-3	TP900 Storage Module Power Switch	5
Figure 1-4	PDU and PDS Circuit Breaker Switches	6
Figure 1-5	L1 Controller Display Location Example	11
Figure 2-1	SGI Altix 3700 Bx2 Server System Control Network (Example) . . .	16
Figure 2-2	Console Controller Network Diagram Example	19
Figure 2-3	L1 Front Panel	22
Figure 2-4	Ethernet Switch System Controller Block Diagram (Example) . . .	23
Figure 3-1	SGI Altix 3700 Bx2 Series Systems	46
Figure 3-2	SGI Altix 3700 Bx2 Server System (Example Configuration)	48
Figure 3-3	Functional Block Diagram of CR-brick	50
Figure 3-4	Altix 3700 Bx2 System Components Example	56
Figure 4-1	Tall (40U) Rack	61
Figure 4-2	Locking the Front Door of the 40U Rack.	62
Figure 5-1	Front and Side View of CR-brick	66
Figure 5-2	CR-brick Block Diagram	68
Figure 5-3	Front View of CR-brick	70
Figure 5-4	Rear View of CR-brick	72
Figure 6-1	Front View of the IX-brick.	76
Figure 6-2	IX-brick Block Diagram	77
Figure 6-3	Front Panel of IX-brick.	80
Figure 6-4	Rear Panel of IX-brick	82
Figure 6-5	Numbering of IX-brick PCI-X Slots	83
Figure 7-1	Front View of the PX-brick	88

Figure 7-2	Block Diagram of PX-brick 89
Figure 7-3	Front Panel and Display on the PX-brick. 91
Figure 7-4	Rear Panel of the PX-brick 93
Figure 7-5	Numbering of PCI and PCI-X Card Slots. 94
Figure 8-1	Front View of the R-Brick 98
Figure 8-2	R-brick Block Diagram 99
Figure 8-3	Rear View of R-brick	102
Figure 9-1	Front and Rear Views of Power Bay	106
Figure 9-2	Front View of Power Bay	107
Figure 9-3	Rear View of Power Bay	108
Figure 9-4	Front and Rear Views of the DPS	109
Figure 10-1	Front View of TP900 Storage Module.	114
Figure 10-2	Drive Carrier LED Indicators	115
Figure 10-3	Rear View of TP900 Storage Module	116
Figure 10-4	SCSI Terminator Plug	118
Figure 11-1	D-brick2 Front and Rear Views	122
Figure 11-2	D-brick2 Drive Carrier Module and “Dummy” Module	124
Figure 11-3	Anti-Tamper Locking on the D-brick2 Disk Drive	125
Figure 11-4	D-brick2 Operators Panel and Power Supply/Cooling Modules	126
Figure 11-5	D-brick2 Loop Resiliency Circuit (LRC) Module	127
Figure 12-1	Removing a Card Carrier	137
Figure 12-2	Extracting the Carrier Metal Filler Plate	138
Figure 12-3	Adjusting the Carrier Guide Bar	139
Figure 12-4	Mounting Card in Carrier	140
Figure 12-5	Mounting Half-height PCI Card into Carrier	141
Figure 12-6	Moving Carrier Guide Bar to Secure Half-height PCI Card	141
Figure 12-7	Installing the Bracket to Secure a Half-height PCI Card	142
Figure 12-8	Installing a Card in a Slot	143
Figure 12-9	Seating the Card in the Slot	143
Figure 12-10	Location of SCSI Disk Drive Bays	144
Figure 12-11	Installing a Disk Drive	146
Figure 12-12	Removing a Disk Drive.	148
Figure 12-13	Unlocking the Anti-Tamper Lock	150

Figure 12-14	Releasing the Carrier Handle151
Figure 12-15	Carrier Camming Lever152
Figure 12-16	Unlocking the Drive Carrier Module154
Figure 12-17	Opening the Module Handle155
Figure 12-18	Inserting the Disk Drive Module in a D-brick2156
Figure 12-19	Locking the Drive Carrier Module157
Figure 12-20	Unlocking the Disk Drive Module158
Figure 12-21	Removing the Drive Carrier Module.159
Figure 13-1	Full Support Sequence166
Figure A-1	Ethernet Port174
Figure A-2	Serial Port175
Figure B-1	VCCI Notice (Japan Only).181
Figure B-2	Chinese Class A Regulatory Notice181
Figure B-3	Korean Class A Regulatory Notice181

Tables

Table 2-1	L1 Controller Functions	20
Table 2-2	Location Guidelines for Connecting USB/Ethernet Adapters.	24
Table 4-1	Tall Rack Technical Specifications	63
Table 5-1	Configurable Items of CR-brick	73
Table 5-2	CR-brick Technical Specifications	74
Table 6-1	Physical Specifications of the IX-brick	85
Table 6-2	Port Specifications of the IX-brick	86
Table 7-1	Physical Specifications of PX-brick	96
Table 7-2	Port Specifications of the PX-brick	96
Table 8-1	R-brick Technical Specifications	103
Table 8-2	R-brick Port Specifications.	103
Table 9-1	Power Supply LED States	110
Table 9-2	Power Bay Technical Specifications	111
Table 9-3	Power Supply Technical Specifications	111
Table 10-1	Status LEDs	115
Table 10-2	TP900 Storage Module Technical Specifications	119
Table 10-3	Configurable Items.	119
Table 11-1	D-brick2 Weight Information	130
Table 11-2	D-brick2 Enclosure Power Specifications	131
Table 11-3	D-brick2 Ambient Temperature and Humidity Requirements	132
Table 11-4	Additional Environmental Requirements for D-brick2	132
Table 12-1	Customer-replaceable Components and Maintenance Procedures	135
Table 13-1	Troubleshooting Chart.	162
Table 13-2	L1 Controller Messages	164
Table A-1	Altix 3700 Bx2 System Configuration Ranges	169

Table A-2	Altix 3700 Bx2 Physical Specifications	170
Table A-3	Environmental Specifications	171
Table A-4	Power Specifications	172
Table A-5	Ethernet Pinouts	174
Table A-6	Serial Port Pinout	176

About This Guide

This guide provides an overview of the architecture and descriptions of the major components that compose the SGI® Altix® 3700 Bx2 family of servers. It also provides the standard procedures for powering on and powering off the system, basic troubleshooting information, and important safety and regulatory specifications.

Audience

This guide is written for owners, system administrators, and users of SGI Altix 3700 Bx2 series computer systems. It is written with the assumption that the reader has a general knowledge of computers and computer systems.

Important Information



Warning: To avoid problems, you must ask your SGI system support engineer (SSE) to perform all the set up, addition or replacement of parts, cabling, and service of your SGI Altix 3700 Bx2 system, with the exception of the following items that you can perform yourself:

- Using your system console and your L1 controller to enter commands and perform system functions such as powering on and powering off, as described in this guide.
- Adding and replacing PCI and PCI-X cards, as described in this guide.
- Adding and replacing disk drives in the TP900 and D-brick2 storage modules, and in the IX-brick, as described in this guide.
- Using the On/Off switch and other switches (the reset and non-maskable interrupt [NMI] switches on the CR-bricks) on the front panel of your system bricks.
- Using the ESI/ops panel (operating panel) on the D-brick2.

Chapter Descriptions

The following topics are covered in this guide:

- Chapter 1, "Operation Procedures," provides instructions for powering on and powering off your system.
- Chapter 2, "System Control," describes the function of the L1 and L2 controllers and provides instructions for operating the controllers.
- Chapter 3, "System Overview," provides environmental and technical information needed to properly set up and configure the Altix 3700 Bx2 series system.
- Chapter 4, "Rack Information," describes the two rack sizes.
- Chapter 5, "CR-brick," describes all the connectors and LEDs located on the front and rear of the CR-brick.
- Chapter 6, "IX-brick," describes all the connectors and LEDs located on the front and rear of the IX-brick.
- Chapter 7, "PX-brick," describes all the connectors and LEDs located on the front and rear of the PX-brick.
- Chapter 8, "R-brick," describes all the connectors and LEDs located on the front and rear of the R-brick.
- Chapter 9, "Power Bay," describes the function and physical components of the power bay.
- Chapter 10, "SGI TP900 Storage Module," provides a brief overview and functional description of the TP900 storage module.
- Chapter 11, "D-brick2 Storage Module," provides a brief overview and functional description of the D-brick2, and also provides powering on and powering off instructions.
- Chapter 12, "Maintenance and Upgrade Procedures," provides instructions for installing or removing the customer-replaceable components of your system.
- Chapter 13, "Troubleshooting and Diagnostics," provides recommended actions if problems occur on your system.
- Appendix A, "Technical Specifications and Pinouts," provides physical, environmental, and power specifications for your system. Also included are the pinouts for the non-proprietary connectors.
- Appendix B, "Safety Information and Regulatory Specifications," lists all regulatory information related to use of the Altix 3700 Bx2 system in the United States and other countries. It also provides a list of safety instructions to follow when installing, operating, or servicing the product.

Related Publications

The following SGI documents are relevant to the Altix 3700 Bx2 series system:

- SGI Total Performance 900 Storage System User's Guide (P/N 007-4428-xxx)

This fully illustrated guide explains how to operate and maintain the SGI Total Performance 900 (TP900) SCSI storage system.

- SGI Total Performance 9100 (2Gb TP9100) Storage System User's Guide (P/N 007-4522-xxx)

This fully illustrated guide explains how to operate and maintain the 2Gb SGI Total Performance 9100 Fibre Channel storage system.

- SGI InfiniteStorage TP9400 and SGI InfiniteStorage TP9500 and TP9500S RAID User's Guide (P/N 007-4304-xxx)

This fully illustrated guide explains how to operate and maintain the 2Gb SGI TP9400 and 2Gb SGI TP9500 Fibre Channel storage systems.

- SGIconsole Hardware Connectivity Guide (P/N 007-4340-xxx)

This fully illustrated guide explains how to connect the SGIconsole to the various SGI server and graphics system configurations. SGIconsole is a multi-server management system that manages and monitors multiple servers throughout a customer's computing environment, whether those servers are on site or remote.

- SGI L1 and L2 Controller Software User's Guide (P/N 007-3938-xxx)

This guide describes how to use the L1 and L2 controller commands at your system console to monitor and manage your SGI system.

- Man pages (online)

Man pages locate and print the titled entries from the online reference manuals.

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.
- The release notes, which contain the latest information about software and documentation in this release, are in a file named README.SGI in the root directory of the SGI ProPack for Linux Documentation CD.
- You can also view man pages by typing `man <title>` on a command line.

SGI systems include a set of Linux man pages, formatted in the standard UNIX “man page” style. Important system configuration files and commands are documented on man pages. These are found online on the internal system disk (or CD-ROM) and are displayed using the `man` command. For example, to display the man page for the `xscsidisktest` command, type the following on a command line:

```
man xscsidisktest
```

References in the documentation to these pages include the name of the command and the section number in which the command is found.

For additional information about displaying man pages using the `man` command, see `man(1)`.

In addition, the `apropos` command locates man pages based on keywords. For example, to display a list of man pages that describe disks, type the following on a command line:

```
apropos disk
```

For information about setting up and using `apropos`, see `apropos(1)`.

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(x)	Man page section identifiers appear in parentheses after man page names.
GUI element	This font denotes the names of graphical user interface (GUI) elements such as windows, screens, dialog boxes, menus, toolbars, icons, buttons, boxes, fields, and lists.

Product Support

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

- If you are in North America, contact the Technical Assistance Center at +1 800 800 4SGI or contact your authorized service provider.
- If you are outside North America, contact the SGI subsidiary or authorized distributor in your country.

Reader Comments

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

You can contact SGI in any of the following ways:

- Send e-mail to the following address: techpubs@sgi.com
- Use the Feedback option on the Technical Publications Library Web page:
<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.

Operation Procedures

This chapter explains how to operate your new system in the following sections:

- “Precautions” on page 1
- “Connecting a System Console” on page 3
- “Powering the System On and Off” on page 3
- “Using Embedded Support Partner (ESP)” on page 10
- “Monitoring Your Server” on page 11
- “Installing Optional Components” on page 13

Precautions

Before operating your system, familiarize yourself with the safety information in the following sections:

- “ESD Precaution” on page 1
- “Safety Precautions” on page 2

ESD Precaution

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

Wear an SGI-approved wrist strap when you handle an ESD-sensitive device to eliminate possible ESD damage to equipment. Connect the wrist strap cord directly to earth ground.

Safety Precautions



Warning: Before operating or servicing any part of this product, read the “Safety Information” on page 177.



Danger: 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.



Caution: Power off the system only after the system software has been shut down in an orderly manner. If you power off the system before you halt the operating system, data may be corrupted.

Note: A lithium battery is installed on the system I/O card located in the IX-brick.



Warning: If a lithium battery is a soldered part, only qualified SGI service personnel should replace this lithium battery. For a battery of another type, replace it only with the same type or an equivalent type recommended by the battery manufacturer, or an explosion could occur. Discard used batteries according to the manufacturer’s instructions.

Connecting a System Console

The system console enables you to perform the following activities:

- Monitor your system by reading the system's status and error message information generated and displayed by the SGI system's L1 controller.
- Enter L1 controller commands to monitor or change particular system functions. You can, for example, monitor the speed of fans for a particular brick. See the *SGI L1 and L2 Controller Software User's Guide* for descriptions of these commands.
- Power on or power off individual bricks or all bricks (except for the storage modules) in your system.

If you connect a console that contains SGIconsole software, you can perform the functions listed above and have boundary scan capabilities. Boundary scan capabilities enable an SGI system support engineer (SSE) to test the functionality of your system.

Powering the System On and Off

This section explains how to power on and power off individual bricks or your entire Altix 3700 Bx2 system, as follows:

- "Powering On the System" on page 4
- "Powering Off the System" on page 8

For servers with a system console, you can power on and power off individual bricks or the entire system at the system console.

If you are using an SGIconsole, you can monitor and manage your server from a remote location. You can also monitor and manage your server with tools such as VACM, Console Manager, and PCP. For details, see the documentation for the particular tool.

The Embedded Support Partner (ESP) program enables you and your SGI system support engineer (SSE) to monitor your server remotely and resolve issues before they become problems. For details on this program, see "Using Embedded Support Partner (ESP)" on page 10.

Powering On the System

This section describes how to prepare to power on your system, and how to power on your system by using either of the following:

- L1 controller
- System console

Preparing to Power On

To prepare to power on your system, follow these steps:

1. Check to ensure that the cabling between the power distribution unit (PDU) and the wall power-plug receptacle is secure.
2. For each individual brick that you want to power on, make sure that the **PWR** (power) switch is set to the **1** (on) position, as shown in Figure 1-1 for the IX-brick. This will start the L1 controller(s) for the brick(s) when the brick(s) are powered on. The 12-VDC LED on an individual brick illuminates green when an L1 controller has powered on successfully.




Figure 1-1 Power Switch “On” Setting

- 3. If you plan to power on a D-brick2, or an entire server that includes a D-brick2, make sure that the power switch on the rear of each PSU/cooling module (two per D-brick2) is in the 1 (on) position. The location of the power switch on the D-brick2 PSU/cooling module is shown in Figure 1-2.




Figure 1-2 D-brick2 PSU/cooling Module Power Switch

The standard configuration of the TP900 storage module has one power supply, as shown in Figure 1-3. You can add an optional second power supply to the TP900 system to provide n+1 redundant power. The second power supply is identical to the first power supply and would be located in the lower-right section of the TP900 storage module. Both power switches must be in the 1 (on) position to provide redundant power.




Figure 1-3 TP900 Storage Module Power Switch

4. Make sure that the PDU and PDS circuit breaker switches shown in Figure 1-4 are turned on to give power to the server system when the system is powered on.




Figure 1-4 PDU and PDS Circuit Breaker Switches

Note: Verify that the LEDs turn on and light green and that your controllers display that the system is powered on for each segment of the procedure, which indicates that the power-on procedure is proceeding properly. If you have a problem while powering on and an error message appears on the L1 controller, or the system console, see your online log files and the information in “L1 Controller Error Messages” on page 164 to learn what the error message indicates and how to resolve the problem.

Powering On at the System Console

The power-on procedure at a system console varies with your server setup, as follows:

- If you have a system console, you can toggle between L1 and console mode. This enables you to power on your server with L1 commands and view the activity by changing to the console mode.
- If you have a system console connected to a server with an ethernet interface, you can toggle between L1 and console mode, and power on your server with L1 controller commands, and monitor the power-on activity by changing to the console mode.

For detailed instructions on using a system console using the L1 mode, see “Operating the L1” on page 25.

See the section “Console Hardware Requirements” in Chapter 2, and the *SGIconsole Hardware Connectivity Guide*, P/N 007-4340-00x for additional information on the optional SGI console.

Powering On in the L1 Mode

The following sections describe how to power on your system in L1 mode.

The L1 controller display, located on the front of each brick, should display `L1 running` once the power-on procedure starts.

Note: If you have a problem while powering on and an error message appears on your console display, see “L1 Controller Error Messages” on page 164 to learn what the error message indicates and how to resolve the problem.

To power on your system while in the L1 mode, follow these steps:

1. The prompt on your system shows the rack and slot number of the CR-brick to which you have connected your console. If you want to power on the CR-brick (001c05 in our example) indicated in the prompt, enter the following command. (If you want to power on the bricks connected to the CR-brick, proceed to the next step.)

```
001c05-L1> power up
```

2. If you want to power on the bricks connected to the CR-brick, enter the following command:

```
003c01-L1> * power up
```

(* indicates all)

3. From the L1 prompt, display the brick configuration information by entering the following command:

```
001c05-L1> config
```

In L1 mode, you can obtain only limited information about the system configuration. A CR-brick only has information about its attached I/O brick and, if another CR-brick is attached to the I/O brick, information about that CR-brick and its attached I/O brick. An I/O brick only has information about its attached CR-brick and an R-brick only has information about itself.

The L1 controller display, located on the front of each brick, should display `L1 running` once the power-on procedure starts.

Powering Off the System

You can power off individual bricks or your entire server system using either the system console interface or the L1 controller of the primary CR-brick.

Powering Off at the System Console

The power-off procedure at a system console varies with your server setup, as follows:

- If you have a system console connected to a server with an ethernet or other hardware interface, you can toggle between L1 and console mode, power off your server with L1 controller commands, and monitor the power-off activity by changing to the console mode.
- If your system console is running on a server via a serial interface connection, you can also toggle between L1 and console mode. This enables you to power off your server with L1 commands and view the activity by changing to the console mode.

For detailed instructions on using a system console in the L1 mode, see “Operating the L1” in Chapter 2.

The following instructions describe how to power off your system in the L1 mode.

Powering Off in the L1 Mode

The L1 controller display for each brick should display `L1 running` once the power-off procedure starts.

Note: If you have a problem while powering off and an error message appears on your console display, see “L1 Controller Error Messages” on page 164 to learn what the error message indicates and how to resolve the problem.

To power off your system while in the L1 mode, follow these steps:

1. The prompt on your system shows the rack and slot number of the CR-brick to which you have connected your console. If you want to power off the CR-brick (001c05 in our example) indicated in the prompt, enter the following command. (If you want to power off the bricks connected to the CR-brick, proceed to the next step.)

```
001c05-L1> power down
```

2. If you want to power off the bricks connected to the CR-brick, enter the following command:

```
001c05-L1> * power down
```

(* indicates all)

3. From the L1 prompt, display the brick configuration information by entering the following command:

```
001c05-L1> config
```

In L1 mode, you can obtain system configuration information within a specific set of inter connectivity limits. A CR-brick only has information about its attached I/O brick and, any other CR-bricks attached to it.

An I/O brick only has information about its attached CR-brick and an R-brick only has information about itself.

Using Embedded Support Partner (ESP)

Embedded Support Partner (ESP) automatically detects system conditions that indicate potential future problems and then notifies the appropriate personnel. This enables you and SGI system support engineers (SSEs) to proactively support systems and resolve issues before they develop into actual failures.

ESP enables users to monitor one or more systems at a site from a local or remote connection. ESP can perform the following functions:

- Monitor the system configuration, events, performance, and availability.
- Notify SSEs when specific events occur.
- Generate reports.

ESP also supports the following:

- Remote support and on-site troubleshooting.
- System group management, which enables you to manage an entire group of systems from a single system.

For additional information on this and other available monitoring services, see the section “SGI Electronic Support” in Chapter 13.

Monitoring Your Server

You can monitor your Altix 3700 Bx2 server from the following sources:

- On the L1 controller's display at the front of each brick (except the TP900 and D-brick2), as shown in Figure 1-5, you can monitor brick items. For example, you can see if the fans on a particular brick are operating properly.




Figure 1-5 L1 Controller Display Location Example

- You can connect an optional SGIconsole via an Ethernet port adapter or simple serial interface. You will need to connect either a local or remote workstation to monitor the servers via Ethernet.

These console connections enable you to view the status and error messages generated by the L1 controllers in your Altix 3700 Bx2 rack. You can also use these consoles to input L1 commands to manage and monitor your system.

Note: The SGIconsole enables you to use various software tools, such as VACM, Console Manager, and PCP, to manage and monitor your system. See the *SGIconsole Start Here* guide for descriptions of these tools and for references to other documents with information about these tools.

- You can connect a dumb terminal to the CR-brick console port. If you have multiple CR-bricks, connect the dumb terminal to the bottom-most CR-brick in the rack. This connection enables you to view the status and error messages generated by the L1 controller and to enter L1 commands to manage and monitor your system.

See the section “Console Hardware Requirements” in Chapter 2, and the *SGIconsole Hardware Connectivity Guide*, P/N 007-4340-00x for additional information on the optional SGI console.

Installing Optional Components

Besides adding a system console, you can add or replace the following hardware items on your Altix 3700 Bx2 series server:

- Peripheral component interface (PCI) cards into your IX-brick or PX-brick.
- Disk drives to your IX-brick, TP900 disk storage module, or D-brick2 storage module.

The sections that follow discuss these activities in more detail.



Warning: You can add or replace only the items listed in this section. For your safety and for the protection of your server system, contact your SGI system support engineer (SSE) to install any hardware items not listed in this section.



Warning: Before installing, operating, or servicing any part of this product, read the “Safety Information” on page 177.

Adding or Removing PCI Cards

The PCI-X based I/O system, an industry standard for connecting peripherals to a processor, is the primary I/O system for the Altix 3700 Bx2 series server.

The IX-brick and PX-brick subsystems provide PCI and PCI-X support for the Altix 3700 Bx2 server system. The IX-brick provides 12 PCI slots. One of the 12 slots is reserved for a system PCI card. If additional PCI slots are required beyond the 11 remaining slots in the IX-brick, a PX-brick is configured into the system.

See “Adding or Replacing a PCI or PCI-X Card” on page 136 for detailed instructions on installing or removing PCI cards.

Adding or Removing Disk Drives

The IX-brick supports the system boot functions and contains one or two low-profile disk drives. See “Installing or Replacing a Disk Drive in the IX-brick” on page 144 for detailed instructions on installing or removing disk drives.

The TP900 is a SCSI-based disk storage module that provides JBOD (just a bunch of disks) mass storage and that has up to 8 disk drives in each storage module. See “Replacing a TP900 Drive Carrier Module” on page 149 for detailed instructions on installing or removing disk drives.

The D-brick2 is a Fibre Channel-based disk storage enclosure that provides JBOD mass storage and that has up to 16 disk drives in each storage enclosure. See “Installing or Replacing a D-brick2 Drive Carrier Module” on page 153 for detailed instructions on installing or removing disk drives.

System Control

This chapter describes the interaction and functions of system controllers in the following sections:

- “Levels of System Control” on page 17
- “System Controller Interaction” on page 18
- “L1 Controller” on page 20
- “Console Hardware Requirements” on page 24
- “Operating the L1” on page 25
- “L2 Emulator Operation” on page 33
- “Upgrading L1 Firmware” on page 43

The control system for the SGI Altix 3700 Bx2 series servers manages power control and sequencing, provides environmental control and monitoring, initiates system resets, stores identification and configuration information, and provides console/diagnostic and scan interface.

Figure 2-1 shows an example system control network using an optional SGIconsole and separate (remote) workstation to monitor a single-rack Altix 3700 Bx2 system.

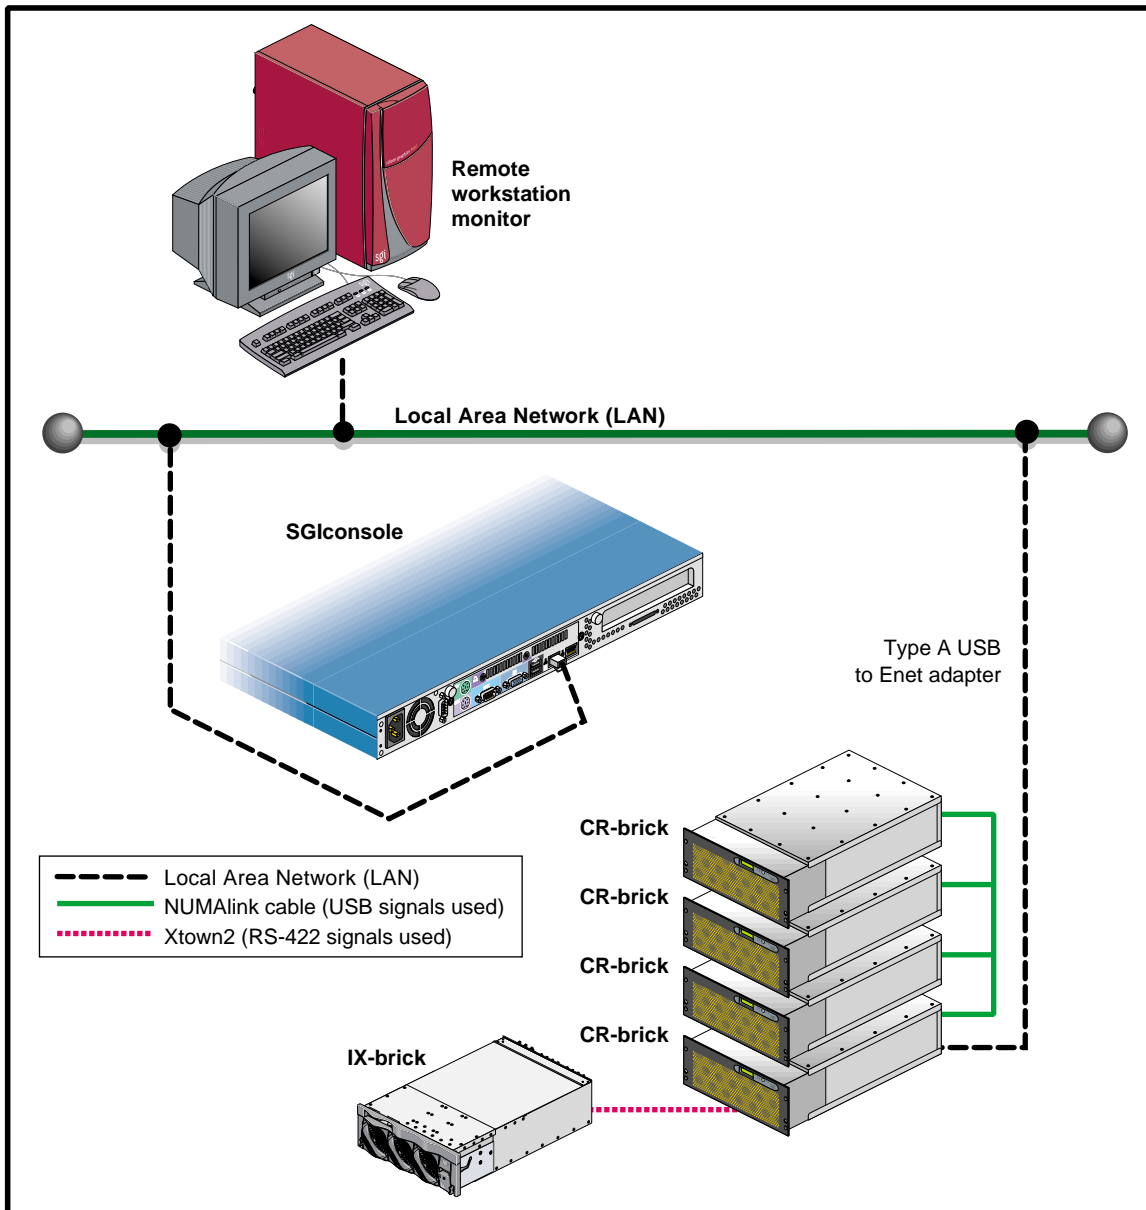


Figure 2-1 SGI Altix 3700 Bx2 Server System Control Network (Example)

Levels of System Control

The system control network configuration of your server will depend on the size of the system and control options selected. The Altix 3700 Bx2 server has optional levels of control, as follows:

- **L1: brick and system-level control.** The L1 system controller is designed into all CR and NUMA4 R-bricks. There is no interconnected L1 function in the TP900 storage module or the D-brick2. Note that the IX-brick is compatible with the new system control functions, but is interconnected to a CR-brick via an SGI Xtown2 interface using RS-422 signals.
- **Optional SGIconsole.** The L1 controllers in the system report and share status information via the NUMALink-4 cables, thus maintaining controller configuration and topology information between all controllers and the system console. For additional information on the SGIconsole options, see:
 - *SGIconsole Hardware Connectivity Guide*, (P/N 007-4340-00x)
 - *SGIconsole 2.x Start Here*, (P/N 007-4356-00x)
 - *Console Manager for SGIconsole Administrator's Guide*, (P/N 007-4477-00x)

Note: The D-brick2, which is not monitored by the L1 controller network, has its own ESI/ops panel module with a microcontroller for monitoring and controlling all elements of the D-brick2.

System Controller Interaction

In all Altix 3700 Bx2 servers all the L1 controllers communicate with each other in the following ways:

- All CR and NUMALink-4 R-bricks communicate with each other through the NUMALink connections using low voltage differential signaling (LVDS).
- An L1 controller of an I/O brick communicates with an L1 controller of a CR-brick via the Crosstown2 connections using serial (RS-422) signals.
- If a USB-to-Ethernet adapter is connected to either the CR or NUMALink-4 R-brick, the L1 system controller on that brick emulates an L2 controller. Access to the L2 functionality is made by way of an Ethernet connection.

Figure 2-2 diagrams the paths for an example console interaction between the system L1s and a remote Ethernet based console interface.

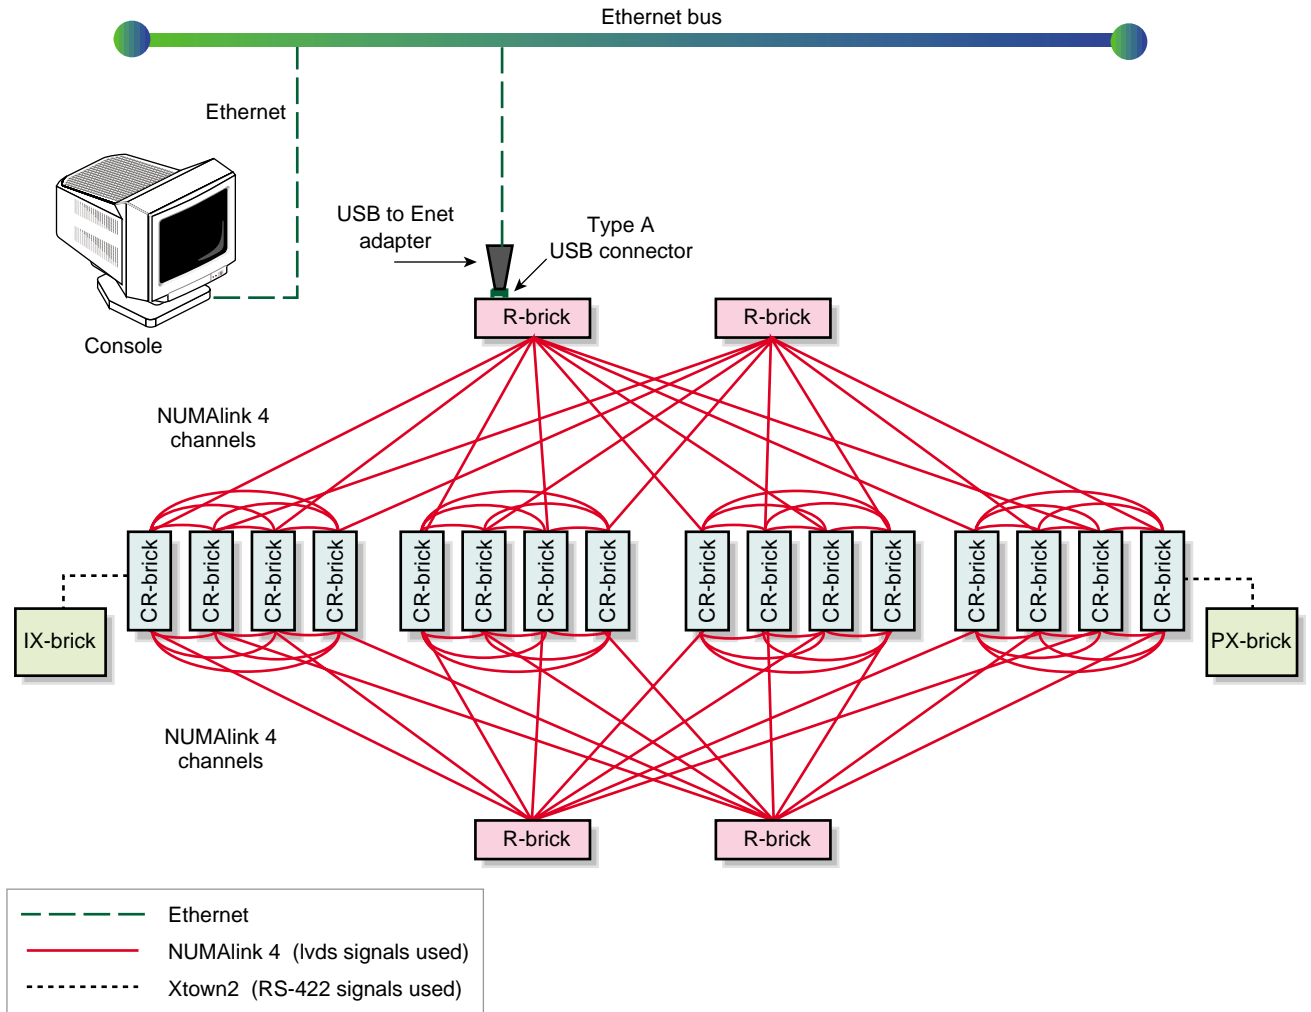


Figure 2-2 Console Controller Network Diagram Example

L1 Controller

All bricks except TP900 storage modules and D-brick2 storage modules have L1 controllers. The following subsections describe the basic features of all L1 controllers:

- “L1 Controller Functions” on page 20
- “L1 Front Panel Display” on page 22

Note: For additional information on L1 controller commands, see the *SGI L1 and L2 Controller Software User’s Guide (007-3938-xxx)*.

L1 Controller Functions

Table 2-1 summarizes the control and monitoring functions that the L1 controller performs. Many of the L1 controller functions are common across all brick types; however, some functions are applicable to a specific brick type.

Table 2-1 L1 Controller Functions

Function	CR-brick	R-brick	IX-brick	PX-brick
Controls voltage regulator modules (VRMs).	X	X	X	X
Controls voltage margining within the brick.	X	X	X	X
Controls and monitors fan speed.	X	X	X	X
Monitors voltage and reports failures.	X	X	X	X
Monitors and reports operating temperature and status of 48-VDC input power.	X	X	X	X
Monitors and controls LEDs.	X	X	X	X
Reads system identification (ID) PROMs.	X	X	X	X
Monitors the On/Off power switch.	X	X	X	X
Monitors the reset switch and the nonmaskable interrupt (NMI) switch.	X			

Table 2-1 L1 Controller Functions **(continued)**

Function	CR-brick	R-brick	IX-brick	PX-brick
Reports the population of the PCI cards and the power levels of the PCI slots.			X	X
Powers on the PCI slots and their associated LEDs.			X	X

L1 Front Panel Display

Figure 2-3 shows the L1 controller front panel on the CR-brick.




Figure 2-3 L1 Front Panel

The front panel display contains the following items:

- 2 x 12 character liquid crystal display (LCD). The display identifies the brick, shows system status, warns of required service, and identifies a failed component.
- Power (On/Off) switch with LED (button with light-emitting diode [LED]).

- Service required LED.
- Failure LED.
- Reset switch and non-maskable interrupt (NMI) switch.

Note: The reset and NMI switch functions are available on the CR-brick only.

Optional Ethernet Switch

Use of an Ethernet switch is the preferred method of connecting remote support hardware to multiple systems.

The optional Ethernet switch provides eight Ethernet connectors. Figure 2-4 shows sample connections between the Ethernet switch, an SGIconsole a CR-brick, and other local SGI systems.




Figure 2-4 Ethernet Switch System Controller Block Diagram (Example)

Console Hardware Requirements

The console type and how these console types are connected to the Altix 3700 Bx2 servers is determined by what console option is chosen.

If you have an Altix 3700 Bx2 server with a dumb terminal, you can connect the terminal via a serial cable to the (DB-9) console port connector on the rear of the CR-brick.

The terminal must be set to the following functional modes:

- Baud rate of 38,400
- 8 data bits
- One stop bit, no parity
- Hardware flow control (RTS/CTS)

You can also connect a USB-to-Ethernet adapter to the CR-brick or R-brick and place it on a network. This starts an L2-level emulator on the brick's system controller. Use the `telnet` command to connect to the CR-brick or R-brick.

Table 2-2 provides information on connection points for the USB-to-Ethernet adapter depending on the system size.

Table 2-2 Location Guidelines for Connecting USB/Ethernet Adapters

System Size	Rack Number	Brick Type	Slot Number (U-Location)
64 Processor or smaller	001	CR	Bottom CR-brick in rack
128 Processor	001	R-brick	U-23
256 Processor	002	R-brick	U-17
	003	R-brick	U-23
512 Processor	003	R-brick	U-05
	003	R-brick	U-35
	009	R-brick	U-05
	009	R-brick	U-35

By default the L2 emulator uses DHCP to obtain an IP address. To set a static IP, access the L1 via the serial port and use the following commands:

```
001c05-L1> L2 ip <ip address> <netmask> <broadcast>
```

```
001c05-L2> reboot_L1
```

These console connections enable you to view the status and error messages generated by the L1 controllers on your system. You can also use the console to input commands to manage and monitor your system(s). For more information on the L2 emulator, see “L2 Emulator Operation” on page 33.

When the system is connected to a network, an optional SGIconsole can also be used to access the system and manage functional processes. For additional information on the SGIconsole options, see:

- *SGIconsole Hardware Connectivity Guide*, (P/N 007-4340-00x)
- *SGIconsole 2.x Start Here*, (P/N 007-4356-00x)
- *Console Manager for SGIconsole Administrator’s Guide*, (P/N 007-4477-00x)

For more details on connecting a console to an Altix 3700 Bx2 series server, see “Connecting a System Console” on page 3. For more information on monitoring your server, see “Monitoring Your Server” on page 11.

Operating the L1

Each CR-brick, and NUMAlink-4 R-brick in the Altix 3700 Bx2 system has an updated and enhanced system control implementation. Note that L1 mode operation is for systems without NUMAlink-4 routers (R-bricks) and when using a serial console connection to one of the CR-bricks.

An L2 emulator mode is used with the system when NUMAlink-4 routers are used and a USB-to-Ethernet adapter is connected. See “L2 Emulator Operation” on page 33 for more information.

Need for a separate optional hardware L2 system controller has been eliminated from the system control network.

The L1 operates in one of these two modes, which are discussed in the sections that follow:

L1 Mode

The L1 prompt is visible and all input is directed to the L1 command processor. The enhanced L1 control system is compatible with the L1 controller in the IX-bricks. The Altix 3700 Bx2 server L1 system control can perform the following:

- Managing power and sequencing control
- Environmental monitoring and control functions
- Initiation of system resets
- Read/write storage for identification and configuration information
- Provides console/diagnostic and scan interface

The L1 controller in each of the CR or NUMA-4 R-bricks is a complete and fully functional system controller. All the bricks are interconnected by NUMALink-4 connectors and each brick shares its system control information with the others.

Console Mode from L1

Output from the system is visible and all input is directed to the system console.

Note: The “console mode from L1” mode is supported only if the system console is connected directly to the console port of the CR-brick.

L1 Mode

If you see a prompt of the following form, the L1 is ready to accept commands.

```
001c19-L1>
```

Common operations are discussed in the following sections:

- “Viewing System Configuration (from a Brick’s Perspective)” on page 27
- “Command Targeting” on page 28
- “Viewing Information, Warnings, and Error Messages” on page 28
- “Powering On, Powering Off, and Resetting the Brick” on page 29

Viewing System Configuration (from a Brick's Perspective)

An L1 has limited knowledge of the system topology, depending on the system's configuration. In configurations without an R-brick, each CR-brick has information about other CR-bricks within the system and its attached I/O brick (if any). An I/O brick only has information about its attached CR-brick.

In some configurations with R-bricks (routers) the L1 will only have knowledge of a portion of the bricks in the system. These R-brick configurations require the use of the emulated L2 (via the use of a USB to ethernet adapter), see "L2 Emulator Operation" on page 33 for further details.

You can view a brick's configuration information with the `config` command as in the following:

```
001c07-L1> config
:0 001c07 LOC
:1 001i11 U-F
:2 001p15 U-G
001c05-L1
```

This example is a system with one CR-brick and two I/O-bricks. The *<number>* that follows the colon (0, 1, and 2, from top to bottom in this example), refers to which local port (NUMAlink or Crosstown) the brick is connected to or accessed through. The local (LOC) brick is the brick that is processing the command.

On all bricks :0 is the local brick, with other values referring to various ports. The specific port description follows the brick's rack/type/slot field: (i.e. LOC, U-F, U-G, etc.)

Command Targeting

All commands entered affect only the local brick. You can target a command to all bricks (including the local brick) by prefixing the command with an asterisk (*).

```
001c05-L1> * version
001c05:
L1 0.7.37 (Image A), Built 11/24/2004 14:59:42 [2MB image]
004i01:
L1 0.7.37 (Image A), Built 11/24/2004 14:59:42 [2MB image]
002c01:
L1 0.7.37 (Image A), Built 11/24/2004 14:59:42 [2MB image]
001x01:
L1 0.7.37 (Image A), Built 11/24/2004 14:59:42 [2MB image]
001c05-L1>
```

Viewing Information, Warnings, and Error Messages

All information, warnings, and error messages generated by any of the system controllers are in the following form:

```
001c05 ERROR: invalid arguments for 'ver' command, try "help ver"
```

The general format of the message includes a brick identification (this is not present if the command was to the local brick only), type of message, and the message. These messages can be the result of an invalid command (as shown in the example) or from tasks running on the L1, such as the environmental monitor.

Each L1 has a log of local events. Use the L1 command `log` to view the event on any of the L1s.

Powering On, Powering Off, and Resetting the Brick

Use a serial terminal cable connected to the console port on any one of the CR bricks. From this perspective all functioning bricks with an L1 can be seen in the output of the L1 config command.

```
001c07-L1> config
:0 001c07 LOC
:1 001c11 U-F
:2 001c15 U-G
:3 001c33 U-H
:10 101i35 IIB
```

To power up/down/reset a non-router configuration from the L1, commands must be targeted to all bricks:

```
001c05-L1> * power up
001c05-L1> * power down
001c05-L1> * reset
```

You can power on and power off a specific brick with the power command:

```
001c05-L1> power up
001c05-L1>
```

This will power on or power off only the locally attached brick. You can target other specific bricks for the power on or power off commands by preceding the power up/down command with a specific rack, and slot destination, such as:

```
001c05-L1> rack 1 slot 11 power up
001c05-L1>
```

or

```
001c05-L1> r 1 s 11 power up
001c05-L1>
```

This command requires from several seconds to several minutes to complete. Remember that TP900 and D-brick2 options will always have to be manually powered on or off.

Console Mode from L1

In console mode, output from the system is visible and all input is directed to the system. To enter console mode, press **Ctrl+D** at the L1 prompt:

```
001c05-L1> Ctrl+D
entering console mode 001c05 console, <CTRL-T> to escape to L1
.
<system output appears here>
.
```

To return to L1 mode, press **Ctrl+T**:

```
Ctrl+T
escaping to L1 system controller
001c05-L1>
```

While in L1 mode, you can enter any L1 command. Once the command is executed, the L1 returns to console mode:

```
re-entering console mode 001c05 console, <CTRL-T> to escape to L1
```

To permanently engage the L1 mode, press **Ctrl+T** and then enter the **l1** command:

```
Ctrl+T
escaping to L1 system controller
001c05-L1> l1
L1 command processor engaged, <CTRL-D> for console mode.
001c05-L1>
```

L1 Console Selection

The brick with which the L1 communicates in console mode is the system console or global master, and you can view and set it with the `select` command. By default, the CR-brick attempts to communicate with its local CPUs when console mode is entered. If the system has been powered on and either one of the bricks received a request to be the system console, then the CR-brick attempts to communicate with that brick. The `select` command by itself shows the current console mode settings:

```
001c05-L1> select
console input: 001c05 console0
console output: not filtered.
```

The following are common subchannels associated with console communications:

- Subchannel 0A specifies Node 0, CPU A.
- Subchannel 0C specifies Node 0 CPU B.
- Subchannel 1A specifies Node 1, CPU A.
- Subchannel 1C specifies Node 1, CPU B.
- Subchannel 2A specifies Node 2, CPU A.
- Subchannel 2C specifies Node 2, CPU B.
- Subchannel 3A specifies Node 3, CPU A.
- Subchannel 3C specifies Node 3, CPU B.
- Node 0 console subchannel.
- Node 1 console subchannel.

The output `console input: 001c05 console0` shows that the system interface console will send input to brick 001c05 and the subchannel to be used is the console0 subchannel.

To change system console status from one brick to the attached CR-brick, use the `select <rack> <slot>` command:

```
001c05-L1> select r 2 s 1
console input: 001c05 console
console output: not filtered.
001c05-L1>
```

To change the subchannel used on the selected brick, use the `select` command followed by the subchannel number or the word `console`:

```
001c05-L1> select sub 0A
console input: 001c05 CPU 0A
console output: not filtered.
001c05-L1>
```

During the boot process on a multi-rack system, there is a window of time in which both CR-bricks are producing output. This output can produce a somewhat jumbled output at the L1. However, you can filter the console output so that the L1 shows output from only the brick chosen to receive console input. You can turn filtering on and off with the `select filter` command.

If you attempt to communicate with a brick that is not responding, a time-out condition results:

```
001c05-L1>
```

```
entering console mode 001c05 console, <CTRL-T> to escape to L1  
no response from 001c05 junk bus console UART:UART_TIMEOUT
```

When this time-out condition occurs, either the brick is hung or the subchannel is incorrect. A brick is identified by its rack, type, and slot (001c05). The structure of the brick location is as follows:

```
rrrbss.p
```

where:

rrr is the rack number.

b is the brick type.

ss is the slot location of the brick.

p is the partition of the brick (not present if the system is not partitioned).
R-bricks are not associated with a partition.

In the example shown above, 001c05 is a CR-brick in rack 001 and slot position 05.

Viewing Information, Warnings, and Error Messages

All information, warnings, and error messages generated by any of the system controllers are in the following form:

```
001c05 ERROR: invalid arguments for 'ver' command, try "help ver"
```

The general format includes a brick identification and the type of message, followed by the message. A message may be the result of an invalid command, as shown in the example, or the result of tasks running on the L1, such as the environmental monitor.

Each L1 has a log of local events. Use the L1 command `log` to view events on any of the L1s.

L2 Emulator Operation

As mentioned in “System Controller Interaction” on page 18 level one system controllers in an Altix 3700 Bx2 system can operate in “L2 emulator mode”. This is to say that they essentially perform the same functions as the L2 system controller on the Altix 3000.

A USB Ethernet adapter can be plugged into the type A connector on the CR-brick. Plugging this adapter causes the system controller to start running as an L2 emulator. The adapter can then be connected to a local area network (LAN) giving network access to the system controller through the emulated L2 interface.

Configuring an Emulated L2's IP Address

This section refers to setting the IP address on the R-brick, but this is also the procedure to set the IP address on the CR-brick when using a USB/ethernet adapter in non-R-brick configurations.

Setting the IP address of the L2 emulator on the target bricks should be done before connecting the bricks to the network as follows:

Connect a serial cable to the console port on the target NUMALink-4 R-brick and get the L1 prompt.

At the L1 prompt type:

```
L1> 12
```

To see if the l2 is running (it will be if the USB/Ethernet adapter is plugged in, even if the adapter is not connected to the network).

If the L2 emulator is not running type:

```
L1>! init 4
```

This switches the system controller to run level 4 and forces the L2 emulator to be started whether or not the USB/ethernet adapter is plugged in. There is a space between the “!” and “init”

Verify the L2 emulator is running again as above.

To set the IP address on the L2 emulator type:

```
L1> 12 ip a.b.c.d 255.255.255.0 a.b.c.255
```

Verifying that the system serial number is set on the L2:

```
L1> l2 serial
```

To set the L2 emulator system serial number:

```
L1> l2 serial set <serial number>
```

Verifying that `msys` is enabled (this allows multiple L2s in a system to exist peacefully with other L2s from another system on the same subnet)

```
L1> l2 msys
```

If `msys` is off, turn it on:

```
L1> l2 msys on
```

Reboot the system controller to make the IP address change take effect.

```
L1> reboot_l1
```

Once this is done for all target R-bricks, connect them to the network (using an optional Ethernet switch if necessary).

The `rackid` on the L2 emulator cannot be set with the L2 `rackid` command, instead it will be inherited from the local L1. As an example: the L2 emulator running on the system controller in 1r14 will have a rack id of 114 (rack * 100 + slot of the local L1).

Once the L2 emulator is running, you can telnet to the L2 emulator, or use an optional SGIconsole. Functionally it is the same as the separate hardware L2 (box) used with older Altix 3000 systems.

Note: The L2 touch displays (used with some older Altix 3000 systems) are not used or supported with the L2 emulator.

After the connection to the L2 emulation controller is established, the following prompt appears, indicating that the L2 emulator is ready to accept commands:

```
L2>
```

Common operations are discussed in the subsections that follow.

Viewing System Configuration

You can use the L2 emulator's `config` command to view the current system configuration from a brick level:

```
L2> config
L2 127.0.0.1: - 001 (LOCAL)
L1 127.0.0.1:0:0 - 001c18
L1 127.0.0.1:1:0 - 001r16
L1 127.0.0.1:2:0 - 001r14
L1 127.0.0.1:3:0 - 001c11
L1 127.0.0.1:4:0 - 001c08
L1 127.0.0.1:5:0 - 001c05
L1 127.0.0.1:5:1 - 001i01
L2>
```

As shown above, `config` produces a list of bricks and their locations in the system and the system controller address of each brick. This is similar to the output from using the `config` command on the L1 with the addition of the emulated L2 IP address and USB port number.

The structure of the brick's address is as follows:

a.b.c.d:x:y

where:

- a.b.c.d is the IP address of the emulated L2. (In the example above, the IP address is 127.0.0.1.)
- x is the USB port number. (In the example above, the port number is 0.)
- y is the L1 index, as follows:
 - 0 is the local brick (the brick to which the USB cable is attached).
 - A number greater than 0 indicates that it is attached directly to or indirectly to the local brick. A brick is identified by its rack, type, and slot (001c05). The structure of the brick location is as follows:

rrrbss.p

where:

- rrr is the rack number.

b	is the brick type.
ss	is the slot location of the brick.
p	is the partition of the brick (not present if the system is not partitioned). R-bricks are not associated with a partition.

In the example shown above, 001c05 is a CR-brick in rack 001 and slot position 05.

Setting Command Targeting

If a command is not understood by the L2 emulator system controller, in general it is passed to the L1 system controllers. The destination determines which L1s receive the command. A destination, specified by the following, is a range of racks and slots:

```
rack <rack list> slot <slot list>
```

The <rack list> specifies a list of racks. This can be a list delimited by commas, such that 2,4,7 specifies racks 2, 4, and 7. You can use a dash to specify a range of racks, such that 2-4 specifies racks 2, 3, and 4. Both nomenclatures can be combined, such that 2-4,7 specifies racks 2, 3, 4, and 7.

You can specify the <slot list> using the same nomenclature. The slot number, sometimes referred to as a bay number, is the unit position number located on the rack, slightly above where the bottom of the brick sits. Each rack unit position number is located toward the top of the two lines that mark the unit position that the number represents. For example, the rack numbering for a brick located in slot 10 would appear on the left front side of the rack.

The slot <slot list> is optional; if not given, then all slots in the specified rack(s) are implied. You should avoid specifying a rack list and a slot list that includes multiple racks and slots, such as rack 2-4,7 slot 1-8,11,13. Generally, you specify a rack and slot together to specify an individual brick.

You can use the aliases *r* and *s* to specify *rack* and *slot*, respectively. You can use the alias *all* or *** in both the <rack list> and the <slot list>, or by themselves, to specify all racks and all slots.

To send a command to all bricks in a partition, enter the following:

```
partition <partition> <cmd>
```

Default Destination

When the L2 emulator starts, the default destination is set to all racks and all slots. You can determine the default destination by using the `destination` command:

```
L2> destination
all racks, all slots
L2>
```

The following command sets the destinations to rack 2 and 3, all slots:

```
L2> r 2,3 destination
2 default destination(s) set
L2>
```

The following example shows what bricks are found in the default destination. If you enter a command not understood by the L2 emulator, the command is sent to these bricks.

Note: In the current implementation, if you add a brick to either rack 2 or 3, it would not be automatically included in the default destination. You would need to reset the default destination.

```
L2> destination
002c05 (127.0.0.1:0:2)
003c05 (127.0.0.1:0:0)
L2>
```

The following command resets the default destination to all racks and all slots:

```
L2> destination reset
default destination reset to all racks and slots
L2>
```

Current Destination

The current destination is a range of racks and slots for a given command. For example, the following command sends the command `<L1 command>` to all bricks in racks 2, 3, 4, and 7:

```
L2> r 2-4,7 <L1 command>
```

This is a one-time destination.

Command Interpretation

Some L2 commands are the same as the L1 commands. In many cases, this is intentional because the L2 emulator provides sequencing that is necessary for a command to function correctly.

When L1 and L2 commands are similar, you can assure that an L1 command is entered for the bricks in the current destination by preceding the command `<L1 command>` with the `l1` command:

```
L2> r 2-4,7 l1 <L1 command>
```

This is a one-time destination.

Viewing Information, Warnings, and Error Messages

All information, warnings, and error messages generated by any of the system controllers are in the following form:

```
001c05 ERROR: invalid arguments for 'ver' command, try "help ver"
```

The general format includes a brick identification and the type of message, followed by the message. A message may be the result of an invalid command, as shown in the example, or the result of tasks running on the L1, such as the environmental monitor.

Each L1 has a log of local events. Use the L1 command `log` to view events on any of the L1s.

Powering On, Powering Off, and Resetting the System From the L2 Emulator

You can power on and power off the system with the `power` command. This command is interpreted by the L2 emulator, because the bricks must be powered on in a specific order.

```
L2> power up  
L2>
```

The `power` command may require several seconds to several minutes to complete. In the example above, all racks and slots in the default destination are affected. Any errors or warnings are reported as described above in “Viewing Information, Warnings, and Error Messages.”

To power on or power off a specific brick, specify a current destination:

```
L2> r 2 s 5 power up
L2>
```

To power on or power off all bricks in a partition, enter the following:

```
L2> partition <partition number> <power up or power down>
```

To reset the system, enter the following:

```
L2> reset
L2>
```

This command restarts the system by resetting all registers to their default settings and rebooting the system controllers. Resetting a running system will cause the operating system to reboot and all memory will be lost.

Console Mode from the L2 Emulator

In console mode, all output from the system is visible and all input is directed to the system.

To enter console mode from L2, press Ctrl+D at the L2 prompt and observe the response:

```
L2> Ctrl+D
entering system console mode (001c05 console0),
<CTRL_T> to escape to L2
.
<system output appears here>
.
```

To return to L2 mode from console mode, press Ctrl+T:

```
Ctrl+T
escaping to L2 system controller
L2>
```

At this point, you can enter any L2 or L1 command. When the command completes, the L2 returns to console mode:

```
Re-entering system console mode (002c05 console0),
<CTRL_T> to escape to L2
```

To permanently engage the L2 mode, press `Ctrl+T` and then enter the `l2` command:

```
Ctrl+T
escaping to L2 system controller
L2> l2
L2 command processor engaged, <CTRL_D> for console mode.
L2>
```

Console Selection

When in console mode, the L2 emulator communicates with the CR-brick set with the `select` command to be the system console or global master. All input from the console is directed to the CR-brick. You can set and view the system console with the `select` command.

The L2 emulator chooses the CR-brick as the default console in the following order of priority:

- The CR-brick in the lowest numbered rack and slot, which has produced console output, and has an attached IX-brick.
- The CR-brick in the lowest numbered rack and slot, which has an attached IX-brick.
- The CR-brick in the lowest numbered rack and slot.

The `select` command by itself shows the current console mode settings:

```
L2> select
known system consoles (non-partitioned)

        001c05-L2 detected

current system console

console input: 001c05 CPU 0A
console output: not filtered
```

The following are six common subchannels associated with console communications:

- Subchannel 0A specifies Node 0, CPU A.
- Subchannel 0C specifies Node 0 CPU B.
- Subchannel 1A specifies Node 1, CPU A.
- Subchannel 1C specifies Node 1, CPU B.

- Node 0 console subchannel.
- Node 1 console subchannel.

The output `console input: 002c05 console0` shows that the L2 emulator will send console input to brick 001c05 and the console subchannel will be used.

To change the brick that will be the system console, use the `select <rack>.<slot>` command, where `<rack>` is the rack and `<slot>` is the slot where the brick is located:

```
L2> select 3.1
console input: 003c01 console
console output: no filtered
console detection: L2 detected
```

To change the subchannel used on the selected brick to be the system console, use the `select subchannel <0A|0C|1A|1C>` command. (Use the `select subchannel console` to select the current console as the subchannel of the brick to be the system console.) For example, to select subchannel b as the subchannel of the brick to be the system console, enter the following:

```
L2> select subchannel 1A
console input: 003c01 console CPU1A
console output: no filtered
```

During the boot process on a multibrick system, there is a window of time in which the CR-bricks are all producing output. This can result in a somewhat jumbled output from the L2 emulator. However, you can filter console output so that the L2 emulator will show output from only the brick chosen to receive console input. You can turn on filtering with the `select filter on` command and turn off filtering with the `select filter off` command.

If you attempt to communicate with a brick chosen to receive console input but that is not responding, a time-out condition results:

```
L2> Ctrl+D
entering console mode 001c05 CPU1A, <CTRL_T> to escape to L2

no response from 001c05 Junk bus CPU1A system not responding
no response from 001c05 Junk bus CPU1A system not responding
```

When this time-out condition occurs, either the brick is hung or the subchannel is not correct.

L1 Mode from L2 Emulation

In L1 mode, the prompt from a single L1 is visible, and all input is directed to that L1 command processor.

To enter L1 mode, enter the rack and a slot followed by `l1` :

```
L2> r 2 s 1 l1
enterling L1 mode 001c05, <CTRL-T> to escape to L2
```

```
001c05-L1>
```

To return to L2 emulation mode, press `Ctrl+T`:

```
001c05-L1> Ctrl+T
escaping to L2 system controller, <CTRL-T> to send escape to L1
L2>
```

At this point, you can enter any L2 command. Once the command is executed, the L2 emulator returns to L1 mode:

```
re-entering L1 mode 002c01, <CTRL-T> to escape to L2
```

```
001c05-L1>
```

To permanently engage the L2 emulation mode, press `Ctrl+T` and enter the `l2` command:

```
002c01-L1> Ctrl+T
escaping to L2 system controller, <CTRL-T> to send escape to L1
L2> l2
L2 command processor engaged, <CTRL-T> for console mode.
```

```
L2>
```

Upgrading L1 Firmware

The L1 firmware is currently distributed as part of the `snxsc_firmware` package. To determine which version of the package is installed on your system console, enter the following command:

```
$> rpm -q snxsc_firmware
```

If the package is installed, the full package name (including the revision) is returned:

```
snxsc_firmware-1.18.3-1
```

The L1 firmware binary and the utilities used to update it are stored in `/usr/cpu/firmware/sysco`.

Upgrading L1 Firmware

Note that a USB-to-Ethernet adapter connection or SGIconsole option is required to execute the commands described in this section. See “Console Hardware Requirements” on page 24 and Figure 2-2 on page 19 for descriptions of the hardware connections.

The L1 firmware consists of three parts:

- Boot image
- A image
- B image

At boot time, the boot image validates the A and B image, and if it is not instructed otherwise, it executes the newer of the two images. Because the L1 is running one of the two images, the image not in use is the image that will be overwritten when the firmware is upgraded. You need to re-boot any L1 update either by power-cycling the brick or by using the L1 command `reboot_l1`.

Typically, you will upgrade the firmware through the network connection from the SGIconsole to the L1:

```
$> /usr/cpu/firmware/sysco/flashsc --12 10.1.1.1 -p  
/usr/cpu/firmware/sysco/l1.bin all
```

This updates all the bricks in the system. The `-p` at the end of the first line instructs the firmware to flash the proms in parallel.

You can update individual bricks by replacing *all* with a rack and slot number:

```
$> /usr/cpu/firmware/sysco/flashsc --12 10.1.1.1  
/usr/cpu/firmware/sysco/l1.bin 1.19
```

This updates only the brick in rack 1, slot 19

System Overview

This chapter provides an overview of the physical and architectural aspects of your SGI Altix 3700 Bx2 series system. The major components of the Altix 3700 Bx2 series systems are described and illustrated.

The Altix 3700 Bx2 series is a family of multiprocessor distributed shared memory (DSM) computer systems that initially scale from 8 to 256 Intel 64-bit processors as a cache-coherent single system image (SSI). Future releases will scale to larger processor counts for SSI applications. Contact your SGI sales or service representative for the most current information on this topic.

In a DSM system, each processor board contains memory that it shares with the other processors in the system. Because the DSM system is modular, it combines the advantages of low entry-level cost with global scalability in processors, memory, and I/O. You can install and operate the Altix 3700 Bx2 series system in a rack in your lab or server room. Each rack holds from 8 to 64 processors in 1 to 8 CR-bricks.

This chapter consists of the following sections:

- “System Models” on page 47
- “System Architecture” on page 49
- “System Features” on page 51
- “System Components” on page 55

Figure 3-1 shows the front views of a multiple-rack system (the Altix 3700 Bx2 system).




Figure 3-1 SGI Altix 3700 Bx2 Series Systems

System Models

The CR-brick contains the processors (8 processors per CR-brick) and two internal high-speed routers. The routers connect to other system bricks via NUMALink cables and expand the compute or memory capacity of the Altix 3700 Bx2. The 40U rack in this server houses all bricks, drives, and other components; up to a 64-processor configuration in a single rack. The Altix 3700 Bx2 server system can expand up to 512 Intel Itanium 2 processors, (a minimum of one IX-brick is required for every 256 processors).

The system requires a minimum of one 40U tall rack with at least one power bay and one single-phase PDU per rack. (The single-phase PDU has two openings with three cables that extend from each opening to connect to the power bay. The three-phase PDU has two openings with six cables that extend from each opening to connect to two power bays.)

You can also add additional racks containing CR-bricks, I/O bricks, and disk storage to your server system.

Figure 3-2 shows an example configuration of a 64-processor Altix 3700 Bx2 server.




Figure 3-2 SGI Altix 3700 Bx2 Server System (Example Configuration)

System Architecture

The Altix 3700 Bx2 computer system is based on a distributed shared memory (DSM) architecture. The system uses a global-address-space, cache-coherent multiprocessor that scales to 64 Intel Itanium 2 processors in a single rack. Because it is modular, the DSM combines the advantages of low entry cost with the ability to scale processors, memory, and I/O independently to a maximum of 256 processors on a single-system image (SSI) at initial release. Larger SSI configurations may be offered in the future, contact your SGI sales or service representative for information.

The system architecture for the Altix 3700 Bx2 system is a fourth-generation NUMAflex DSM architecture known as NUMALink-4. In the NUMALink-4 architecture, all processors and memory are tied together into a single logical system with special crossbar switches (routers). This combination of processors, memory, and crossbar switches constitute the interconnect fabric called NUMALink. There are two router switches in each CR-brick.

The basic building block for the NUMALink interconnect is the CR-brick. A CR-brick contains up to four processor nodes; each processor node consists of a Super-Hub (SHub) ASIC and two 64-bit processors with three levels of on-chip secondary caches. The two Intel 64-bit processors are connected to the SHUB ASIC via a single high-speed front side bus.

The SHub ASIC is the heart of the CR-brick. This specialized ASIC acts as a crossbar between the processors, local SDRAM memory, the network interface, and the I/O interface. The SHub ASIC memory interface enables any processor in the system to access the memory of all processors in the system. Its I/O interface connects processors to system I/O, which allows every processor in a system direct access to every I/O slot in the system.

Another component of the NUMALink-4 architecture is the router ASIC. The router ASIC is a custom designed 8-port crossbar ASIC in the CR-brick. Using the router ASICs with highly specialized NUMALink-4 cables provides a high-bandwidth, extremely low-latency interconnect between all CR-bricks in the system. This interconnection can create a single contiguous system memory of up to 6 TB (terabytes) for a 512 processor system.

Figure 3-3 on page 50 shows a functional block diagram of the Altix 3700 Bx2 series system CR-brick including nodes, routers, and other major components.




Figure 3-3 Functional Block Diagram of CR-brick

System Features

The main features of the Altix 3700 Bx2 series server systems are introduced in the following sections:

- “Modularity and Scalability” on page 51
- “Distributed Shared Memory (DSM)” on page 51
- “Distributed Shared I/O” on page 52
- “Reliability, Availability, and Serviceability (RAS)” on page 53

Modularity and Scalability

The Altix 3700 Bx2 series systems are modular systems. The components are housed in building blocks referred to as bricks. You can add different brick types to a system to achieve the desired system configuration. You can easily configure systems around processing capability, I/O capability, memory size, or storage capacity. You place individual bricks that create the basic functionality (compute/memory, I/O, and power) into custom 19-inch racks. The air-cooled system has redundant, hot-swap fans at the brick level and redundant, hot-swap power supplies at the rack level.

Distributed Shared Memory (DSM)

In the Altix 3700 Bx2 series server, memory is physically distributed among the CR-bricks (compute/router nodes); however, it is accessible to and shared by all CR-bricks within the system. A CR-brick with configured memory but zero processors is generally referred to as a “memory node” or M-brick.

Note the following types of memory:

- If a processor accesses memory that is connected to the same SHUB ASIC on a compute node, the memory is referred to as the node’s *local memory*.
- If processors access memory located in other nodes, the memory is referred to as *remote memory*.
- The total memory within the system is referred to as *global memory*.

Memory latency is the amount of time required for a processor to retrieve data from memory. Memory latency is lowest when a processor accesses local memory.

Distributed Shared I/O

Like DSM, I/O devices are distributed among the compute nodes (each compute node has two I/O ports) and are accessible by all compute nodes through the NUMALink interconnect fabric.

ccNUMA Architecture

As the name implies, the cache-coherent non-uniform memory access (ccNUMA) architecture has two parts, *cache coherency* and *nonuniform memory access*, which are discussed in the sections that follow.

Cache Coherency

The Altix 3700 Bx2 server series use caches to reduce memory latency. Although data exists in local or remote memory, copies of the data can exist in various processor caches throughout the system. Cache coherency keeps the cached copies consistent.

To keep the copies consistent, the ccNUMA architecture uses directory-based coherence protocol. In directory-based coherence protocol, each block of memory (128 bytes) has an entry in a table that is referred to as a directory. Like the blocks of memory that they represent, the directories are distributed among the compute nodes. A block of memory is also referred to as a cache line.

Each directory entry indicates the state of the memory block that it represents. For example, when the block is not cached, it is in an unowned state. When only one processor has a copy of the memory block, it is in an exclusive state. And when more than one processor has a copy of the block, it is in a shared state; a bit vector indicates which caches contain a copy.

When a processor modifies a block of data, the processors that have the same block of data in their caches must be notified of the modification. The Altix 3700 Bx2 server series use an invalidation method to maintain cache coherence. The invalidation method purges all unmodified copies of the block of data, and the processor that wants to modify the block receives exclusive ownership of the block.

Non-uniform Memory Access (NUMA)

In DSM systems, memory is physically located at various distances from the processors. As a result, memory access times (latencies) are different or “non-uniform.” For example, it takes less time for a processor to reference its local memory than to reference remote memory.

Reliability, Availability, and Serviceability (RAS)

The Altix 3700 Bx2 server series components have the following features to increase the reliability, availability, and serviceability (RAS) of the systems.

- **Power and cooling:**
 - Power supplies are redundant and can be hot-swapped.
 - Bricks have overcurrent protection.
 - Fans are redundant and can be hot-swapped.
 - Fans run at multiple speeds in all bricks except the optional R-brick. Speed increases automatically when temperature increases or when a single fan fails.
- **System monitoring:**
 - System controllers monitor the internal power and temperature of the bricks, and automatically shut down bricks to prevent overheating.
 - Memory, L2 cache, L3 cache, and all external bus transfers are protected by single-bit error correction and double-bit error detection (SECDED).
 - The NUMALink interconnect network is protected by cyclic redundancy check (CRC).
 - The L1 primary cache is protected by parity.
 - Each brick has failure LEDs that indicate the failed part; LEDs are readable via the system controllers.
 - Systems support the optional Embedded Support Partner (ESP), a tool that monitors the system; when a condition occurs that may cause a failure, ESP notifies the appropriate SGI personnel.
 - Systems support remote console and maintenance activities.

- **Power-on and boot:**
 - Automatic testing occurs after you power on the system. (These power-on self-tests or POSTs are also referred to as power-on diagnostics or PODs).
 - Processors and memory are automatically de-allocated when a self-test failure occurs.
 - Boot times are minimized.
- **Further RAS features:**
 - Systems support partitioning.
 - Systems have a local field-replaceable unit (FRU) analyzer.
 - All system faults are logged in files.
 - Memory can be scrubbed when a single-bit error occurs.

System Components

The Altix 3700 Bx2 series system features the following major components:

- **40U rack.** This is a custom rack used for both the compute and I/O rack in the Altix 3700 Bx2 system. The power bays are mounted vertically on one side of the rack.
- **CR-brick.** This contains the compute power, standard routers and DIMM memory for the Altix 3700 Bx2 series systems. The CR-brick is 4U high and contains four Super-Hub (SHUB) ASICs, eight Intel Itanium 2 processors, eight NUMALink-4 router ports, and up to 48 memory DIMMs.
- **IX-brick.** This 4U-high brick provides the boot I/O functions and 12 PCI-X slots.
- **PX-brick.** This 4U-high brick provides 12 PCI-X slots on 6 buses for PCI expansion.
- **R-brick.** This is a 2U-high, 8-port metarouter brick.
- **Power bay.** The 3U-high power bay holds a maximum of six power supplies that convert 200-240 VAC to 48 VDC. The power bay has eight 48-VDC outputs.
- **D-brick2.** This is a 3U-high disk storage enclosure that holds a maximum of 16 low-profile Fibre Channel disk drives.
- **TP900 disk storage module.** This is a 2U-high disk storage enclosure that holds a maximum of eight low-profile Ultra320 SCSI disk drives.
- **SGIconsole.** This is an optional combination of hardware and software that allows you to manage multiple SGI servers.

Figure 3-4 on page 56 shows the Altix 3700 Bx2 system components.




Figure 3-4 Altix 3700 Bx2 System Components Example

Bay (Unit) Numbering

Bays in the racks are numbered using standard units. A standard unit (SU) or unit (U) is equal to 1.75 inches (4.445 cm). Because bricks occupy multiple standard units, brick locations within a rack are identified by the bottom unit (U) in which the brick resides. For example, in a tall 40U rack, the CR-brick positioned in U05, U06, and U07 is identified as C05.

Rack Numbering

A rack is numbered with a three-digit number. Compute racks are numbered sequentially beginning with 001. A compute rack is a rack that contains CR-bricks. I/O racks are numbered sequentially and by the physical quadrant in which the I/O rack resides. In a single compute rack system, the rack number is always 001.

Optional System Components

The Altix 3700 Bx2 series system has the following external storage options:

- **Host bus adapter interfaces (HBA)**
 - 2Gbit Fibre Channel, 200MB/s peak bandwidth
 - Ultra320 SCSI, 320 MB/s peak bandwidth
 - Gigabit Ethernet copper and optical
 - 10-Gigabit Ethernet PCI-X
- **JBOD (just a bunch of disks)**
 - SGI TP900 Ultra320 SCSI
- **RAID**
 - SGI TP9300, TP9300S, TP9500 or TP9500S, 2 Gbit Fibre Channel
 - D-brick2, 2 Gbit Fibre Channel
- **Data servers**
 - SGI InfiniteStorage NAS 2000 Fileserver
 - SGI SAN Server 2000/San Gateway
- **Tape libraries**
 - STK L20, L40, L80, L180, and L700
 - ADIC Scalar 100, Scalar 1000, Scalar 10000, and ADIC Scalar i2000
- **Tape drives**
 - STK 9840B, 9940B, LTO, SDLT, and DLT

Availability of optional components for the SGI 3700 Bx2 systems may vary based on new product introductions or end-of-life components. Check with your SGI sales or support representative for the most current information on available product options.

Rack Information

This chapter describes the physical characteristics of the tall (40 U) racks in the following sections:

- “Overview” on page 59
- “Tall Rack (40U)” on page 60
- “Technical Specifications” on page 63

Overview

At the time this document was published only the tall (40U) rack shown in Figure 4-1 was approved for use with the Altix 3700 Bx2 systems.

Tall Rack (40U)

The tall rack (shown in Figure 4-1 on page 61) has the following features and components:

- **Front and rear door.** Both front and rear doors lock to prevent unauthorized access to the system. You lock the front door from the rear of the rack by depressing and turning the actuator handle (see “Locking the Front Door” on page 62). Once the front door is locked, you lock the rear door with a key.
- **Cable entry/exit area** located in the rear floor of the rack. All cables are attached to the rear of the bricks. The rack is mounted on four casters; the two rear casters swivel. The base of the rack has leveling pads, a ground strap, and seismic tie-downs.

The tall rack also has cable entry/exit areas at the top, bottom, and sides of the rack. I/O and power cables pass through the bottom of the rack. NUMALink cables pass through the top and sides of the rack. Cable management occurs in the rear of the rack.

- **One, two, or three power bays**, depending on your computing needs. Each power bay in a compute rack will typically have three to six power supplies. Each power bay in an I/O rack will typically have one to three power supplies.
- **One, two, or three power distribution units (PDUs)** per rack, depending on the number of power bays. A PDU can be a single-phase power or three-phase power distribution unit.



Figure 4-1 Tall (40U) Rack

Locking the Front Door

To lock the front door of the tall rack, follow these steps:

1. Close the front door.
2. In the rear of the rack, press in the door lock actuator and rotate the lever to point downward (see Figure 4-2).




Figure 4-2 Locking the Front Door of the 40U Rack

Technical Specifications

Table 4-1 lists the technical specifications of the tall rack.

Table 4-1 Tall Rack Technical Specifications

Characteristic	Specification
Height	74.25 in. (1885.95 mm)
Width	30.875 in. (784.23 mm)
Depth	53 in. (1346.2 mm)
Weight (empty)	425 lb (192.78 kg)

CR-brick

This chapter describes the function and physical components of the CR-brick in the following sections:

- “Overview” on page 65
- “External Components” on page 69
- “Internal Components” on page 72
- “Pinouts” on page 73
- “Product Options” on page 73
- “Important Notes” on page 74
- “Technical Specifications” on page 74

Overview

The CR-brick is a 4U-high enclosure that contains the compute and memory functionality for the SGI Altix 3700 Bx2 system. Figure 5-1 on page 66 shows a front and side view of the CR-brick.

The CR-brick is divided into separate nodes that are connected internally by high-speed NUMalink channels. Each node has two 64-bit processors and 12 memory DIMM slots connected to a custom-designed application-specific integrated circuit (ASIC). This ASIC is the heart of the CR-brick, and it provides an intelligent interface between the processors, memory, network fabric, and peripheral I/O.




Figure 5-1 Front and Side View of CR-brick

The CR-brick has the following features:

- Up to eight 64-bit RISC processors (two processors per node)
- Large secondary caches (on each processor)
- Up to 48 memory DIMM slots (12 slots per node)
- Up to eight internal NUMAlink channels (two per node)
- Up to eight external NUMAlink channels (two per node)
- Two 1.2-GB/s (each direction) Xtown2 I/O channels
- One USB port for system controller support
- One serial console port
- One L1 controller and LCD display
- Three hot-pluggable fans

Figure 5-2 shows a block diagram of the CR-brick.

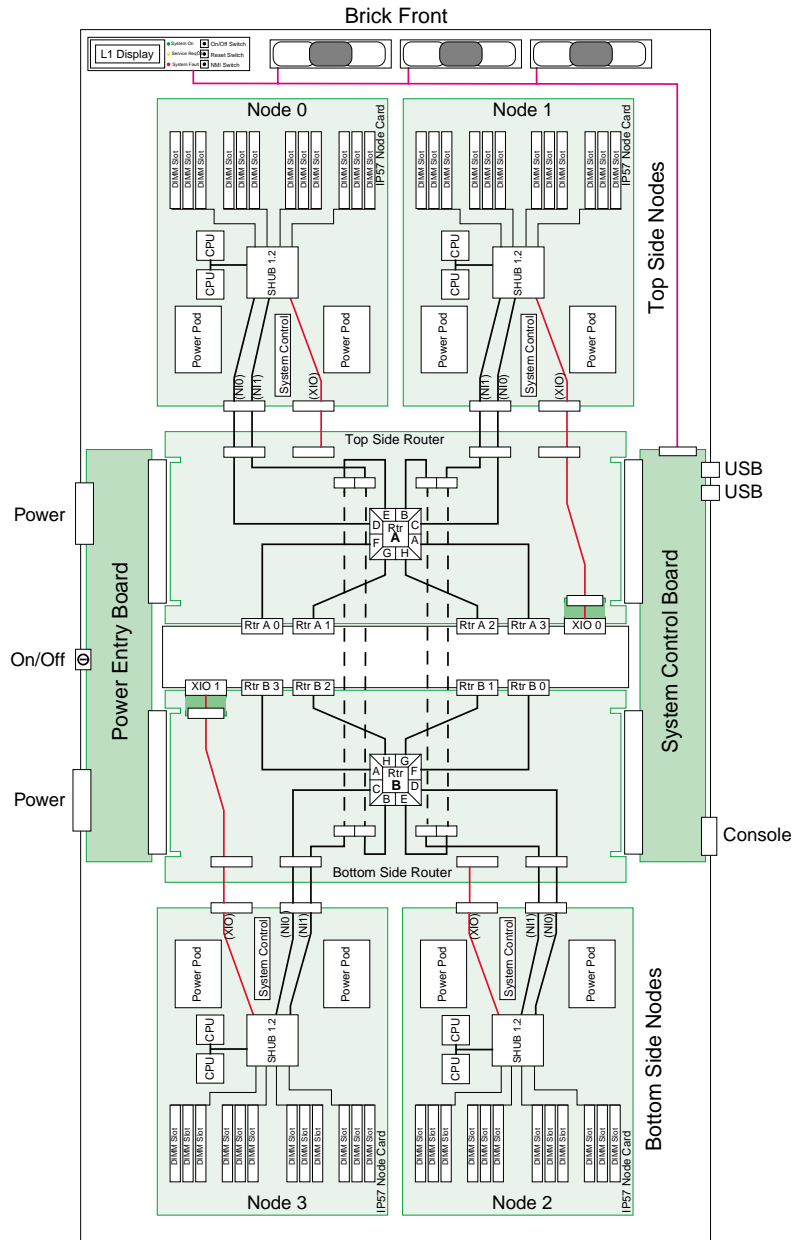


Figure 5-2 CR-brick Block Diagram

External Components

This section describes the external components located on the front and rear panels of the CR-brick.

Front Panel Components

The CR-brick contains the following front panel items (see Figure 5-3 on page 70):

- Three hot-pluggable fans.
- L1 controller display. The display is a 55.7 mm X 32mm backlit liquid crystal display (LCD) that displays system messages. It will display two lines with a maximum of 12 characters on each line.
- On/Off switch with LED. Press this switch to turn on the CR-brick internal components. You can also turn these on at a system console.
- L1 controller switches and LEDs:
 - **On/Off LED.** This LED illuminates green when the CR-brick internal components are on and turns off when they are off.
 - **Service required LED.** This LED illuminates orange to indicate that an item is broken or not operating properly (for example, a fan is off), but the CR-brick is still operating.
 - **Failure LED.** This LED illuminates red to indicate that a system failure has occurred and the CR-brick system is down.
 - **Reset switch.** Press this switch to reset the internal processors of the CR-brick. All register values will be reset to their default states and the operating system will be rebooted. (See the non-maskable interrupt [NMI] to perform a reset without losing the register data.)
 - **Non-maskable interrupt [NMI] switch.** Press this switch to force the CR-brick into power-on diagnostics (POD) mode. The PROM saves the register state for each CPU. On the next reboot, the operating system retrieves the information from PROM and logs it in the system log. SGI service personnel can also perform an error dump or other maintenance action before rebooting the partition. The NMI command is necessary to troubleshoot a system.




Figure 5-3 Front View of CR-brick

Rear Panel Components

The CR-brick has the following rear panel items (see Figure 5-4):

- **Power switch.** Move the power switch to the **I** position to power on the L1 controller within the CR-brick, and to the **0** position to power off the L1 controller. Powering on the L1 controller illuminates the 12-VDC LED green.
- **Two PWR (power) connectors.** These connect the CR-brick to the power bay, to provide 12-VDC and 48-VDC power to the CR-brick.
- **48-VDC and 12-VDC LEDs.** The power switch must be in the ON (**I**) position for these LEDs to be on. The 12-VDC LED illuminates green when the L1 controller is powered on and operating, and the 48-VDC LED illuminates green when the rest of the CR-brick internal components are powered on and operating. You can power on the internal components by pressing the On/Off switch (brick reset button) on the L1 controller panel.
- **NUMAlink connectors A, F, G, H (or Network Interface (NI) connectors).** These upper and lower connectors attach the CR-brick to an R-brick or to other CR-bricks. These are connected with a NUMAlink cable at 3.2 GB/s in each direction.
- **Heartbeat LEDs.** The four heartbeat LEDs turn on and off at a preset rate when the Linux operating system is running. The heartbeat LEDs indicate that the processor is functioning and can process an interrupt from the operating system. The LEDs for each node are arranged in pairs. The lower LED represents the processor 0 heartbeat, and the upper LED represents the processor 1 heartbeat.
- **Node 1 and Node 3 XIO (II - I/O Interface) connectors.** These Xtown2 connectors connect a CR-brick to an IX-brick or PX-brick. These are connected with a NUMAlink cable at 1.2 GB/s in each direction.
- **Console connector.** This is an RS-232 serial port (console and diagnostic port) that connects the L1 controller within a CR-brick to a system console. The CR-brick L1 controller communicates with a system console through the console connector, which provides a serial bus that uses the standard RS-232 protocol.
- **L1 host port connector.** This universal serial bus (USB) connector connects to the L1 controller.
- **Processor status LEDs.** The 32 processor status LEDs (8 for each processor) are used by SGI service engineers for diagnostic purposes.
- **Link and XIO connector LEDs.** Each connector has an LED that illuminates yellow and an LED that illuminates green. One LED illuminates yellow to indicate that both the CR-brick and the brick to which it is connected are powered on. The other

LED illuminates green when the link between the CR-brick and the other brick is established.




Figure 5-4 Rear View of CR-brick

Internal Components

The node electronics are connected internally to two NL4 router ASICs. Each of the two router ASICs connects to four Dnet connectors on the rear panel of the CR-brick (8 total Dnet connectors). Two XIO ports from the node boards (node 1 and node 3 are brought out to the rear of the brick). Each node board contains two processors and 4, 8, or 12 memory DIMMs. The processors, memory, I/O, and network connections on each node board are interconnected via a SHub ASIC. This ASIC acts as a crossbar between the processor interface, local memory interface, network interface, and I/O interfaces.

Note: The internal components of the CR-brick can be serviced only by trained SGI technicians.

Processors

Eight 64-bit Intel Itanium 2 processors with secondary cache are mounted on half panel printed circuit boards (referred to as node boards). To reduce memory latency, each processor can access three levels of on-chip cache memory. L3 cache sizes and types vary with processor speed.

Memory DIMMs

The Altix 3700 Bx2 systems use commodity off-the-shelf memory DIMMs. The DIMM card is a JEDEC standard 184-pin card.

Pinouts

See “I/O Port Specifications” on page 173 for pinout specifications for the non-proprietary connectors on the CR-brick.

Product Options

See Table 5-1 for configurable items in the CR-brick.

Table 5-1 Configurable Items of CR-brick

Configurable Item	Options
Memory size	4, 8, or 12 DIMMs per node
Processors	0 or 8 per CR-brick

Note: A CR-brick with 0 processors on the internal node boards is generally referred to as a “memory only” brick or M-brick.

Important Notes

The CR-brick has the following restrictions:

- All processors within the CR-brick must be the same frequency; however, CR-bricks within a partition or system can have different processor speeds.
- All processor revisions must be the same within a processor node.
- The processor revisions of CPUs between processor nodes can be no greater than n+1.
- All memory DIMMs within a memory bank (group) must be the same speed, capacity, and use the same chip technology.
- Different logical banks within a CR-brick can have different DIMM capacities and chip technologies.
- Memory DIMMs must be added in groups of four DIMMs per node board.

Technical Specifications

Table 5-2 lists the technical specifications of the CR-brick.

Table 5-2 CR-brick Technical Specifications

Characteristic	Specification
Height	6.67 in. (16.9 cm)
Width	19 in. (49.26 cm)
Depth	29 in. (73.66 cm)
Weight	68 lb (30.9 kg) maximum
Input power	48 VDC (~1300 W)

IX-brick

The IX-brick is an I/O expansion subsystem that connects I/O devices to your system by using either the PCI or PCI-X protocol. The PCI-X protocol enables I/O devices to operate at clock speeds of up to 133 MHz, or 1 Gbit/s. This protocol also enables I/O devices to operate more efficiently, thereby providing a higher sustained bandwidth at any clock frequency. By supporting this protocol, the IX-brick addresses the need for increased bandwidth of PCI devices.

The IX-brick also has the components (disk drives and DVD-ROM) that are required to run your operating system or install other software applications or upgrades.

This chapter describes the function and physical components of the IX-brick, and it describes how to install and replace PCI cards and disk drives. Specifically, it includes the following information:

- “Product Overview” on page 76
- “External Components” on page 79
- “PCI and PCI-X Card Configuration Guidelines” on page 83
- “Technical Specifications” on page 85

Note: Throughout this chapter, the term “PCI card” refers to cards with PCI or PCI-X capabilities. When necessary, distinctions between the cards are noted.

Note: For information about installing or replacing a PCI card, see “Adding or Replacing a PCI or PCI-X Card” on page 136.

Product Overview

The 4U-high IX-brick, shown in Figure 6-1, provides 12 PCI-X slots that support up to 12 PCI or PCI-X cards. The 12 slots are configured as six 2-slot buses.

Note: You can install PCI cards in 11 of the 12 PCI slots. One PCI-X slot (the leftmost slot) is reserved for a base I/O PCI card. This card is required for the base I/O functions, including system disk(s) and DVDROM, Ethernet interface, and serial connections.




Figure 6-1 Front View of the IX-brick

The IX-brick offers the following features:

- The IX-brick supports both PCI and PCI-X cards.
- The IX-brick can have up to four serial ports supported by the base I/O card. Additional optional serial ports reside on a daughtercard that you can purchase as an optional component.

Note: The serial-port daughtercard requires the space of one PCI-X slot, but it does not connect to the PCI-X bus.

Three PIC (PCI interface chip) ASICs are key components of the IX-brick architecture. These ASICs support two 1200- or 800-MB/s Xtown2 XIO ports and six PCI-X buses (see Figure 6-2). Each bus has two card slots in which you can install PCI cards. (Slot 1 of bus 1, however, seats the base I/O card only.) For best base I/O performance, leave the PCI slot (bus 1 slot 2) next to the base I/O card empty, or install only a 66 MHz PCI card.

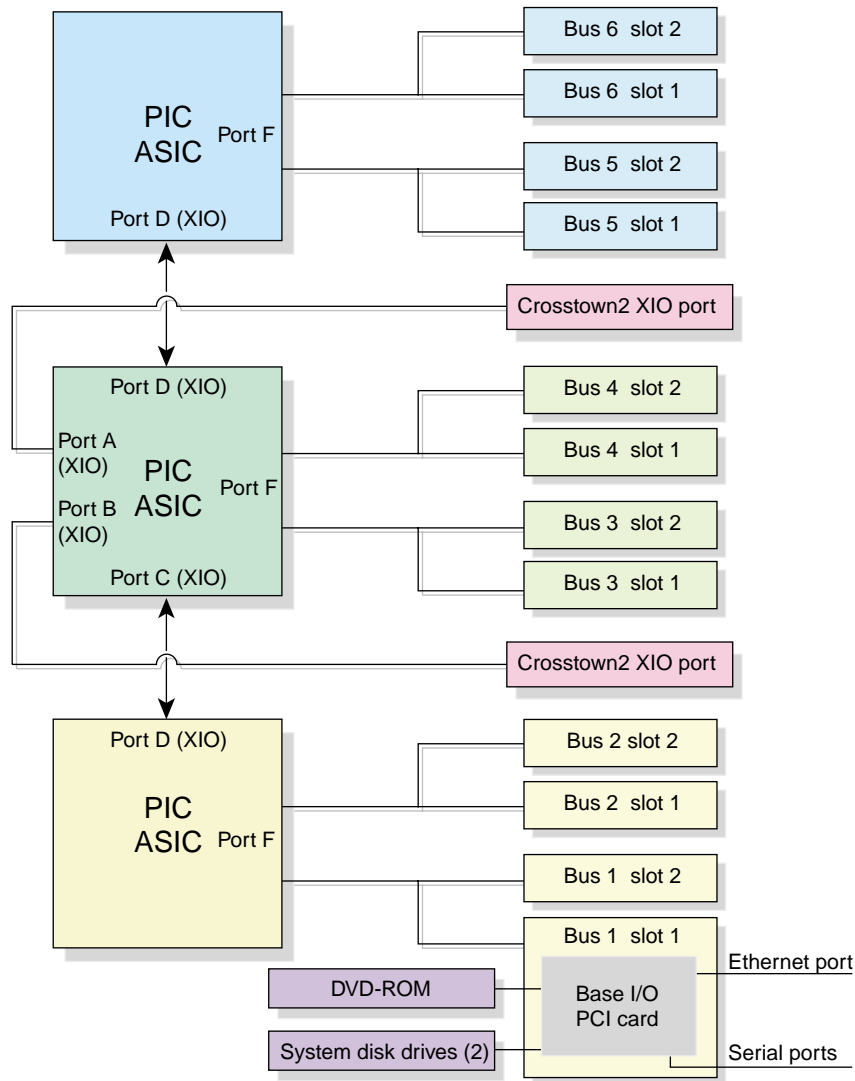


Figure 6-2 IX-brick Block Diagram

Also important to the IX-brick architecture is the base I/O PCI card. This card contains logic that controls the DVD-ROM and internal system disk drives, and it provides the following connectors (see Figure 6-4):

- **Internal system disk connector** that connects to up to two internal disk drives.
- **Ethernet port** (10/100/1000 Mbits). This autonegotiating 10BaseT/100BaseT/1000BaseT twisted-pair Ethernet port connects the system to an Ethernet network.
- **External SCSI connector.** This 68-pin VHDCI external SCSI port enables you to connect your system to an external SCSI device.

For an updated list of supported SCSI devices, see the SGI Supportfolio at <http://support.sgi.com>.

- **Real-time interrupt input and output.** The RT interrupt input and RT interrupt output functionality of the IO9 PCI card is not supported under SGI Linux + ProPack. These connectors are not usable with this system.

External Components

This section describes the external components that are located on the front and rear panels of the IX-brick.

Front Panel Components

The IX-brick has the following front panel components (see Figure 6-3):

- **Two system disk drives.** These customer-removable, sled-mounted disk drives are used to house your operating system and other application software. (See “Installing or Replacing a Disk Drive in the IX-brick” on page 144) for instructions for installing or removing the SCSI disk drives.)
- **DVD-ROM device.** This device loads software onto your IX-brick. (It is used for text reading only in CD-ROM mode.)
- **L1 controller display.** This liquid crystal display (LCD) displays status and error messages that the L1 controller generates.

Note: For more information about the L1 controller, see the *SGI L1 and L2 Controller Software User's Guide*.

- **On/Off switch with LED.** Press this button to turn on the internal components of the IX-brick. Alternatively, you can turn on the internal components at a system console. When the internal components are on, the LED illuminates green.
- **LEDs:**
 - **On/Off switch LED.** This green LED illuminates when the internal components of the IX-brick are on and turns off when they are off.
 - **Service required LED.** This LED illuminates yellow to indicate that a component is broken or is not operating properly (for example, if a fan is off), but the IX-brick is still operating.
 - **Failure LED.** This LED illuminates red to indicate that a system failure has occurred and the IX-brick is not operating.
- **Fans.** Three hot-swappable fans provide the required cooling for your IX-brick.



Warning: To prevent personal injury, or damage to the IX-brick, the hot-swappable fans can be installed only by a trained SGI system support engineer (SSE).




Figure 6-3 Front Panel of IX-brick

Rear Panel Components

The IX-brick has the following rear panel components (see Figure 6-4):

- **Power switch.** Moving the power switch to the **1** position powers on the L1 controller of the IX-brick, and moving it to the **0** position powers off the L1 controller. The 12-VDC LED illuminates green when the L1 controller is powered on.
- **12-VDC LED.** The 12-VDC LED illuminates green when the L1 controller is powered on and operating. The L1 controller is powered on by turning on the power switch.
- **48-VDC LED.** The power switch must be in the ON (**1**) position for this LED to light. The 48-VDC LED illuminates green when the rest of the IX-brick internal components are powered on and operating. The rest of the internal components are powered on by pressing the On/Off switch on the front panel of the brick, or via the L2 controller display or system console.
- **PWR (power) connector.** This connector connects to a power bay, which provides power to the IX-brick.
- **PCI-X slots.** These slots seat the PCI cards. See “Adding or Replacing a PCI or PCI-X Card” on page 136. The card slots are numbered bus 1 through bus 6. Each bus has two slots, labeled 1 and 2, as shown in Figure 6-4.
- **PCI-X slot LEDs.** Each PCI-X slot has the following LEDs:
 - **PWR (power) LED.** This LED illuminates green when the PCI card is installed securely and is getting power.
 - **Fault LED.** This LED illuminates yellow when a fault occurs with the PCI card.
- **XIO 10 and XIO 11 connectors.** Each connector can connect the IX-brick to one C-brick. As an option, the second connector can be connected to another C-brick to create a dual-ported IX-brick, which would provide greater bandwidth.
- **XIO 10 and XIO 11 connector LEDs.** Each XIO connector has two LEDs, as follows:
 - The yellow LED illuminates to indicate that both the IX-brick and the brick to which it is connected are powered on.
 - The green LED illuminates when a link has been established between the IX-brick and the brick to which it connects.
- **Two DB-9 RS-232 serial port connectors.** These ports can be used as COM ports to connect to serial devices. They require the presence of the base I/O card.

The leftmost PCI-X slot (bus 1, slot 1) seats a base I/O card that has the following connectors:

- **External SCSI connector.** This 68-pin VHDCI external SCSI port enables you to connect your system to an external SCSI device.
For an updated list of supported SCSI devices, see the SGI Supportfolio at <http://support.sgi.com>.
- **Ethernet RJ45 connector.** This autonegotiating 10/100/1000BaseT Ethernet port connects the system to an Ethernet network.
- **RTO and RTI stereo jack connectors (Not used).** The RTO (output) and RTI (input) connectors are not supported under the current SGI Linux ProPack release.




Figure 6-4 Rear Panel of IX-brick

PCI and PCI-X Card Configuration Guidelines

The PCI-X slots are numbered bus 1 through bus 6. Each bus has two slots, labeled 1 and 2, as shown in Figure 6-5. Separate buses enable the IX-brick to run cards of different frequencies at the same time. For example, bus 1 can have one 133-MHz card, bus 2 can have two 66-MHz cards, bus 3 can have two 33-MHz cards, and so on.

The PCI-X bus supports 32-bit and 64-bit PCI or PCI-X cards at the same time.

For maximum bandwidth, PCI cards are distributed across all six buses when they are integrated at the factory.

This section contains the following information:

- “Important Installation Considerations” on page 84
- “Supported PCI Cards” on page 84
- “PCI Card Carrier” on page 85




Figure 6-5 Numbering of IX-brick PCI-X Slots

Important Installation Considerations

To maximize the operating efficiency of the PCI cards, consider the following configuration guidelines before you install the cards.

- You can place one or two PCI cards on one bus, or one or two PCI-X cards on one bus.
- You should avoid mixing cards that operate at different frequencies or in different modes. If you have two cards of different speeds on the same bus, both cards operate at the lower speed. If a PCI card and PCI-X card are on the same bus, both cards operate in PCI mode. Note the following examples:
 - When one 133-MHz PCI-X card resides on a bus, the card operates at 133 MHz in PCI-X mode.
 - When two 133-MHz PCI-X cards reside on a bus, the cards operate at 100 MHz in PCI-X mode.
 - When two 66-MHz PCI-X cards reside on a bus, the cards operate at 66 MHz in PCI-X mode.
 - When two 66-MHz PCI cards reside on a bus, the cards operate at 66 MHz in PCI mode.
 - When one 66-MHz PCI card and one 66 MHz PCI-X card reside on a bus, the cards operate at 66 MHz in PCI mode.
 - When two 33-MHz PCI cards reside on a bus, the cards operate at 33 MHz in PCI mode.
 - When one 66-MHz PCI card and one 33-MHz PCI card reside on the same bus, the cards operate at 33 MHz in PCI mode.

Note: When installing a PCI card, you need to shut down the operating system and power off the IX-brick before installing the card.

Supported PCI Cards

SGI supports various PCI cards. These cards can be purchased from SGI or another manufacturer. Ask your SGI sales representative for a current list of PCI cards that SGI supports.

PCI Card Carrier

Each PCI card is mounted on a carrier so that you can slide the cards into and out of the brick. This carrier supports most PCI cards; it can be adjusted to accommodate cards of different sizes. (To learn how to adjust a carrier, see “Adding or Replacing a PCI or PCI-X Card” on page 136.)

When the IX-brick is shipped, any card that was ordered is installed with a carrier, and any unoccupied slot is populated with an empty carrier. A carrier must be present in an unpopulated slot to maintain an even airflow through the brick and to protect against electromagnetic interference (EMI).

Technical Specifications

Table 6-1 lists the physical specifications of the IX-brick.

Table 6-1 Physical Specifications of the IX-brick

Characteristics	Specifications
Height	6.64 in. (168.65 mm)
Width	17.5 in. (444.5 mm)
Depth	27.74 in. (704.59 mm)
Weight	65 lb (29.5 kg)
Input Power	+48 VDC (250 watts)

Table 6-2 shows the port specifications of the IX-brick.

Table 6-2 Port Specifications of the IX-brick

Port	Quantity	Connector Type
Power entry	One	21-pin Foxcon
XIO	Two	Proprietary 100-pin dual row
The following external connectors are located on the base I/O card:		
SCSI connector	One external	68-pin VHDCI external
Ethernet	One	RJ-45
RT interrupt input and output	One input and one output - (not used)	Stereo jack (not supported under SGI ProPack for Linux)

PX-brick

The PX-brick is an optional PCI-X based I/O expansion subsystem that connects I/O devices to your system. The PX-brick supports both PCI and PCI-X protocol. PCI-X protocol enables I/O devices to operate at clock speeds of up to 133 MHz, or 1 Gbit/s. This new protocol also enables I/O devices to operate more efficiently, thereby providing a higher sustained bandwidth at any clock frequency.

This chapter describes the function and physical components of the PX-brick, and it provides guidelines for configuring PCI cards in the brick. Specifically, it includes the following information:

- “Product Overview” on page 88
- “External Components” on page 90
- “PCI and PCI-X Card Configuration Guidelines” on page 94
- “Technical Specifications” on page 96

Note: Throughout this chapter, the term “PCI card” refers to cards with PCI or PCI-X capabilities. When necessary, distinctions between the cards are noted.

For information about installing or replacing a PCI card, see “Adding or Replacing a PCI or PCI-X Card” on page 136.

Product Overview

The 4U-high PX-brick shown in Figure 7-1 provides 12 card slots to support up to 12 PCI or PCI-X cards. The 12 slots are configured as six 2-slot buses.

The PX-brick offers the following:

- The PX-brick supports both PCI cards and PCI-X cards.
- The PX-brick can connect an InfiniteReality graphics pipe to a CR-brick.




Figure 7-1 Front View of the PX-brick

Three PIC (PCI interface chip) ASICs are key components of the PX-brick architecture. The PIC ASICs support the following (see Figure 7-2):

- Two 1200- or 800-MB/s Xtown2 XIO ports. (You can select the MB/s setting with the L1 controller command `XIO`. For more information, see the *SGI L1 and L2 Controller Software User's Guide*.)
- Six PCI/PCI-X buses. Each bus has two card slots in which you can install PCI or PCI-X cards.

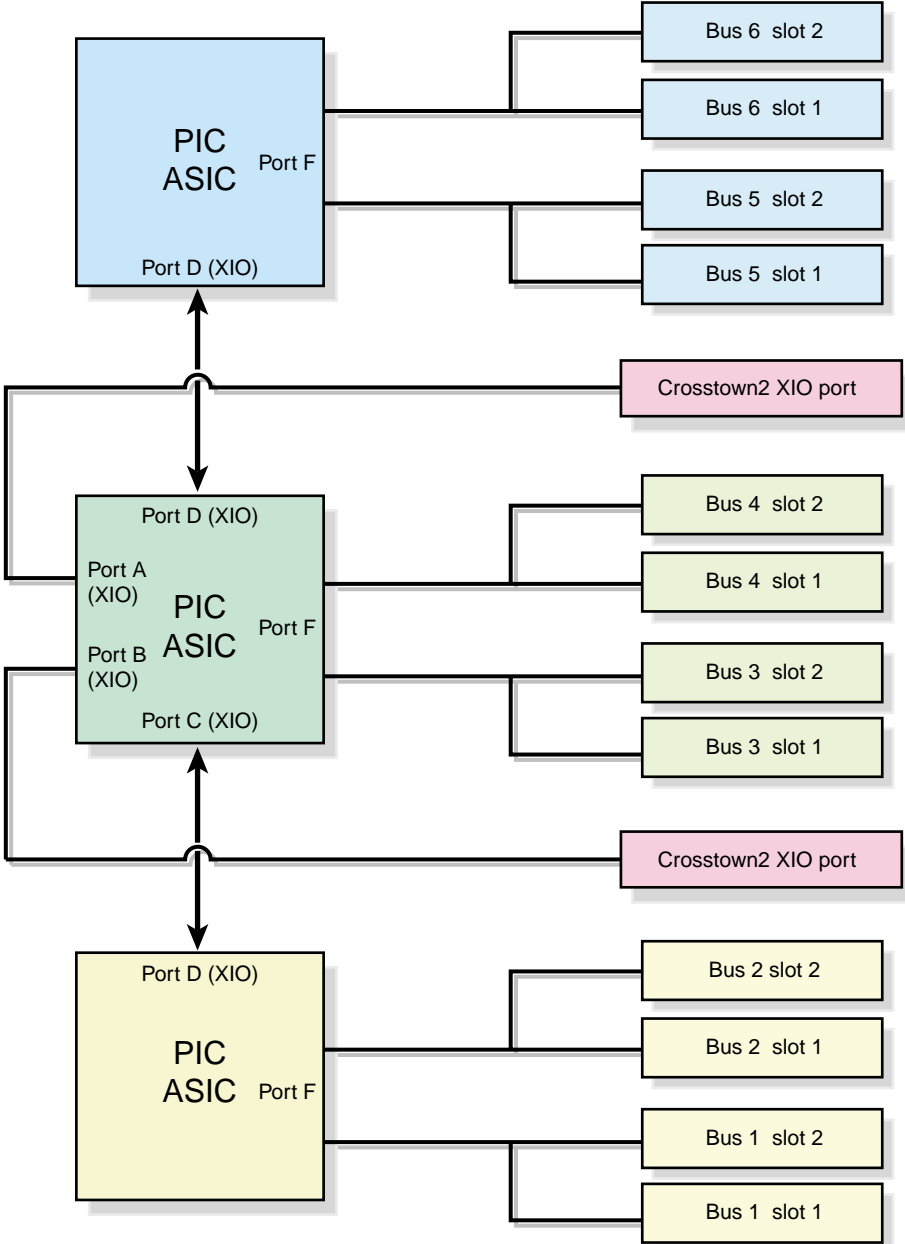


Figure 7-2 Block Diagram of PX-brick

External Components

This section describes the external components located on the front and rear panels of the PX-brick.

Front Panel Components

The following are the front panel components of the PX-brick (see Figure 7-3):

- **L1 controller and display.** The L1 controller generates PX-brick status and error messages that appear on the liquid crystal display (LCD).

Note: For more information about the L1 controller, see the *SGI L1 and L2 Controller Software User's Guide*.

- **On/Off switch with LED.** Press this button to turn on the internal components of the PX-brick. Alternatively, you can turn on the internal components at an L2 controller display or at a system console. When the internal components are on, the LED illuminates green.
- **LEDs:**
 - **On/Off switch LED.** This green LED illuminates when the internal components of the PX-brick are on and turns off when they are off.
 - **Service required LED.** This LED illuminates yellow to indicate that a component is broken or is not operating properly (for example, if a fan is off), but the PX-brick is still operating.
 - **Failure LED.** This LED illuminates red to indicate that a system failure has occurred and the PX-brick is not operating.
- **Fans.** Three hot-swappable fans provide the required cooling for your PX-brick.



Warning: To prevent personal injury, or damage to the PX-brick, the hot-swappable fans can be installed only by a trained SGI system support engineer (SSE).




Figure 7-3 Front Panel and Display on the PX-brick

Rear Panel Components

The PX-brick has the following rear panel components (see Figure 7-4):

- **Power switch.** Moving the power switch to the **1** position powers on the L1 controller of the PX-brick, and moving it to the **0** position powers off the L1 controller. The 12-VDC LED illuminates green when the L1 controller is powered on.
- **12-VDC LED.** The 12-VDC LED illuminates green when the L1 controller is powered on and operating. The L1 controller is powered on by turning on the power switch.
- **48-VDC LED.** The power switch must be in the ON (**1**) position for this LED to illuminate. The 48-VDC LED illuminates green when the rest of the PX-brick internal components are powered on and operating. The rest of the internal components are powered on by pressing the On/Off switch on the front panel of the brick, or via the L2 controller display or system console.
- **PWR (power) connector.** This connector connects to a power bay, which provides power to the PX-brick.
- **PCI/PCI-X slots.** These slots support PCI or PCI-X cards. See “Adding or Replacing a PCI or PCI-X Card” on page 136. The card slots are numbered bus 1 through bus 6. Each bus has two slots, labeled 1 and 2, as shown in Figure 5-4.
- **PCI-X Slot LEDs.** Each slot has the following LEDs:
 - **PWR (power) LED.** This LED illuminates green when the PCI card carrier is installed securely and is getting power.
 - **Fault LED.** This LED illuminates yellow when a fault occurs with the PCI card.
- **XIO 10 and XIO 11 connectors.** Each connector can connect the PX-brick to one compute brick (CR-brick). As an option, the second connector can be connected to another compute brick to create a dual-ported PX-brick, which would provide greater bandwidth.
- **XIO 10 and XIO 11 connector LEDs.** Each connector has two LEDs, as follows:
 - The yellow LED illuminates to indicate that both the PX-brick and the compute brick to which it is connected are powered on.
 - The green LED illuminates when a link has been established between the PX-brick and the brick to which it connects.

Figure 7-4 shows the location of the rear panel components of the PX-brick.

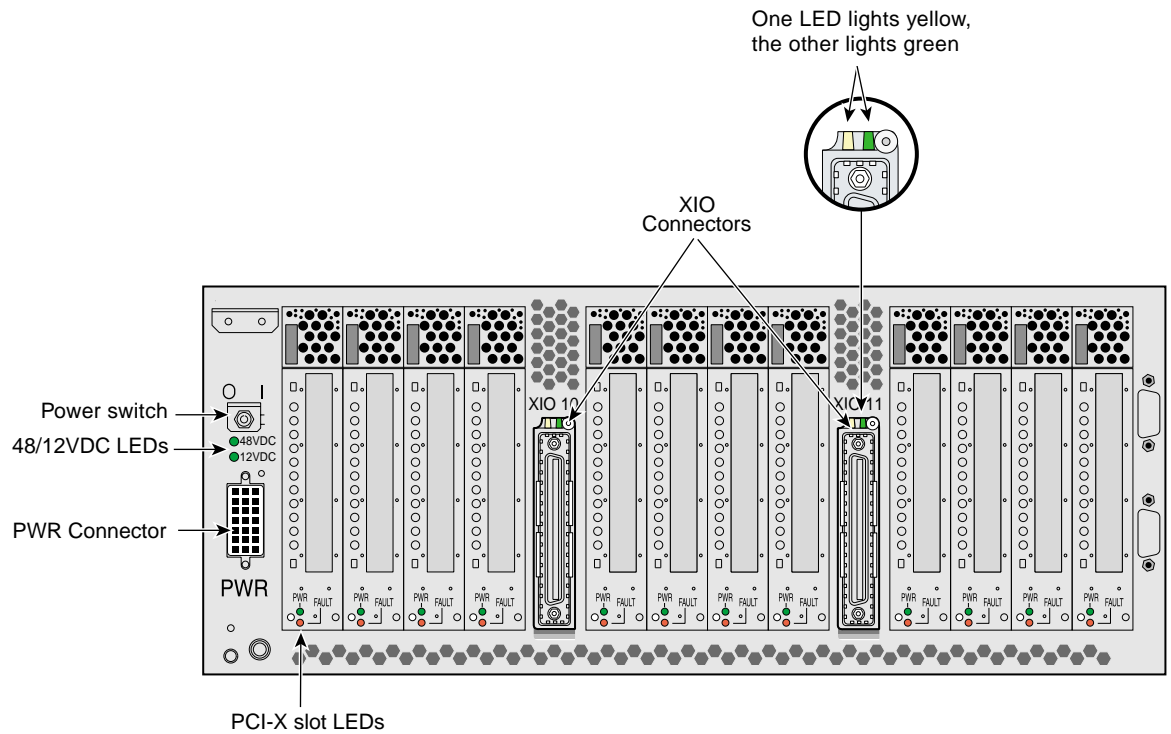


Figure 7-4 Rear Panel of the PX-brick

PCI and PCI-X Card Configuration Guidelines

The PCI and PCI-X card slots on the motherboard are numbered bus 1 through bus 6. Each bus has two slots, labeled 1 and 2, as shown in Figure 7-5. Separate buses enable the PX-brick to run cards of different frequencies at the same time. As a result, the same PX-brick can run 133-MHz, 100-MHz, 66-MHz, and 33-MHz cards at the same time.

The PCI/PCI-X bus supports 32-bit and 64-bit PCI or PCI-X cards at the same time.

For maximum bandwidth, PCI/PCI-X cards are distributed across all six buses when they are integrated at the factory.

This section contains the following information:

- “Important Installation Considerations” on page 95
- “Supported PCI and PCI-X Cards” on page 95
- “PCI Card Carrier” on page 96




Figure 7-5 Numbering of PCI and PCI-X Card Slots

Important Installation Considerations

To maximize the operating efficiency of the PCI cards, consider the following configuration guidelines before you install the cards.

- You can place one or two PCI cards on one bus, or one or two PCI-X cards on one bus.
- You should avoid mixing cards that operate at different frequencies or in different modes. If you have two cards of different speeds on the same bus, both cards operate at the lower speed. If a PCI card and a PCI-X card are on the same bus, both cards operate in PCI mode. Note the following examples:
 - When one 133-MHz PCI-X card resides on a bus, the card operates at 133 MHz in PCI-X mode.
 - When two 133-MHz PCI-X cards reside on a bus, the cards operate at 100 MHz in PCI-X mode.
 - When two 66-MHz PCI-X cards reside on a bus, the cards operate at 66 MHz in PCI-X mode.
 - When two 66-MHz PCI cards reside on a bus, the cards operate at 66 MHz in PCI mode.
 - When one 66-MHz PCI card and one 66 MHz PCI-X card reside on a bus; the cards operate at 66 MHz in PCI mode.
 - When two 33-MHz PCI cards reside on a bus, the cards operate at 33 MHz in PCI mode.
 - When one 66-MHz PCI card and one 33-MHz PCI card reside on the same bus, the cards operate at 33 MHz in PCI mode.

Note: When installing a PCI card, you need to shut down the operating system and power off the PX-brick before installing the card.

Supported PCI and PCI-X Cards

SGI supports various PCI and PCI-X cards. These cards can be purchased from SGI or another manufacturer. Ask your SGI sales representative for a current listing of PCI and PCI-X cards supported by SGI.

PCI Card Carrier

Each PCI card is mounted on a carrier so that you can slide the cards into and out of the brick. This carrier supports most PCI cards; it can be adjusted to accommodate cards of different sizes. (To learn how to adjust a carrier, see “Adding or Replacing a PCI or PCI-X Card” on page 136.)

When the PX-brick is shipped, any card that was ordered is installed with a carrier, and any unoccupied slot is populated with an empty carrier. A carrier must be present in an unpopulated slot to maintain an even airflow through the brick and to protect against electromagnetic interference (EMI).

Technical Specifications

Table 7-1 lists the physical specifications of the PX-brick.

Table 7-1 Physical Specifications of PX-brick

Characteristic	Specification
Height	6.64 in. (168.65 mm)
Width	17.5 in. (444.5 mm)
Depth	27.74 in. (704.59 mm)
Weight	60 lb (27.2 kg)
Input power	+48 VDC (225 watts)

Table 7-2 shows the port specifications of the PX-brick.

Table 7-2 Port Specifications of the PX-brick

Port	Quantity	Connector Type
Power entry	One	21-pin Foxcon
XIO	Two	Proprietary 100-pin dual row

R-brick

This chapter describes the function and physical components of the optional R-brick in the following sections:

- “Overview” on page 97
- “External Components” on page 100
- “Technical Specifications” on page 103

Overview

The R-brick (router brick) is an eight-port optional router that functions as a high-speed switch to route network packets from one CR-brick to another throughout the NUMAlink interconnect fabric. Figure 8-1 shows a front view of the R-brick.

The key component within the R-brick 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.




Figure 8-1 Front View of the R-Brick

The R-brick has the following features:

- Eight NUMAlink channels
- One USB port for system controller support
- One L1 controller and LCD display
- Two hot-pluggable cooling fans

Figure 8-2 shows a block diagram of the R-brick.

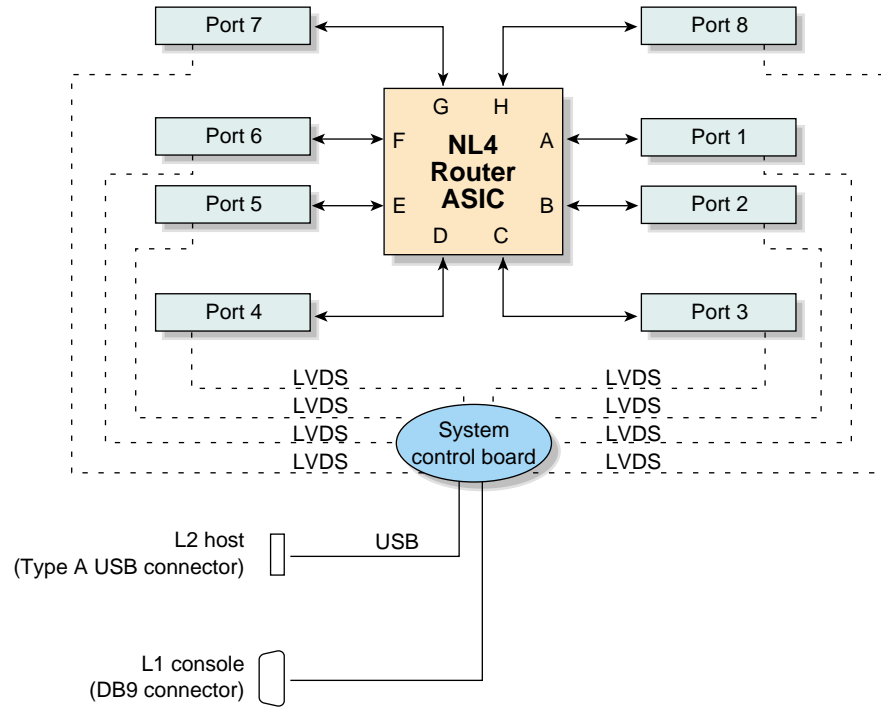


Figure 8-2 R-brick Block Diagram

External Components

This section describes the external components that are located on the front and rear panels of the R-brick.

Front Panel Components

The R-brick contains the following front panel items (as shown in Figure 8-1 on page 98):

- **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 R-brick internal components. You can also turn on the R-brick internal components at a system console.
- **Three LEDs:**
 - On/Off switch LED. This LED illuminates green when the R-brick internal components are on and turns off when they are off.
 - Service required LED. This LED illuminates orange to indicate that an item is broken or not operating properly (for example, a fan is off), but the R-brick is still operating.
 - Failure LED. This LED illuminates red to indicate that a system failure has occurred and the R-brick is down.
- **Fans.** Two hot-pluggable fans provide N+1 redundant cooling.

Rear Panel Components

The R-brick has the following rear panel items (see Figure 8-3):

- **Power switch.** Move the power switch to the I position to power on the L1 controller within the R-brick, and to the 0 position to power off the L1 controller. Powering on the L1 controller illuminates the 12-VDC LED green.
- **PWR (power) connector.** This connects the R-brick to the power bay, to provide 12-VDC and 48-VDC power to the R-brick.
- **48-VDC and 12-VDC LEDs.** The power switch must be in the ON (I) position for these LEDs to be on. The 12-VDC LED illuminates green when the L1 controller is powered on and operating, and the 48-VDC LED illuminates green when the rest of the R-brick internal components are powered on and operating. You power on the internal components by pressing the On/Off switch on the L1 controller panel.
- **Links R TO R connectors** (1, 6, 7, and 8, or A, F, G, and H). These link connectors connect the R-brick to other R-bricks in the network fabric.
- **Links R to R and CR to R connectors** (2, 3, 4, and 5 or B, C, D, and E). These link connectors normally connect the R-brick to CR-bricks; however, in larger systems they are also used to connect R-bricks to R-bricks.
- **L1 port connector.** This connects the internal USB hub of the R-brick 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 R-brick.
- **Link connector LEDs.** Each NUMAlink connector has two LEDs, as follows:
 - The yellow LED illuminates to indicate that both the R-brick and the brick to which it is connected are powered on.
 - The green LED illuminates when a link has been established between the R-brick and the brick to which it is connected.




Figure 8-3 Rear View of R-brick

Technical Specifications

Table 8-1 lists the technical specifications of the R-brick.

Table 8-1 R-brick 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)
Input power	48 VDC (~ 60 W)

Table 8-2 lists the specifications of the R-brick ports.

Table 8-2 R-brick Port Specifications

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

Power Bay

This chapter describes the function and physical components of the power bay in the following sections:

- “Overview” on page 105
- “External Components” on page 107
- “Distributed Power Supplies” on page 109
- “Technical Specifications” on page 111

Overview

The power bay is a 3U-high enclosure that holds a maximum of six hot-swappable distributed power supply modules (DPSs). The power bay monitors, controls, and supplies AC power to the DPSs. Although the power bay can hold a maximum of six DPSs, in this system the power bay contains three or five DPSs. The compute racks require power bays with five DPSs and the I/O racks require power bays with three DPSs.

Each DPS inputs single-phase AC voltage and outputs 950 W at 48 VDC and 42 W at 12 VDC. The outputs of the DPSs are bused together. For example, when the power bay contains three DPSs, the DPSs are bused together to provide approximately 2,760 W of 48 VDC power and 135 W of 12 VDC power.

Figure 9-1 shows front and rear views of the power bay.




Figure 9-1 Front and Rear Views of Power Bay

The power bay has the following features:

- 3U-high enclosure
- Six AC input connectors
- Eight DC output connectors
- Fault reset button

External Components

This section describes the external components located on the front and rear panels of the power bay.

Front Panel Components

The front of the power bay seats the distributed power supplies (see Figure 9-2). When the power bay contains three power supplies, the supplies reside in locations 4, 5, and 6.




Figure 9-2 Front View of Power Bay

Rear Panel Components

The power bay module has eight DC output connectors (see Figure 9-3). A power cord connects one output connector to a C-brick, R-brick, IX-brick, or PX-brick. This connection provides 12-VDC standby power, 48-VDC power, and monitoring signals.

Note: The L1 controller of a connecting module can monitor the status and enable the output ports of the power bay.

The power bay module also has six AC input connectors, one for each power supply location. For example, when the power bay contains three power supplies, connectors 4, 5, and 6 connect to the power distribution strip (PDS). When the power bay contains six power supplies, connectors 1, 2 and 3 also connect to the PDS. The reset button is used to reset the power bay when a fault condition occurs.

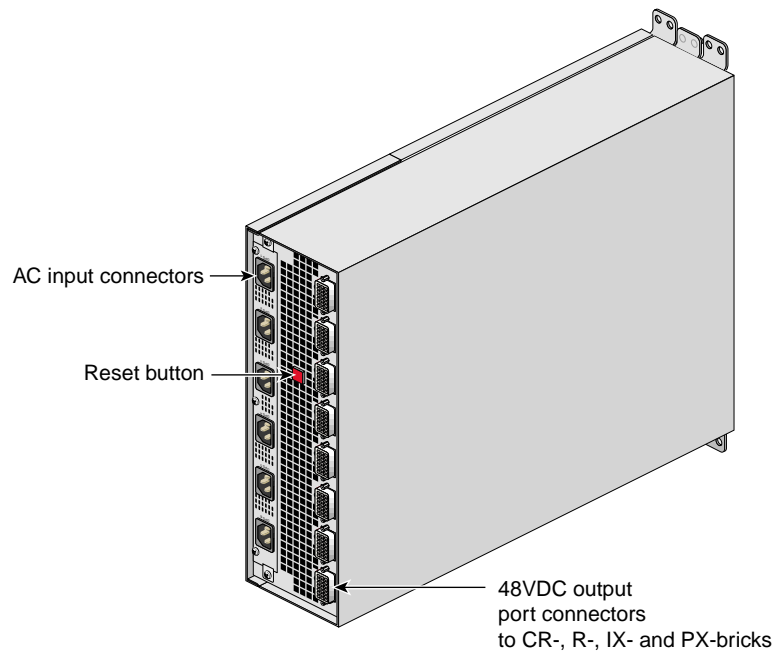


Figure 9-3 Rear View of Power Bay

Distributed Power Supplies

The distributed power supplies (DPSs) are air-cooled devices; each DPS has two fans that move air from the front of the rack to the rear of the rack. Figure 9-4 shows front and rear views of a DPS.

Each DPS has the following LEDs:

- Power
- Predictive fan fail (PFAIL)
- Power supply fail (FAIL)




Figure 9-4 Front and Rear Views of the DPS

The input and output power connectors shown in the rear view picture of the power supply interconnect with the power bay through a motherboard located inside the power bay.

Table 9-1 lists conditions of the power supplies and the corresponding states of the LEDs.

Table 9-1 Power Supply LED States

Power Supply Condition	LED States		
	Power (Green)	PFAIL (Amber)	FAIL (Amber)
AC voltage not applied to all power supplies	Off	Off	Off
AC voltage not applied to this power supply	Off	Off	On
AC voltage present; standby voltage on	Blinking	Off	Off
Power supply DC outputs on	On	Off	Off
Power supply failure	Off	Off	On
Current limit reached on 48-VDC output	On	Off	Blinking
Predictive failure	On	Blinking	Off

Each power supply also contains a serial ID EEPROM that identifies the model and serial number of the supply. The L1 controller of a connecting module reads this information.

Technical Specifications

Table 9-2 lists the technical specifications of the power bay.

Table 9-2 Power Bay Technical Specifications

Characteristic	Specification
Height	5.12 in. (130.04 mm)
Width	17.5 in. (444.5 mm)
Depth	23.87 in. (606.3 mm)
Weight (with five power supplies)	64.5 lb (29.26 kg)
Input voltage	220 VAC
Output voltage	12 VDC and 48 VDC

Table 9-3 lists the technical specifications of the distributed power supplies.

Table 9-3 Power Supply Technical Specifications

Characteristic	Specification
Height	5 in. (127 mm)
Width	2.8 in. (71.12 mm)
Depth	13 in. (330.2 mm)
Weight	7.5 lb (3.40 kg)
Input voltage	220 VAC
Output voltage	12 VDC and 48 VDC

SGI TP900 Storage Module

This chapter describes the function and physical components of the optional SGI TP900 storage module in the following sections:

- “Overview” on page 113
- “External Components” on page 114
- “Technical Specifications” on page 119
- “Product Options” on page 119

Overview

The SGI TP900 storage module is a rackmountable, 2U-high, 8-drive storage system that provides JBOD (“just a bunch of disks”) storage for the SGI Altix 3700 Bx2 system. The drive backplane of the TP900 connects the 8 drives on one SCSI bus. As an option, the storage system can also be configured on two SCSI buses (2 strings of 4 drives).

Note: For more information about the TP900 storage module, see the *SGI Total Performance 900 Storage System User’s Guide*.

The TP900 storage module has the following features:

- Supports 8 1.0-inch high, 3.5-in. disk drives.
- Anti-tamper locks on each disk drive to prevent accidental removal.
- Sled-mounted drives for easy removal and insertion.
- Redundant power and cooling.
- Single- or dual-channel operation.

External Components

This section describes the external components located on the front and rear panels of the TP900 storage module.

Front Panel Components

The front of the SGI TP900 chassis consists of 8 drive bays that contain either disk drive carrier modules or dummy carrier modules, which are discussed in the subsections that follow. The front of the TP900 is 4 drive bays wide by 2 drive bays high. The bays are numbered 1 and 2 from top to bottom, and 1 to 4 from left to right. See Figure 10-1.




Figure 10-1 Front View of TP900 Storage Module

The TP900 houses the following front panel components, which are discussed in the subsections that follow:

- Drive carrier module
- Dummy carrier module

The front of the chassis also contains three status LEDs that are described in Table 10-1.

Table 10-1 Status LEDs

LED	Status
System status	Green signifies that power is applied to the enclosure. Amber signifies that the ESI processor has failed.
Power status	Green signifies that the power supplies are functioning normally. Amber signifies a power supply failure.
Cooling status	Green signifies that all fans are functioning normally. Amber signifies a fan failure.

Disk Drive Carrier Module

The disk drive carrier module is a die-cast aluminum carrier that houses a single 1.0-in. high, 3.5-in. disk drive.

Each drive carrier has two LEDs: an upper green LED and a lower amber LED. In normal operation, the green illuminates and flickers as the drive operates. The amber LED illuminates when a fault occurs. See Figure 10-2 for the locations of LEDs.




Figure 10-2 Drive Carrier LED Indicators

This carrier has a handle that enables you to insert and remove the carrier from the drive bay. The handle also has an anti-tamper lock that locks the handle in place; therefore, the carrier cannot be removed from the drive bay. When the handle is locked, a red indicator is visible in the handle. When carrier is unlocked, a white indicator is visible.

Dummy Carrier Module

Dummy carrier modules are provided for installation in all unused drive bays. They are designed as integral drive module front caps with handles. They must be installed in unused drive bays to maintain a balanced airflow.

Rear Panel Components

The rear of the TP900 storage module contains 6 bays that house power supply, cooling, and SCSI I/O modules. See Figure 10-3. The rear of the TP900 storage module is 3 bays wide by 2 bays high. The bays are numbered 1 and 2 from top to bottom, and 1 to 3 from right to left.




Figure 10-3 Rear View of TP900 Storage Module

The TP900 rear bays house the following components, which are discussed in the subsections that follow:

- Power supply module
- Input/output (I/O) module
- Loopback module
- Cooling module
- Blank module
- Terminator plug

Power Supply Module

The SGI TP900 storage system requires one 350-W power supply module, which is mounted in the rear of the system (location 2/3). Power supply voltage operating ranges are nominally 115 V or 230 V and are selected automatically.

Optionally, a second power supply module (location 2/1) can be added to the TP900 system to provide n+1 redundant power. In this n+1 redundant configuration, the power supply modules operate together; if one power supply module fails, the other module supplies power until the faulty unit is replaced.

The power supply module contains two LEDs that are located on the front panel of the power supply. The green LED indicates power output. The amber LED indicates a power supply failure.

Input/Output (I/O) Module

The I/O module contains two VHDCI connectors and a switch to set the SCSI ID range. The SGI TP900 storage module supports one I/O module for a 1 x 8 configuration and 2 I/O modules for a 2 x 4 configuration. The 1 x 8 configuration creates a SCSI bus structure with 1 string of 8 drives. The 2 x 4 configuration creates a SCSI bus structure with 2 strings of 4 drives.

Note: When only one I/O module is present in the system, the other I/O module location contains a loopback module that has no external connectors.

Loopback module

The backplane of the TP900 consists of two 4-drive channels. A loopback module is used to connect the two internal drive channels together when the TP900 is configured as one channel with 8 drives.

Cooling Module

The cooling module consists of fans that pull warm air from a plenum behind the drive backplane and exhaust the warm air at the rear of the enclosure. The cooling modules connect to the backplane for power and status signal connections.

The TP900 system requires one cooling module, which is mounted in the rear of the system. Optionally, a second cooling module can be added to the system to provide redundant cooling. In this redundant configuration, the cooling modules operate together; if one module fails, the other module adequately cools the system until the faulty unit is replaced.

The front panel of the cooling module has two LEDs. The green LED indicates that the cooling module is functioning properly. The amber LED indicates a fan failure.

Blank Module

Blank modules are placed in all vacant power supply and cooling module bays to maintain proper airflow through the system.

Terminator Plug

Each I/O module has two 68-pin VHDCI connectors. A host bus adapter is connected to one of the VHDCI connectors and a terminator plug is placed in the remaining connector. Figure 10-4 shows the terminator plug.




Figure 10-4 SCSI Terminator Plug

Technical Specifications

Table 10-2 lists the technical specifications of the TP900 storage module.

Table 10-2 TP900 Storage Module Technical Specifications

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

Product Options

Table 10-3 lists the available configurable items for the TP900 storage module.

Table 10-3 Configurable Items

Configurable Item	Options
Disk drives	(1, minimum, 8 maximum)
I/O module	1 minimum, 2 maximum
Cooling module	1 minimum, 2 maximum
Power supply module	1 minimum, 2 maximum

D-brick2 Storage Module

This chapter describes the functional and physical components of the D-brick2 non-RAID mass storage option brick. It also explains how to power the D-brick2 on and off, and how to replace disk drive modules in the D-brick2. For troubleshooting and additional detailed information on the D-brick2, see *SGI Total Performance 9100 (2Gb TP9100) Storage System User's Guide (P/N 007-4522-00x)*, which is available online at <http://docs.sgi.com>.

Note: If you require RAID storage, ask your SGI sales representative about SGI RAID storage products, such as the SGI TP9400 and SGI TP9500 or TP9500S.

This chapter is divided into the following five major sections:

- “D-brick2 Overview” on page 122
- “D-brick2 External Components” on page 123
- “Powering the D-brick2 On or Off” on page 128
- “D-brick2 Technical and Environmental Specifications” on page 130

For instructions on removing or replacing the drive carrier modules, see “Installing or Replacing a D-brick2 Drive Carrier Module” on page 153.

D-brick2 Overview

The optional D-brick2 module is a high performance, large-scale non-RAID storage system for your SGI rackmounted system. Each enclosure contains a minimum of 2 and maximum of 16 disk drives, and the component modules that handle I/O, power and cooling, and operations. Optional RAID storage systems are available for your system through your SGI sales representative. Figure 11-1 shows front and rear views of the D-brick2.

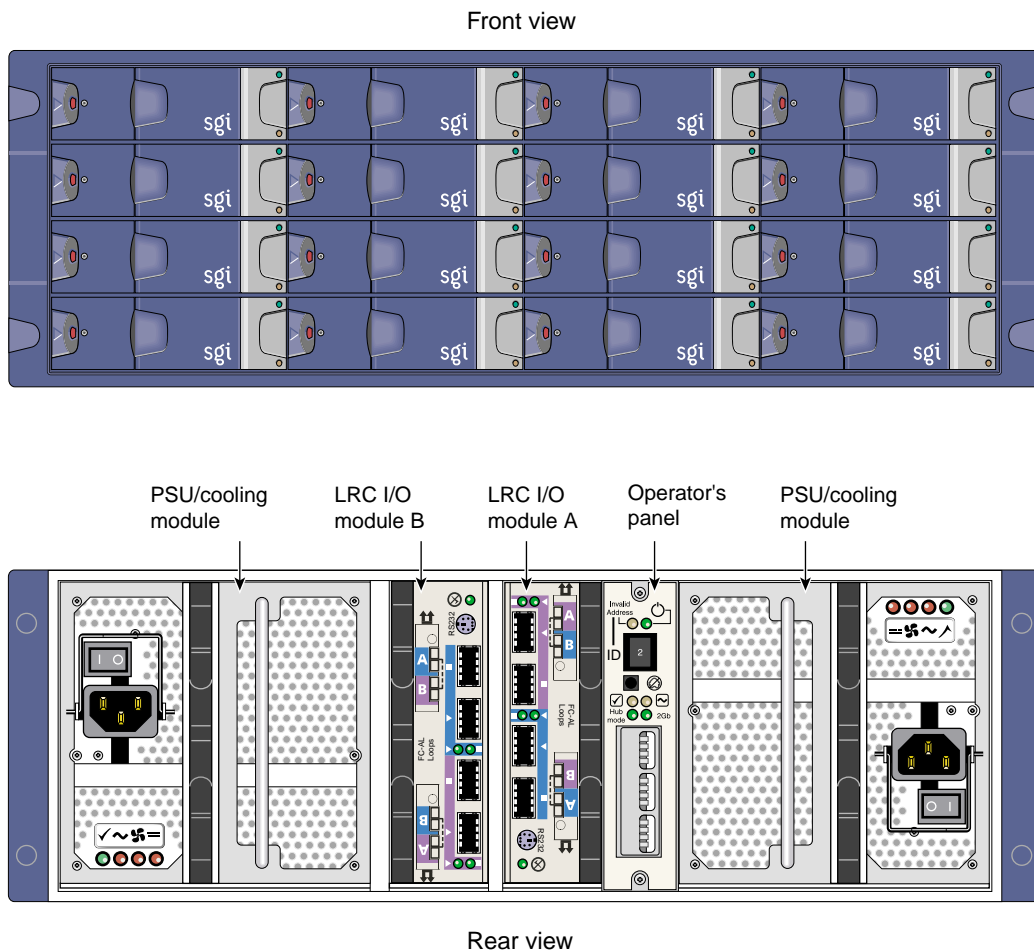


Figure 11-1 D-brick2 Front and Rear Views

Functional Description

The modular design of the D-brick2 expands easily to meet your mass storage needs. This storage system provides compact, high-capacity JBOD ("just a bunch of disks") storage for supported SGI systems. Each D-brick2 is connected to one or more Fibre Channel boards (host bus adapters, or HBAs) in the SGI host system, either separately or in a combination (loop).

D-brick2 Features

The optional D-brick2 storage system has the following features:

- Maximum configuration of up to 96 drives (six D-brick2 units)
- 1x16 (more storage) and 2x8 (more bandwidth) disk topologies in each brick
- Dual power feeds with dual power supplies
- Redundant cooling
- Non-disruptive component replacement
- Enclosure services interface (ESI) for SCSI enclosure services (SES)

D-brick2 External Components

This section discusses the external D-brick2 components that you can access, add, replace, or upgrade. These include the following:

- Drive carrier modules
- Dummy" drive modules
- Power supply and cooling modules
- Operator's panel
- Loop resiliency circuit (LRC) I/O modules

The first two items are accessed at the front of the module; the last three items are located in the rear of the D-brick2 unit.

Drive Carrier Modules and “Dummy” Modules

The disk drive carrier module consists of a hard disk drive mounted in a die-cast aluminum carrier (see Figure 11-2). The carrier protects the disk drive from radio frequency interference, electromagnetic induction, and physical damage. It also provides a means for thermal conduction.

Dummy drive carrier modules must be installed in all unused drive bays. They are designed as integral drive module front caps with handles and must be fitted to all unused drive bays to maintain a balanced airflow. Figure 11-2 shows examples.




Figure 11-2 D-brick2 Drive Carrier Module and “Dummy” Module

A key (Torx screwdriver) to lock or unlock the disk drives is provided with each D-brick2. A red indicator is visible in the center rectangular aperture in the handle if the anti-tamper lock is locked, and a black indicator is visible when the lock is unlocked (see Figure 11-3).




Figure 11-3 Anti-Tamper Locking on the D-brick2 Disk Drive

Power Supply/Cooling Modules

Two power supply/cooling modules (PSUs) are mounted in the rear of the D-brick2. These modules supply redundant cooling and power to the enclosure and are auto-ranging. The supplies should be plugged into your SGI rack's power distribution strip (PDS). Four LEDs mounted on the rear panel of the D-brick2's PSU/cooling module indicate the status of the power supply and the fans (see Figure 11-4 on page 126).

Note: If a power supply fails, do not remove it from the enclosure until you have a replacement power supply. The cooling fans in the supply will continue to operate.



Caution: You must finish replacing a PSU/cooling module within 10 minutes after removing one. Otherwise, thermal overload of the D-brick2 can result.




Figure 11-4 D-brick2 Operators Panel and Power Supply/Cooling Modules

Operator's Panel

The operators panel (“ops panel”) contains an enclosure services processor that monitors and controls D-brick2 functions. The ops panel contains LEDs that show the status of all modules, an audible alarm that indicates a fault state is present, a push-button alarm mute switch, and a thumbwheel enclosure ID address range selector switch. When the D-brick2 is powered on, the audible alarm sounds for one second, and the power-on LED illuminates.

Figure 11-4 identifies all controls and indicators on the ops panel. Note that the operator’s panel configuration switches slide left for “on” and right for “off.” For detailed descriptions of the LEDs and configuration switch information, see Chapter 5 in the *SGI Total Performance 9100 (2Gb TP9100) Storage System User’s Guide (P/N 007-4522-00x)*.

Loop Resiliency Circuit (LRC) Modules

The D-brick2 uses a Fibre Channel arbitrated loop (FC-AL) to interface with the host computer system. The FC-AL backplane incorporates two independent loops formed by port bypass circuits within the LRC I/O modules. Figure 11-5 shows an LRC module.

Processors housed on the LRC modules provide enclosure management and interface to devices on the backplane, a PSU/cooling module, and an ops panel to monitor internal functions. These processors operate in a master/slave configuration to allow failover. See the *SGI Total Performance 9100 (2Gb TP9100) Storage System User's Guide (P/N 007-4522-00x)* for details.

Note: The JBOD LRC I/O module can address up to 96 disk drives; a maximum of six enclosures can be cabled together.

The enclosure may be configured with either one or two LRC I/O modules. If only one module is installed, an I/O blank module must be installed in the unused bay.




Figure 11-5 D-brick2 Loop Resiliency Circuit (LRC) Module

Powering the D-brick2 On or Off

This section explains how to power the D-brick2 on and off.

Powering On the D-brick2

Before powering on the D-brick2, confirm that the following is true:

- Drives are seated in the correct bays and blank plates are fitted in any empty bays.
- Ambient temperature is within the specified range of 10 °C to 40 °C (50 °F to 104 °F).
- The system power distribution unit (PDU) and power distribution strip (PDS) are on.

To power on the D-brick2, follow these steps:

1. Connect an AC power cord to each PSU/cooling module.
2. Connect the AC power cords to the system PDS.
3. Turn the power switch on each PSU/cooling module to the “on” position (I = on, O = off). See Figure 11-4 on page 126 for switch locations.

The green “PSU good” LED illuminates. Also, the “power on” LED on the ESI/ops panel (see Figure 11-4 on page 126) of each module turns green when AC power is present.

If the “power on” LED on the ESI/ops panel does not illuminate, or if the amber “system/ESI fault” LED illuminates, verify that you followed all steps. For troubleshooting tips, see the *SGI Total Performance 9100 (2Gb TP9100) Storage System User’s Guide* or contact your service provider.

Powering Off the D-brick2

Before powering off the D-brick2, confirm that the following has occurred:

- All system users have been notified and are logged off.
- Disk data has been backed up as appropriate.

To power off a D-brick2 enclosure, follow these steps:

1. Move the power switch on the rear of each PSU/cooling module to the “off” position (position **I** = on, **O** = off). See Figure 11-4 on page 126 for switch locations.
2. Unplug the power cable from the PSU/cooling module(s) as appropriate. For example, you would follow this step if you are replacing an individual module.
3. If you are shutting down multiple bricks in the system, you may want to switch the PDS or PDU breaker switch to the “off” position.

The LEDs on the back of the unit should turn dark a few seconds after you power off the PSU/cooling module.

D-brick2 Technical and Environmental Specifications

The rackmounted D-brick2 disk enclosure is 13.4 cm (5.3 in.) high, 50 cm (19.7 in.) deep, and 44.6 cm (17.5 in.) wide. Component weights, power requirements, and environmental information for the D-brick2 are provided in the following tables.

Note: You should always follow system-level technical, operational, and environmental specifications, regardless of the tolerances of individual system bricks.

Table 11-1 shows the weights of the brick and the various D-brick2 components.

Table 11-1 D-brick2 Weight Information

Component	Weight
D-brick2 enclosure, fully populated	32.3 kg (71 lb)
D-brick2 enclosure, empty	17.9 kg (39.4 lb)
Power supply/cooling module	3.6 kg (7.9 lb)
Disk carrier module with drive	0.88 kg (1.9 lb)
LRC I/O module	1.2 kg (2.6 lb)

Table 11-2 shows the power requirements and specifications of the D-brick2.]

Table 11-2 D-brick2 Enclosure Power Specifications

Specification	Value
Voltage range for rack	200-240 VAC
Voltage range selection	Automatic
Frequency	50-60 Hz
Power factor	>0.98
Maximum power consumption	700 VA
Typical power consumption	400 VA or less
Inrush current (25 °C [77 °F] cold-start 1 PSU)	100 A maximum peak for 4 ms, 25 A thereafter at maximum voltage
Harmonics	Meets EN61000-3-2
Power cord:	
Cord type	SV or SVT, 18 WG minimum, 3 conductor
Plug	250 V, 10 A
Socket	IEC 320 C-14, 250 V, 15 A

Table 11-3 provides temperature and humidity requirements for D-brick2 modules.]

Table 11-3 D-brick2 Ambient Temperature and Humidity Requirements

Factor	Temperature	Relative Humidity	Maximum Wet Bulb
Operating temperature	5 °C to 40 °C (41 °F to 104 °F)	20% to 80% noncondensing	23 °C (73 °F)
Non-operating temperature	0 °C to 50 °C (32 °F to 122 °F)	8% to 80% noncondensing	27 °C (80 °F)
Storage temperature	1 °C to 60 °C (34 °F to 140 °F)	8% to 80% noncondensing	29 °C (84 °F)
Shipping temperature	-40 °C to +60 °C (-40 °F to 140 °F)	5% to 100% nonprecipitating	29 °C (84 °F)

Table 11-4 provides additional environmental specifications for D-brick2 units.

Table 11-4 Additional Environmental Requirements for D-brick2

Environmental Factor	Requirement
Altitude, operating	0 to 3047 m (0 to 10,000 ft)
Altitude, non-operating	-305 to 12,192 m (-1000 to 40,000 ft)
Shock, operating	Vertical axis 5 g peak 1/2 sine, 10 ms
Shock, non-operating	30 g 10 ms 1/2 sine
Vibration, operating	0.21 grms 5-500 Hz random
Vibration, non-operating	1.04 grms 2-200 Hz random
Acoustics	Less than 6.0 B LwA operating at 20 °C
Safety and approvals	CE, UL, cUL
EMC	EN55022 (CISPR22-A), EN55024 (CISPR24), FCC-A

Maintenance and Upgrade Procedures

This chapter provides information about installing or removing components from your SGI system, as follows:

- “Maintenance Precautions and Procedures” on page 133
- “Adding or Replacing a PCI or PCI-X Card” on page 136
- “Installing or Replacing a Disk Drive in the IX-brick” on page 144
- “Replacing a TP900 Drive Carrier Module” on page 149
- “Installing or Replacing a D-brick2 Drive Carrier Module” on page 153

Maintenance Precautions and Procedures

This section describes how to open the system for maintenance and upgrade, protect the components from static damage, and return the system to operation. The following topics are covered:

- “Preparing the System for Maintenance or Upgrade” on page 134
- “Returning the System to Operation” on page 134
- “Installing or Removing Internal Parts” on page 135

Preparing the System for Maintenance or Upgrade

To prepare the system for maintenance, follow these steps:

1. If you are logged on to the system, log out.
At your system console, switch to the L1 mode by entering the following command:
`$> Ctrl+T`
From the L1 prompt (L1>), power off the system with the following command:
`L1> power down`
2. Locate the PDU(s) in the rear of the rack and turn off the circuit breaker switches on each PDU.

Returning the System to Operation

When you finish installing or removing components, return the system to operation as follows:

1. Turn each of the circuit breaker switches to the “on” position.
2. At your system console, enter the following command:
`L1> power up`
3. Verify that the LEDs turn on and illuminate green and that your controllers display that the system is powered on for each segment of the procedure, which indicates that the power-on procedure is proceeding properly.

If your system does not boot correctly, see “Troubleshooting Chart” in Chapter 13, for troubleshooting procedures.

Installing or Removing Internal Parts



Caution: The components inside the system are extremely sensitive to static electricity. Always wear a wrist strap when you work with parts inside your system.

To use the wrist strap, follow these steps:

1. Unroll the first two folds of the band.
2. Wrap the exposed adhesive side firmly around your wrist, unroll the rest of the band, and then peel the liner from the copper foil at the opposite end.
3. Attach the copper foil to an exposed electrical ground, such as a metal part of the chassis.



Caution: Do not attempt to install or remove components that are not listed in Table 12-1. Components not listed must be installed or removed by a qualified SGI field engineer.

Table 12-1 lists the customer-replaceable components and the page on which you can find the instructions for installing or removing the component.

Table 12-1 Customer-replaceable Components and Maintenance Procedures

Component	Procedure
PCI and PCI-X cards	"Adding or Replacing a PCI or PCI-X Card" on page 136
IX-brick disk drives	"Installing or Replacing a Disk Drive in the IX-brick" on page 144
TP900 drive carrier modules	"Replacing a TP900 Drive Carrier Module" on page 149
D-brick2 disk drives	"Installing or Replacing a D-brick2 Drive Carrier Module" on page 153

Adding or Replacing a PCI or PCI-X Card



Warning: Before installing, operating, or servicing any part of this product, read the “Safety Information” on page 177.

This section provides instructions for adding or replacing a PCI or PCI-X card in the IX-brick or PX-brick. To maximize the operating efficiency of your cards, be sure to read all the introductory matter in the “PCI and PCI-X Card Configuration Guidelines” on page 94 before beginning the installation.



Caution: To protect the PCI cards from ESD damage, SGI recommends that you use a grounding wrist strap while installing a PCI card.

To add or replace a PCI card, follow these steps:

1. Shut down the operating system. (See your software guide if you need instructions to do this.)
2. Power off the I/O brick by following the power-off instructions in the “Powering Off the System” on page 8.
3. To extract it from the slot, pull the handle on the selected card carrier straight out, as shown in Figure 12-1.




Figure 12-1 Removing a Card Carrier

4. If you are replacing a card, gently remove the existing card from the carrier.

If you are adding a card, extract the carrier metal filler plate by pushing down on it, as shown in Figure 12-2. This filler plate covers the area where your card connectors will protrude out to the rear panel of the I/O brick.




Figure 12-2 Extracting the Carrier Metal Filler Plate

Note: If you are adding or replacing a half-height PCI card, skip step 5 and proceed to step 6.

5. If the full-height card that you want to install does not fit in the carrier, follow these steps:
 - a. Loosen the three screws on the guide bar with a Phillips screwdriver (if your carrier uses T8 Torx screws, you will need a Torx screw driver).
 - b. Adjust the carrier guide bar (also known as the alignment rail) up or down as necessary (see Figure 12-3).




Figure 12-3 Adjusting the Carrier Guide Bar

- c. Mount the card in the carrier, as shown in Figure 12-4, so that the following occurs:
 - The card connectors protrude through the bulkhead where the metal filler plate was located.
 - The edge connectors fit between the bottom guides of the carrier.
 - The top of the card fits under the clips on the guide bar.
- d. Tighten the three screws on the guide bar. Then skip step 6 and proceed to step 7.




Figure 12-4 Mounting Card in Carrier

6. To install a half-height PCI card in a carrier, follow these steps:
 - a. Remove the two screws from the carrier guide bar, as shown in Figure 12-5.

Note: Take care not to lose the threaded inserts in the carrier guide bar.

- b. Loosen the third screw of the carrier guide bar.
 - c. Place the half-height PCI card in the carrier so that the connector edge of the card rests against the lower guide tab of the carrier (see Figure 12-5).




Figure 12-5 Mounting Half-height PCI Card into Carrier

- d. Move the carrier guide bar (adjustment rail) so that it holds the card firmly in place (see Figure 12-6).




Figure 12-6 Moving Carrier Guide Bar to Secure Half-height PCI Card

- e. Place the PCI card bracket so that the screw holes of the bracket align with the empty screw holes of the carrier guide bar, as shown in Figure 12-7.
- f. Secure the bracket to the carrier guide bar with two T8 Torx screws.
- g. Tighten the screw that you loosened in step 6b.




Figure 12-7 Installing the Bracket to Secure a Half-height PCI Card

7. Insert the carrier-mounted PCI card into the vacant slot, using the slot guide, as shown in Figure 12-8.




Figure 12-8 Installing a Card in a Slot

8. Push horizontally on the carrier seating bar, as shown in Figure 12-9, to seat the card securely in the slot.




Figure 12-9 Seating the Card in the Slot

9. Power on the I/O brick by following the power-on instructions in the “Powering On the System” on page 4.

10. Boot your operating system software. (See your software operation guide if you need instructions to boot your operating system.)
11. Run the `hinv` hardware inventory command to verify the installation. This command lists hardware that the operating system discovered during the boot operation.

Installing or Replacing a Disk Drive in the IX-brick

The IX-brick has two system disk drive bays as shown in Figure 12-10. This section describes how to install or remove the drives.




Figure 12-10 Location of SCSI Disk Drive Bays

Installing a System Disk Drive

Before you install a disk drive, you must power off the IX-brick according to the instructions in the “Powering Off the System” on page 8.



Caution: To protect the system from ESD damage, SGI recommends that you use a grounding wrist strap while installing the disk drive.

To install a disk drive, see Figure 12-11 and follow these steps:

1. Remove the drive bay cover to access the two disk drive bays (see panel 1 in Figure 12-11).

Note: If you are installing only one disk drive, install it in the rightmost bay.

2. Position the drive assembly so that it engages the bay guide rails, and then gently push the drive into the bay. Do not use the locking handle to push the drive into the bay.
3. Swing the locking handle towards the chassis until the locking handle engages the latch.
4. Snap on the drive bay cover.

After you finish installing the disk drive, power on the IX-brick according to the instructions in the “Powering On the System” on page 4.




Figure 12-11 Installing a Disk Drive

Removing a System Disk Drive

Before you remove a disk drive, you must power off the IX-brick according to the instructions in the “Powering Off the System” on page 8.

To remove a disk drive, follow these steps (see Figure 12-12):

1. Remove the drive bay cover to access the two disk drive bays, as shown in panel 1 of Figure 12-12.
2. Remove the disk drive by depressing its handle lock with your thumb and pulling the handle away from the chassis until the handle disengages the disk drive connector from the backplane connector.
3. Carefully slide the disk drive out of the bay and place it on an ESD-safe surface. Do not use the handle to pull the disk drive out of the bay.

Note: When the IX-brick has only one disk drive installed, it should reside in the rightmost bay.

4. Install the drive bay cover.

After you remove the disk drive, power on the IX-brick according to the instructions in the “Powering On the System” on page 4.




Figure 12-12 Removing a Disk Drive

Replacing a TP900 Drive Carrier Module



Caution: Follow electrostatic discharge (ESD) precautions while replacing the drive carriers. Avoid contact with the backplane components and module connectors.

To remove and replace a drive carrier module from the SGI TP900 enclosure, follow these steps:



Caution: Do not remove a drive carrier unless a replacement or a dummy carrier can be immediately added. The system must not be run without all modules in place. Any unused drive bays must be fitted with a dummy carrier module.

1. Spin down the drive prior to removal, using the operating system.



Caution: Damage can occur to a drive if it is removed while still spinning. If you cannot use the operating system to spin down the drives prior to removal, perform all steps of the following procedure to ensure that the drive has stopped prior to removal.

Note: The anti-tamper lock must be off. If the drive module is locked, insert the Torx screwdriver (included with the disk drive) into the socket in the lower part of the handle trim. Then turn it 90 degrees counterclockwise until the indicator visible in the center aperture of the handle shows black. See Figure 12-13.




Figure 12-13 Unlocking the Anti-Tamper Lock

2. Release the carrier handle by pressing the latch in the handle towards the right.
3. Gently withdraw the drive carrier module approximately 1 inch (25 mm), and then wait 30 seconds for the drive to spin down.
4. After 30 seconds, withdraw the module from the drive bay.

5. Release the handle of the replacement carrier by pressing the latch in the handle towards the right, as shown in Figure 12-14. Insert the carrier into the enclosure.




Figure 12-14 Releasing the Carrier Handle

Important: Ensure that the carrier is oriented so that the drive is uppermost and the handle opens from the left.

6. Gently slide the carrier all the way into the enclosure until it is stopped by the camming lever on the right of the carrier, as shown in Figure 12-15.




Figure 12-15 Carrier Camming Lever

7. Push the carrier handle into place—the camming foot on the base of the carrier will engage in a slot in the enclosure. Continue to push firmly until the handle fully engages. You should hear a click as the latch engages and holds the handle closed.

Installing or Replacing a D-brick2 Drive Carrier Module

This section explains how to install or replace a drive carrier in the D-brick2 storage module, as follows:

- “Installing a Drive Carrier Module” on page 153
- “Replacing a Drive Carrier Module” on page 158



Caution: Observe all ESD precautions when handling modules and components. Avoid contact with backplane components and module connectors. Failure to observe ESD precautions could damage the equipment.



Caution: Do not remove a drive carrier unless a replacement or dummy carrier can be immediately added. The system must not be run without all modules in place. Any unused drive bays must be fitted with a dummy carrier module.

Installing a Drive Carrier Module

Note the following:

- All disk drive bays must be filled with either a drive carrier module or a dummy drive module; no bay should be left completely empty.
- The drives in bays 1/1 and 4/4 of the D-brick2 storage module are required for enclosure management; these bays must always be occupied.

To install a new disk drive module in the storage system, follow these steps:

1. Ensure that you have enough drive carrier modules and dummy modules to occupy all bays.
2. Carefully open the bag containing the drive carrier module.



Warning: The disk drive handle might have become unlatched in shipment and might spring open when you open the bag. As you open the bag, keep it a safe distance from your face.

3. Place the drive carrier module on an antistatic work surface and ensure that the anti-tamper lock is disengaged (unlocked). The indicator window is red when the drive is locked. A disk drive module cannot be installed if its anti-tamper lock is activated (locked) outside the enclosure.

Drives are shipped with their locks set in the unlocked position. However, if a drive is locked, insert the Torx screwdriver (included with the disk drive) into the socket in the lower part of the handle trim and turn it 90 degrees counterclockwise until the indicator visible in the center aperture of the handle shows black. See Figure 12-16.




Figure 12-16 Unlocking the Drive Carrier Module

4. Open the handle of the replacement carrier by pressing the latch handle towards the right (see Figure 12-17).




Figure 12-17 Opening the Module Handle

5. Remove the dummy drive carrier module from the target drive bay.

6. Orient the new drive carrier module so that the hinge of the handle is on the right. Then slide the drive carrier module into the chassis until it is stopped by the camming lever on the right of the module (see Figure 12-18).




Figure 12-18 Inserting the Disk Drive Module in a D-brick2

7. Swing the drive handle shut and press it to seat the drive carrier module. The camming lever on the right of the module will engage with a slot in the chassis. Continue to push firmly until the handle fully engages with the module cap. You should hear a click as the latch engages and holds the handle closed.
8. Repeat steps 2 through 7 for all drive modules to be installed.

- When you finish installing the drive carrier module(s), lock the drive carrier module(s) in the chassis by inserting the Torx screwdriver in the drive module and turning it 90 degrees clockwise. The indicator in the drive carrier module turns red when the drive is locked. See Figure 12-19.




Figure 12-19 Locking the Drive Carrier Module

- Install dummy drive carrier modules in all empty drive bays. The drive handle and camming mechanisms operate the same as those in a standard drive carrier module.

Replacing a Drive Carrier Module

This section explains how to replace a defective drive carrier module.



Caution: Follow electrostatic discharge (ESD) precautions while replacing the drive carriers. Avoid contact with the backplane components and module connectors.

Note: Replace disk drive modules one at a time.

To remove a drive carrier module from the D-brick2 enclosure, follow these steps:

1. Make sure that enough disk drives and dummy drives are available to occupy all bays.
2. Ensure that users are logged off of the affected systems; back up data if necessary.
3. If the drive module is locked, insert the Torx screwdriver into the anti-tamper lock and turn it 90 degrees counterclockwise. The indicator in the drive carrier module turns black when the drive is unlocked. See Figure 12-20.




Figure 12-20 Unlocking the Disk Drive Module

4. Open the handle by pressing the latch on the module handle towards the right. Then gently slide the module out of the enclosure approximately 1 inch (25 mm), and wait 30 seconds for the drive to stop spinning. See Figure 12-21.



Caution: Damage can occur to a drive if it is removed while still spinning.




Figure 12-21 Removing the Drive Carrier Module

5. After 30 seconds, withdraw the module from the drive bay. Replace it immediately; follow the instructions in “Installing a Drive Carrier Module” on page 153.

Troubleshooting and Diagnostics

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

- “Troubleshooting Chart” on page 162
- “L1 Controller Error Messages” on page 164
- “SGI Electronic Support” on page 166

Troubleshooting Chart

Table 13-1 lists recommended actions for problems that can occur. To solve problems that are not listed in this table, use the SGI Electronic Support system or contact your SGI system support engineer (SSE). For more information about the SGI Electronic Support system, see the “SGI Electronic Support” on page 166.

Table 13-1 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 brick will not power on.	Ensure that the power switch (if applicable) at the rear of the brick is on (1 position). View the L1 display; see Table 13-2 if an error message is present. If the L1 controller is not running, contact your SSE. Check the connection between the brick and its power source.
The system will not boot the operating system.	Contact your SSE.
The Service Required LED illuminates on a CR-brick, an R-brick, an IX-brick, or a PX-brick.	View the L1 display of the failing brick; see Table 13-2 for a description of the error message.
The Failure LED illuminates on a CR-brick, an R-brick, an IX-brick, or a PX-brick.	View the L1 display of the failing brick; see Table 13-2 for a description of the error message.
The green or yellow LED of a NUMALink port (rear of R-brick) is not illuminated.	Ensure that the NUMALink cable is seated properly on the R-brick and the destination brick.
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.

Table 13-1 Troubleshooting Chart (continued)

Problem Description	Recommended Action
The System Status LED of the TP900 is amber.	Contact your SSE.
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 13-2 lists error messages that the L1 controller generates and displays on the L1 display. This display is located on the front of the CR-bricks, R-bricks, IX-brick, and PX-bricks.

Note: In Table 13-2, 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 13-2 L1 Controller Messages

L1 System Controller Message	Message Meaning and Action Needed
Internal voltage messages:	
ATTN: <power VRM description> high fault limit reached @ x.xxV	30-second power-off sequence for the brick.
ATTN: <power VRM description> low fault limit reached @ x.xxV	30-second power-off sequence for the brick.
ATTN: <power VRM description> high warning limit reached @ x.xxV	A higher than nominal voltage condition is detected.
ATTN: <power VRM description> low warning limit reached @ x.xxV	A lower than nominal voltage condition is detected.
ATTN: <power VRM description> level stabilized @ x.xxV	A monitored voltage level has returned to within acceptable limits.
Fan messages:	
ATTN: FAN <fan description> 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 <fan description> 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 <fan description> stabilized @ xx RPM	An increased fan RPM level has returned to normal.

Table 13-2 L1 Controller Messages (continued)

L1 System Controller Message	Message Meaning and Action Needed
Temperature messages: low alt.	
ATTN: <temp sensor description> advisory temperature reached @ xxC xxF	The ambient temperature at the brick's air inlet has exceeded 30° C.
ATTN: <temp sensor description> critical temperature reached @ xxC xxF	The ambient temperature at the brick's air inlet has exceeded 35 °C.
ATTN: <temp sensor description> fault temperature reached @ xxC xxF	The ambient temperature at the brick's air inlet has exceeded 40 °C.
Temperature messages: high alt.	
ATTN: <temp sensor description> advisory temperature reached @ xxC xxF	The ambient temperature at the brick's air inlet has exceeded 27 °C.
ATTN: <temp sensor description> critical temperature reached @ xxC xxF	The ambient temperature at the brick's air inlet has exceeded 31 °C.
ATTN: <temp sensor description> fault temperature reached @ xxC xxF	The ambient temperature at the brick's air inlet has exceeded 35 °C.
Temperature stable message:	
ATTN: <temp sensor description> stabilized @ xxC/xxF	The ambient temperature at the brick'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.
Brick appears to have been powered down	The L1 controller has registered a fault and has shut down.

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 problems.

Figure 13-1 shows the sequence of events that occurs if you use all of the SGI Electronic Support capabilities.

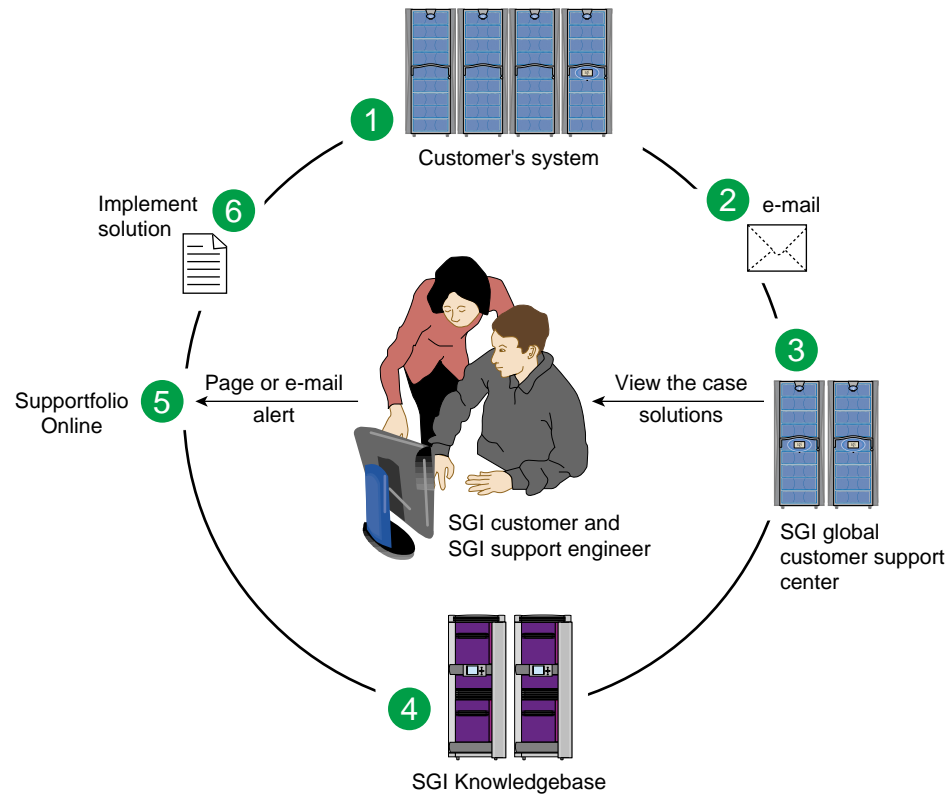


Figure 13-1 Full Support Sequence

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 SGI Linux ProPack release. 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 Web page:

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

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

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

Technical Specifications and Pinouts

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

- “System-level Specifications” on page 169
- “Physical Specifications” on page 170
- “Environmental Specifications” on page 171
- “Power Specifications” on page 172
- “I/O Port Specifications” on page 173

System-level Specifications

Table A-1 summarizes the Altix 3700 Bx2 system configuration ranges.

Table A-1 Altix 3700 Bx2 System Configuration Ranges

Category	Minimum	Maximum
Processors	8	512
CR-brick DIMM capacity	16	48
System main memory capacity	8 GB (4-node bricks)	6 TB (64-node bricks)
Available XIO I/O channels	2	128
Number of I/O bricks	One IX-brick	One IX-brick and seven PX-bricks

Physical Specifications

Table A-2 shows the physical specifications of the Altix 3700 Bx2 system.

Table A-2 Altix 3700 Bx2 Physical Specifications

Feature	Specification
Dimensions for a single tall rack, including doors and side panels	Height: 74.25 in. (1886 mm) Width: 30.88 in. (784 mm) Depth: 53.27 in. (1353 mm)
Shipping dimensions	Height: 80 in. (2032 mm) Width: 41.5 in. (1054 mm) Depth: 62.75 in. (1594 mm)
Weight (maximum)	
Compute rack	1400 lb (635 kg)
I/O rack	1225 lb (556 kg)
Disk rack	1400 lb (635 kg)
Shipping weight (maximum)	
Compute rack	1785 lb (810 kg)
I/O rack	1588 lb (720 kg)
Disk rack	1785 lb (810 kg)
Access requirements	
Front	48 in. (1219 mm)
Rear	48 in. (1219 mm)
Side	None

Environmental Specifications

Table A-3 lists the environmental specifications of the system.

Table A-3 Environmental Specifications

Feature	Specification
Temperature tolerance (operating)	+5 °C (41 °F) to +35 °C (95 °F) (up to 1500 m / 5000 ft.) +5 °C (41 °F) to +30 °C (86 °F) (1500 m to 3000 m / 5000 ft. to 10,000 ft.)
Temperature tolerance (non-operating)	-40 °C (-40 °F) to +60 °C (140 °F)
Relative humidity	10% to 80% operating (no condensation) 10% to 95% non-operating (no condensation)
Heat dissipation to air Altix 3700 Bx2 (rack)	37.2 kBTU/hr maximum (based on 10.9 kW)
Cooling requirement	Ambient air
Air flow: intake (front), exhaust (rear)	Approximately 3200 CFM
Maximum altitude	10,000 ft. (3,049 m) operating 40,000 ft. (12,195 m) non-operating
Acoustical noise level	Less than 65 dBa maximum

Power Specifications

Table A-4 shows the power specifications for the system.

Table A-4 Power Specifications

Feature	Specification
Single-phase power requirements	
Voltage	200-240V (180-254 VAC min/max)
Frequency	50-60 Hz
Power	11.2 kW per rack
Hold-up time	20 ms
Total harmonic distortion	Less than 14% at full load

I/O Port Specifications

This section contains specifications and port pinout information for the base I/O ports of your system, as follows:

- “Ethernet Port” on page 174
- “Serial Ports” on page 175

Ethernet Port

The system auto-selects the Ethernet port speed and type (duplex vs. half-duplex) when the server is booted, based on what it is connected to. Figure A-1 shows the Ethernet port.




Figure A-1 Ethernet Port

Table A-5 shows the cable pinout assignments for the Ethernet port operating in 10/100-Base-T mode and also operating in 1000Base-T mode.

Table A-5 Ethernet Pinouts

Ethernet 10/100Base-T Pinouts		Gigabit Ethernet Pinouts	
Pins	Assignment	Pins	Assignment
1	Transmit +	1	Transmit/Receive 0 +
2	Transmit -	2	Transmit/Receive 0 -
3	Receive +	3	Transmit/Receive 1 +
4	NU	4	Transmit/Receive 2 +
5	NU	5	Transmit/Receive 2 -
6	Receive -	6	Transmit/Receive 1 -
7	NU	7	Transmit/Receive 3 +
8	NU	8	Transmit/Receive 3 -

NU = Not used

Serial Ports

The IX-brick features two standard available 9-pin serial ports via the base I/O PCI card. These ports are capable of transferring data at rates as high as 230 kbps. Other features of the ports include the following:

- Programmable data, parity, and stop bits
- Programmable baud rate and modem control

Figure A-2 shows a serial port.

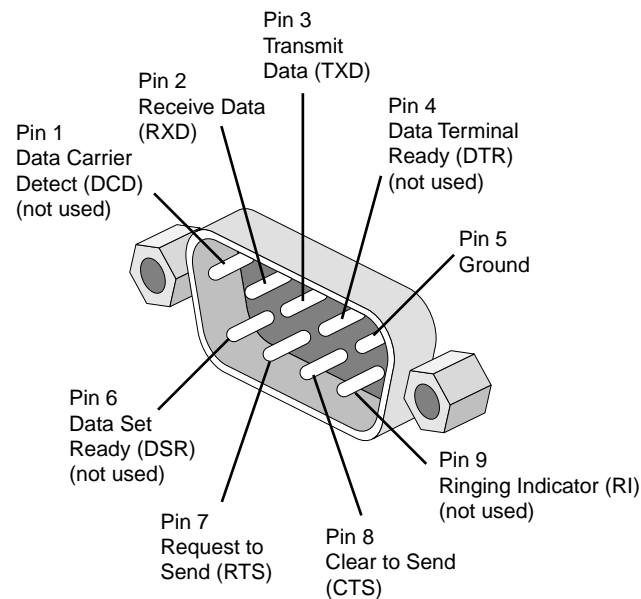


Figure A-2 Serial Port

Table A-6 shows pinout assignments for the 9-pin male DB-9 connector.

Table A-6 Serial Port Pinout

Pin	Assignment	Description
1	DCD	Data carrier detect
2	RXD	Receive data
3	TXD	Transmit data
4	DTR	Data terminal ready
5	GND	Signal ground
6	DSR	Data set ready
7	RTS	Request to send
8	CTS	Clear to send
9	RI	Ring indicator

Safety Information and Regulatory Specifications

This appendix provides safety information and regulatory specifications for your system in the following sections:

- “Safety Information” on page 177
- “Regulatory Specifications” on page 179

Safety Information

Read and follow these instructions carefully:

1. Follow all warnings and instructions marked on the product and noted in the documentation included with this product.
2. Unplug this product before cleaning. Do not use liquid cleaners or aerosol cleaners. Use a damp cloth for cleaning.
3. Do not use this product near water.
4. 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.
5. Slots and openings in the system are provided for ventilation. To ensure reliable operation of the product and to protect it from overheating, these openings 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.
6. This product should be operated from 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.
7. Do not allow anything to rest on the power cord. Do not locate this product where people will walk on the cord.
8. Never push objects of any kind into this product through cabinet slots as 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.

9. Do not attempt to service this product yourself except as noted in this guide. Opening or removing covers of node and switch internal components may expose you to dangerous voltage points or other risks. Refer all servicing to qualified service personnel.
10. Unplug this product from the wall outlet and refer servicing to qualified service personnel under the following conditions:
 - When the power cord or plug is damaged or frayed.
 - If liquid has been spilled into the product.
 - If the product has been exposed to rain or water.
 - If the product does not operate normally when the operating instructions are followed. Adjust only those controls that are covered by the operating instructions since 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 the cabinet has been damaged.
 - If the product exhibits a distinct change in performance, indicating a need for service.
11. If a lithium battery is a soldered part, only qualified SGI service personnel should replace this lithium battery. For other types, replace it only with the same type or an equivalent type recommended by the battery manufacturer, or the battery could explode. Discard used batteries according to the manufacturer's instructions.
12. Use only the proper type of power supply cord set (provided with the system) for this unit.
13. Do not attempt to move the system alone. Moving a rack requires at least two people.
14. Keep all system cables neatly organized in the cable management system. Loose cables are a tripping hazard that cause injury or damage the system.

Regulatory Specifications

The following topics are covered in this section:

- “CMN Number” on page 179
- “CE Notice and Manufacturer’s Declaration of Conformity” on page 179
- “Electromagnetic Emissions” on page 180
- “Shielded Cables” on page 182
- “Electrostatic Discharge” on page 182
- “Laser Compliance Statements” on page 183
- “Lithium Battery Statements” on page 184

This SGI system conforms to several national and international specifications and European Directives listed on the “Manufacturer’s Declaration of Conformity.” The CE mark insignia displayed on each device is an indication of conformity to the European requirements.



Caution: This product 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.

CMN Number

The model number, or CMN number, for the system is on the system label, which is mounted inside the rear door on the base of the rack.

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 from 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。

Figure B-1 VCCI Notice (Japan Only)

Chinese Class A Regulatory Notice

Figure B-2 Chinese Class A Regulatory Notice

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.

Figure B-3 Korean Class A Regulatory Notice

Shielded Cables

This SGI system is FCC-compliant under test conditions that include the use of shielded cables between the system and its peripherals. Your system and any peripherals 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 need to be shielded.

Optional monitor cables supplied with your system use additional filtering molded into the cable jacket to reduce radio frequency interference. Always use the cable 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 immune 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 unit 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. The wrist strap is used during the installation of these upgrades to prevent the flow of static electricity, and it should 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's classification label is located on the drive.



Warning: Avoid exposure to the invisible laser radiation beam when the device is open.



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 såteeseenstirra 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 Statements



Warning: If a lithium battery is a soldered part, only qualified SGI service personnel should replace this lithium battery. For other types, replace the battery only with the same type or an equivalent type recommended by the battery manufacturer, or the battery could explode. Discard used batteries according to the manufacturer's instructions.



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 suositteluun tyypin. 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

A

Altix 3700 Bx2 server
 illustration, 48
Altix 3700 Bx2 servers
 monitoring locations, 11
 powering off, 8
 powering on, 4
apropos command, xxii
audible alarm, 126

B

Base IO card
 location, 76
battery statements, 184
block diagram
 system, 50
brick
 monitored at L1 controller, 11
 powering off, 8
 powering on, 4
brick slot numbering, 57

C

cache coherency, 52
carrier module
 replacement procedure, 153

CMN number, 179
command targeting
 L1 mode, 28
 L2 mode, 36
console
 L1 mode power off, 9
 L1 mode power on, 7
 powering off bricks, 9
 powering on bricks, 7
console communications, 40
 subchannels, 40
console mode
 from L1 mode, 30
 from L2 emulator mode, 39
console selection, 40
controller
 L1, 17
cooling module LED
 TP900 storage module, 118
CR-brick
 block diagram, 68
 front panel components, 69
 important notes, 74
 overview, 65
 rear panel components, 71
 technical specifications, 74
customer service, xxiii

D

D-brick2, 131
 components, 123
 cooling module, 125
 drive carrier modules, 124
 dummy modules, 124
 features, 123
 front view, 122
 LRC module, 127
 overview, 122
 power supply module, 125
 powering off, 129
 powering on, 128
 rear view, 122
 technical specifications, 130-132

D-brick2 components
 weight, 130

D-brick2 operators panel, 126

disk drive carrier module
 replacement procedure, 153

disk drive installation, 144

disk drive LED
 TP900 storage module, 115

disk drive module
 adding, 153-157
 antitamper lock
 disengaging, 154
 replacing, 158-159

distributed power supplies, 109
 front view, 109
 LEDs, 109
 rear view, 109
 technical specifications, 111

documentation
 available via the World Wide Web, xxii
 conventions, xxiii
 man pages, xxii

drive carrier module

antitamper lock
 disengaging, 158

E

embedded support partner (ESP), 168
 system support, 10

environmental specifications, 171

error messages, 164

ESI/ops panel module, 17

ESP (embedded support partner)
 system support, 10

Ethernet connector, 78

external SCSI connector, 78, 82

F

front and rear panels of the CR-brick, 69

front panel display
 L1 controller, 22

H

hot swap disk drive replacement, 153

I

information, warning, and error message
 viewing in L1 mode, 28

information, warning, and error messages
 viewing in L2 mode, 32, 38

install
 disk drive, 144

IX-brick
 Base IO card, 78

- block diagram, 77
- front panel components, 79
- overview, 76
- PCI and PCI-X card slots, 83
- PCI and PCI-X slot rules, 84
- PCI and PCI-X slots, 81
- PCI card configuration guidelines, 83
- rear panel, 82
- rear panel components, 81
- RT0 and RT1, 82
- technical specifications, 85
- XIO 11 and XIO 10 connectors, 81

K

- Knowledgebase, 168

L

- L1 console selection, 30
- L1 controller, 11
 - description, 20
 - error messages, 164
 - front panel display, 22
 - front panel display items, 22
 - functions, 20
- L1 firmware upgrade, 43
- L1 mode
 - command targeting, 28
 - from L1 software, 26
 - from L2 emulation mode, 42
 - powering off, 29
 - powering off at console, 9
 - powering on, 29
 - powering on at console, 7
 - resetting the brick, 29
 - to console mode, 30
 - viewing information, warning, and error

- messages, 28, 32
 - viewing system configuration, 27
- L1 software
 - modes of operation, 25
 - to L1 mode, 26
- L2 emulation mode
 - to L1 mode, 42
 - viewing system configuration, 35
- L2 emulator mode
 - powering off, 38
 - powering on, 38
 - resetting system, 38
 - setting command targeting, 36
 - to console mode, 39
 - viewing information, warning, and error messages, 38
- laser compliance statements, 183
- LED
 - cooling module
 - TP900 storage module, 118
 - disk drive
 - TP900 storage module, 115
 - power supply
 - TP900 storage module, 117
 - status
 - TP900 storage module, 115
- lithium battery warning statements, 2, 184

M

- modes of operation
 - L1 software, 25
- monitoring
 - bricks with L1 controller, 11
 - embedded support partner, 10
 - server, 11

N

non-uniform memory access (NUMA), 53

numbering

bricks in a rack, 57

racks, 57

P

PCI and PCI-X cards

installation procedure, 136

IX-brick PCI and PCI-X slots, 81

IX-brick slot rules, 84

PX-brick PCI and PCI-X slots, 92

PX-brick slot rules, 95

slots in IX-brick, 83

slots in PX-brick, 94

PCI and PCI-X slots

IX-brick, 81

PX-brick, 92

PCI card configuration guidelines

IX-brick, 83

PCI cards

cards supported, 84, 95

carrier, 85, 96

PCI carrier

for PCI cards, 85, 96

PCI-X cards

cards supported, 84, 95

carrier, 85, 96

PCI-X carrier

for PCI-X cards, 85, 96

physical specifications

Altix 3700 Bx2 Physical Specifications, 170

pinouts

Ethernet connector, 174

serial connector, 175

power

connector

TP900 storage module, 117

power bay

features, 106

front panel components, 107

front view, 106, 107

overview, 105

rear panel components, 108

rear view, 106, 108

technical specifications, 111

power specifications, 172

power supply LED

TP900 storage module, 117

powering off

at system console, 9

bricks and server, 8

in L1 mode, 9

L1 mode, 29

L2 mode, 38

powering on

at system console, 7

bricks and server, 4

in L1 mode, 7

L1 mode, 29

L2 emulator mode, 38

preparation, 4

product support, xxiii

PX-brick

block diagram, 89

front panel components, 90

overview, 88

PCI and PCI-X card installation, 136

PCI and PCI-X card slots, 94

PCI and PCI-X slot rules, 95

PCI and PCI-X slots, 92

rear panel, 93

rear panel components, 92

technical specifications, 96

XIO 11 and XIO 10 connectors, 92

R

- RAS features, 53
- R-brick
 - block diagram, 99
 - features, 98
 - front panel components, 100
 - overview, 97
 - port specifications, 103
 - rear panel components, 101
 - rear view, 102
 - technical specifications, 103
- resetting system
 - L1 mode, 29
 - L2 mode, 38
- RT0 and RT1
 - IX-brick, 82
- RTO and RTI connectors, 78

S

- SCSI connector, external, 78, 82
- server
 - monitoring locations, 11
 - powering off, 8
 - powering on, 4
- SGI Knowledgebase. See Knowledgebase
- status LED
 - TP900 storage module, 115
- subchannels
 - associated with console communications, 40
- support
 - embedded support partner (ESP), 10
- Supportfolio Online, 168
- system architecture, 49
- system block diagram, 50
- system components
 - Altix 3700 Bx2 server, 56

- list of, 55
- optional, 58
- system controller, 15
- system features, 51
- system overview, 45

T

- tall rack
 - features, 60
 - locking front door, 62
- technical specifications
 - CR-brick, 74
 - distributed power supplies, 111
 - IX-brick, 85
 - power bay, 111
 - PX-brick, 96
 - R-brick, 103
 - system level, 169
- technical support, xxiii
- to monitor bricks, 11
- TP900 storage module
 - power connector, 117
- troubleshooting
 - L1 controller error messages, 164
 - problems and recommended actions, 162

U

- upgrading L1 firmware, 43

V

- viewing system configuration
 - L1 mode, 27

W

warning and error messages
 viewing in L1 mode, 28
warning messages
 viewing in L2 mode, 32, 38

X

XIO 11 and XIO 10
 IX-brick, 81
 PX-brick, 92