

SGI® Altix® 4700 System User's Guide

Document Number 007-4823-002

CONTRIBUTORS

Written by Mark Schwenden

Illustrated by Chrystie Danzer

Production by Mark Schwenden

Engineering contributions by Rich Altmaier, Steve Bowen, Michael T. Brown, Dick Brownell, Holly Brunell, David Collins, Thomas Crapisi, Steve Dean, Hartmut Gottwald, Paul Wiley, Gary Meyer, Galen Flunker, Tom Hotle, Mark Koneazny, William Kellerman, Paul Kinyon, Paul Pedersen, Gregory Thorson, and Gary Spilde.

COPYRIGHT

© 2007 SGI 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 SGI, 1140 E. Arques Avenue, Sunnyvale, CA 94085.

TRADEMARKS AND ATTRIBUTIONS

SGI, the SGI logo, and Altix are registered trademarks, and NUMALink is a trademark, of SGI, 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	October, 2006 First release
-002	March, 2007 Engineering updates describing 2D graphics card, IA2 base I/O blade and RAID 1 option.

Contents

- List of Figures xi**
- List of Tables xiii**
- About This Guide xv**
 - Audience. xv
 - Important Information xv
 - Chapter Descriptions xvi
 - Related Publications xvii
 - Conventions xix
 - Product Support xix
 - Reader Comments xx
- 1. Operation Procedures 1**
 - Precautions 1
 - ESD Precaution 1
 - Safety Precautions 2
 - System Controller Network 3
 - Console Connections 4
 - Connecting to the System Controller 4
 - Communicating with the System Controller 5
 - Powering the System On and Off 5
 - Preparing to Power On 6
 - Powering On and Off at the Console 9
 - Powering On and Off from the L1. 9
 - Powering-On and Off From the L2 11
 - Using Embedded Support Partner (ESP) 12

Monitoring Your Server	13
2D Graphics Video Interface.	15
Optional Components	16
PCI Cards	17
Adding or Removing Disk Drives	18
Reconfigure the System Disk Drive(s) for RAID 1	18
2. System Control	25
Levels of System Control	27
System Controller Interaction	27
L1 Controller	28
L1 Controller Functions	28
L1 Front Panel Display	29
Ethernet Switch	30
Console Hardware Requirements	30
Operating the L1	31
L1 Mode	32
Viewing System Configuration (from an IRU's Perspective)	33
Command Targeting	33
Console Mode from L1	34
L1 Console Selection	35
Viewing Information, Warnings, and Error Messages	36
L2 Operation	36
Configuring an L2's IP Address.	37
Viewing System Configuration	38
Setting Command Targeting.	40
Viewing Information, Warnings, and Error Messages	42
Powering On, Powering Off, and Resetting the System From the L2	42
Console Mode from the L2	43
Console Selection	44
L1 Mode From L2	45
Upgrading L1 Firmware	47
Upgrading L1 Firmware	47

3.	System Overview	49
	System Models	50
	System Architecture	52
	System Features	55
	Modularity and Scalability	55
	Distributed Shared Memory (DSM)	55
	Distributed Shared I/O	56
	ccNUMA Architecture	56
	Cache Coherency	56
	Non-uniform Memory Access (NUMA)	57
	Reliability, Availability, and Serviceability (RAS)	57
	System Components	59
	Bay (Unit) Numbering	61
	Rack Numbering	61
	Optional System Components	61
4.	Rack Information	63
	Overview	63
	Tall Altix Rack (42U)	64
	Technical Specifications	66
5.	Optional Dense Router Module	67
	Overview	67
	External Components	69
	Front Panel Components	69
	Rear Panel Components	71
	Technical Specifications	72
6.	Maintenance and Upgrade Procedures	73
	Maintenance Precautions and Procedures	73
	Preparing the System for Maintenance or Upgrade	74
	Returning the System to Operation	74
	Installing or Removing Internal Parts	75

Replacing IRU Components	76
Remove and Replace the System Control Display Panel	76
Removing or Adding IRU Power Supplies	79
Removing and Replacing IRU Fans (Blowers)	80
Removing or Replacing a System Blade From the IRU.	84
Removing and Replacing a Compute/Memory Blade	84
Memory DIMM Group Placement and Guidelines	87
Remove and Replace the DVD Drive	90
Overview of PCI Express Operation	93
Adding or Replacing PCI/PCI-X Cards	94
Installing PCI Cards in the Three-Slot Blade	95
Add or Remove a PCI Card in Base IA Blades	102
Installing or Replacing a Disk Drive in the IA/IA2 Blade	104
Removing a System Disk Drive	105
Installing a System Disk Drive	106
7. Troubleshooting and Diagnostics	109
Troubleshooting Chart	110
L1 Controller Error Messages	111
LED Status Indicators	113
IRU Power Supply LEDs	113
IRU NUMAlink Router Port LEDs	113
Compute/Memory Blade LEDs	114
SGI Electronic Support	115
A. Technical Specifications and Pinouts	119
System-level Specifications	119
Physical Specifications	120
Environmental Specifications	121
Power Specifications	122

I/O Port Specifications123
Ethernet Port124
Serial Ports125
External SAS/SATA Port Connector127
USB Type A Connector128
B. Safety Information and Regulatory Specifications129
Safety Information129
Regulatory Specifications131
CMN Number131
CE Notice and Manufacturer’s Declaration of Conformity131
Electromagnetic Emissions132
FCC Notice (USA Only)132
Industry Canada Notice (Canada Only)133
VCCI Notice (Japan Only).133
Chinese Class A Regulatory Notice133
Korean Class A Regulatory Notice133
Shielded Cables134
Electrostatic Discharge134
Laser Compliance Statements135
Lithium Battery Statements.136
Index137

List of Figures

Figure 1-1	IRU Power Supply Cable Location Example	6
Figure 1-2	Single-Phase Rack PDU Circuit Breaker Switches	7
Figure 1-3	Three-Phase PDU Example	8
Figure 1-4	IRU L1 Controller Display Location Example	14
Figure 1-5	2D Graphics Card Location in IA2 Blade	16
Figure 2-1	SGI Altix 4700 L1/L2 System Control Network (Example).	26
Figure 2-2	L1 Front Panel	29
Figure 2-3	Ethernet Switch System Controller Block Diagram (Example)	30
Figure 3-1	SGI Altix 4700 Systems	50
Figure 3-2	SGI Altix 4700 Server System (Single Rack)	51
Figure 3-3	Blade, IRU and Rack Components	52
Figure 3-4	Functional Block Diagram of the Individual Rack Unit (IRU)	54
Figure 3-5	Altix 4700 IRU System Components Example	60
Figure 4-1	Tall (42U) Altix Rack and Front Lock	65
Figure 5-1	Dense Router Functional Block Diagram	68
Figure 5-2	Dense Router Module Front and Rear Panels	69
Figure 5-3	Dense Router Functionality	70
Figure 6-1	Removing the L1 Controller Panel	77
Figure 6-2	Replacing the L1 Controller Panel	78
Figure 6-3	Removing an IRU Power Supply	79
Figure 6-4	Replacing an IRU Power Supply.	80
Figure 6-5	IRU Fan Assembly (Blowers)	81
Figure 6-6	Removing a Fan From the IRU	82
Figure 6-7	Replacing an IRU Fan	83
Figure 6-8	Removing a Compute/Memory Blade From the IRU	85

Figure 6-9	Replacing a Compute/Memory Blade in the IRU	86
Figure 6-10	DIMM Group Locations Diagram	88
Figure 6-11	DIMM Memory Slot Location Example	89
Figure 6-12	DVD Removal from IA Blade	91
Figure 6-13	DVD Drive Replacement in the IA Blade	92
Figure 6-14	Comparison of PCI/PCI-X Connector with PCI Express Connectors	93
Figure 6-15	Removing a Card Carrier From the Three-Slot Blade	95
Figure 6-16	Extracting the Carrier Metal Filler Plate	96
Figure 6-17	Adjusting the Carrier Guide Bar	97
Figure 6-18	Mounting Card in Carrier	98
Figure 6-19	Mounting Half-height PCI Card into Carrier	99
Figure 6-20	Moving Carrier Guide Bar to Secure Half-height PCI Card	99
Figure 6-21	Installing the Bracket to Secure a Half-height PCI Card	100
Figure 6-22	Installing a Card in a Slot	101
Figure 6-23	Seating the Card in the Slot	101
Figure 6-24	Removing the Card Access Panel From the IA Blade	102
Figure 6-25	Replacing the IA Blade PCI Access Door	103
Figure 6-26	Location of Disk Drive Bays	104
Figure 6-27	Removing a Disk Drive	105
Figure 6-28	Replacing a Disk Drive	107
Figure 7-1	Compute Blade Status LED Locations	114
Figure 7-2	Full Support Sequence	115
Figure A-1	Ethernet Port	124
Figure A-2	Serial Port Connector	125
Figure A-3	Pin Number Locations for External SAS/SATA Port	127
Figure A-4	Pin Number Locations for USB Type A Connector	128
Figure B-1	VCCI Notice (Japan Only)	133
Figure B-2	Chinese Class A Regulatory Notice	133
Figure B-3	Korean Class A Regulatory Notice	133

List of Tables

Table 4-1	Tall Altix Rack Technical Specifications	66
Table 5-1	Dense Router Technical Specifications	72
Table 5-2	Dense Router Port Specifications	72
Table 6-1	Customer-replaceable Components and Maintenance Procedures	75
Table 6-2	SGI PCIe Support Levels	94
Table 7-1	Troubleshooting Chart110
Table 7-2	L1 Controller Messages111
Table 7-3	Power Supply LED States113
Table A-1	Altix 4700 System Configuration Ranges119
Table A-2	Altix 4700 Physical Specifications120
Table A-3	Environmental Specifications121
Table A-4	Power Specifications122
Table A-5	Ethernet Pinouts124
Table A-6	Serial Port Pinout.126
Table A-7	SAS/SATA Pin Assignments127
Table A-8	Pin Assignments for USB Type A Connector128

About This Guide

This guide provides an overview of the architecture, general operation and descriptions of the major components that compose the SGI® Altix® 4700 family of servers. It also provides the standard procedures for powering on and powering off the system, basic troubleshooting and maintenance information, and important safety and regulatory specifications.

Audience

This guide is written for owners, system administrators, and users of SGI Altix 4700 computer systems. It is written with the assumption that the reader has a good working knowledge of computers and computer systems.

Important Information



Warning: To avoid problems that could void your warranty, your SGI or other approved system support engineer (SSE) should perform all the set up, addition, or replacement of parts, cabling, and service of your SGI Altix 4700 system, with the exception of the following items that you can perform yourself:

- Using your system console and your L1/L2 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.
- Replacing fan modules and power supplies in the individual rack unit (IRU) enclosures.
- Adding and replacing disk drives in I/O blades.
- Using the On/Off switch and other switches (the reset and non-maskable interrupt [NMI] switches) on the IRU.
- Using the ESI/ops panel (operating panel) on optional mass storage bricks.

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 SGI Altix 4700 systems.
- Chapter 4, “Rack Information,” describes the rack sizes and features.
- Chapter 5, “Optional Dense Router Module,” describes all the connectors and LEDs located on the front and rear of the optional four-board, 32-port router enclosure.
- Chapter 6, “Maintenance and Upgrade Procedures,” provides instructions for installing or removing the customer-replaceable components of your system.
- Chapter 7, “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 regulatory information related to use of the Altix 4700 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 4700 series system:

- *SGI L1 and L2 Controller Software User's Guide*
(P/N 007-3938-00x)

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

- *Reconfigurable Application-Specific Computing User's Guide*
(P/N 007-4718-00x)

This guide discusses the optional SGI reconfigurable application-specific software computing (RASC) program that delivers scalable, configurable computing elements for the SGI Altix family of servers and superclusters.

- 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 DVD) 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 bold 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
- 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
1140 East Arques Avenue, M/S 50-1-946
Sunnyvale, California 94085

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
- “System Controller Network” on page 3
- “Powering the System On and Off” on page 5
- “Using Embedded Support Partner (ESP)” on page 12
- “Monitoring Your Server” on page 13
- “Optional Components” on page 16

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 129.



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.



Warning: If a lithium battery is installed in your system as 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.

System Controller Network

All Altix 4700 enclosures contain an embedded microprocessor board and display assembly known as the system controller. This microprocessor runs an embedded version of the Linux operating system. The system controller runs off standby power and is running as long as the enclosure is connected to an active power source.

There are two primary applications that run on the system controller. The L1, or Level 1 system controller, is an application that provides control and monitoring functionality for each individual rack unit (IRU) enclosure, and communication to other L1s in adjacent enclosures connected via NUMALink 4 cables. The L1 is always resident.

The Level 2 (L2) system controller is an application that provides control over multiple L1s and communication to other L2s. The L2 is resident when the enclosure is connected by an Ethernet connection to a Local Area Network (LAN).

The system controller network provides the following functionality:

- Powering the entire system on and off.
- Powering individual IRUs and Dense routers on and off.
- Monitoring the environmental state of the system
- Viewing 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 individual rack unit (IRU) enclosure. See the *SGI L1 and L2 Controller Software User's Guide* for a complete list of commands.
- Provides access to the system OS console allowing you to run diagnostics and boot the system.

Console Connections

A console is defined as a connection to the system that provides access to the system controller network. A console can be a LAN-attached personal computer, laptop or workstation (Ethernet connection) or a dumb terminal (serial connection).

The Altix 4700 series supports two types of console connections:

- An ethernet connection on the system control board of an IRU or on a Dense router - an RJ45 connection is typically labeled “L2 host”
- A serial connection to the serial console port on the system control board of an IRU or the serial console port on a Dense router - a DB9 connector typically labeled “Console”

Connecting to the System Controller

The ethernet connection is the preferred method of accessing the system console. Depending on the size of the system, there will be one or more ethernet connections utilized.

When an enclosure is connected to the LAN, the system controller will start the ethernet interface with either the pre-assigned IP address or will attempt to acquire an IP address via DHCP, and will then spawn the L2 application. See the section on setting up an IP address for further details. Once the system controller is connected to the LAN, it can be accessed via a simple telnet session.

The serial connection is used to communicate directly with the L1 system controller. This connection is typically used for service purposes or for system controller and system console access in small systems where an ethernet connection is not used or available.

Communicating with the System Controller

Once a connection to the console is established, the user will be presented with either an L2 prompt (ethernet connection), known as L2 mode, or and L1 prompt (serial connection), known as L1 mode. From either of these prompts various system controller commands can be entered.

```
001c01-L1> <11 command>
```

```
olympic-101-L2> <12 command>
```

To access the system console, known as console mode, a control-d is entered at either the L1 prompt or the L2 prompt. To return to or escape back to the L1 or L2 mode, a control-t is entered. This escape is only temporary and you will be returned to console mode once the “Return/Enter” key is pressed. To re-engage L1 or L2 mode, enter control-t followed by either “11” or “12” depending on the original mode. System control commands are always entered in lower case unless otherwise specified.

Powering the System On and Off

This section explains how to power on and power off individual rack units, Dense routers, or your entire Altix 4700 system, as follows:

- “Preparing to Power On” on page 6
- “Powering On and Off at the Console” on page 9

For servers with a system console, you can power on and power off individual IRUs and Dense routers 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 may 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 12.

Preparing to Power On

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

1. Check to ensure that the cabling between the rack's power distribution unit (PDU) and the wall power-plug receptacle is secure.
2. For each individual IRU that you want to power on, make sure that the power cables are plugged into all the IRU power supplies correctly, as shown in Figure 1-1. Setting the circuit breakers on the PDUs to the "On" position will apply power to the IRU and will start the system controller(s) in the IRUs. Note that the system controller in each IRU and Dense router stays powered on as long as there is power coming into the unit.




Figure 1-1 IRU Power Supply Cable Location Example

3. If you plan to power on a server that includes optional mass storage enclosures, make sure that the power switch on the rear of each PSU/cooling module (one or two per enclosure) is in the **1** (on) position.
4. Make sure that all PDU circuit breaker switches (see the examples in Figure 1-2 and Figure 1-3) are turned on to provide power to the server system when the system is powered on.




Figure 1-2 Single-Phase Rack PDU Circuit Breaker Switches

Figure 1-3 shows an example of the three-phase PDU.




Figure 1-3 Three-Phase PDU Example

Powering On and Off at the Console

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

- If you have a console connected to a server with a serial interface, 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.

For detailed instructions on using a system console using the L1 mode, see “Operating the L1” on page 31.

- If you have a system console connected to a server with an ethernet interface, you can toggle between L2 and console mode, and power on your server with L2 controller commands, and monitor the power-on activity by changing to the console mode.

See the section “Console Hardware Requirements” in Chapter 2 for additional information on optional consoles.

Powering On and Off from the L1

Commands issued at the L1 prompt typically only affect the local enclosure. The following sections describe how to power on and power off your system in L1 mode.

The L1 controller display, located on the front of each IRU, should display `L1 running` once the power-on procedure starts (storage modules do not have L1s). The prompt on your console screen shows the rack and slot number of the IRU to which you have connected your console.

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 111 to learn what the error message indicates and how to resolve the problem.

If you want to power on the IRU (001c01 in our example) indicated in the prompt, enter the following command.

```
001c01-L1> power up
```

To power off:

```
001c01-L1> power down
```

If you are attempting to power on or power off a system with multiple enclosures from the L1 prompt, you will need to prepend the command with an asterisk as follows:

```
oo1c01-L1> * power up  
or  
oo1c01-L1> * power down  
(* indicates all)
```

Note: If you are accessing a large system via the serial connection, the entire system may not be accessible from this point and as a result only a portion of the system may be affected.

From the L1 prompt, display the system configuration information by entering the following command:

```
001c01-L1> config  
:0 001c01 L0C  
:2 002r01 L1H  
:8 002r03 R0H  
:4 002r05 L3H  
:6 002r07 R2H  
001c01-L1>
```

In L1 mode, you can obtain only limited information about the system configuration. An IRU has information about its internal blades, and also if other IRUs are NUMALink attached to the IRU, information about those IRUs.

Note: Verify that the power LED on the L1 displays turns on and lights 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 111 to learn what the error message indicates and how to best resolve the problem.

Powering-On and Off From the L2

Commands issued at the L2 prompt are typically sent to all enclosures in the system.

To power on a system:

```
system-101-L2> power up
```

To power off a system:

```
system-101-L2> power down
```

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 7.

Monitoring Your Server

You can monitor your Altix 4700 server from the following sources:

- On the L1 controller's display at the front of each IRU as shown in Figure 1-4 on page 14, you can monitor system operational status. For example, you can monitor error messages that warn of power or temperature values that are out of tolerance.
- A DVI or VGA monitor with USB keyboard/mouse can be connected for basic monitoring and administration of the Altix system. See the section "2D Graphics Video Interface" on page 15 for more information. SLES 10 is required for this card. RHEL 5 is *not* supported.
- You can connect an optional console via an Ethernet port adapter. You will need to connect either a local or remote workstation/PC to monitor the servers via Ethernet.

These console connections enable you to view the status and error messages generated by the L1/L2 controllers in your Altix 4700 rack. You can also use these consoles to input L1/L2 commands to manage and monitor your system. See the section "Console Hardware Requirements" in Chapter 2, for additional information on the L1/L2 console.




Figure 1-4 IRU L1 Controller Display Location Example

2D Graphics Video Interface

If your system was ordered in January 2007 or later, it may come equipped with an optional 2D graphics board interface (not supported with SLES9 or RHEL5 systems). This low-profile PCI interface card is installed in the IA2 (base I/O) blade in the Altix system. One 2D card is supported per system or partition; no audio function is supported. The 2D video card has the following features:

- 64 MB DDR graphics memory
- DMS-59 display connector for two display support (no synchronization between monitors)
- Simultaneous display to DVI and VGA monitors using an adapter
- Analog resolution up to 2048 x 1536 per display
- DVI resolution up to 1600 x 1200 per display

Note: A single USB keyboard/mouse is supported by the IA2 (base I/O) blade.

The 2D video interface card can be used for all basic interaction with your Altix system. You should note that it does **not** provide an interface to the L1/L2 controller and L1/L2 administrative commands cannot be issued through the 2D video interface.




Figure 1-5 2D Graphics Card Location in IA2 Blade

Optional Components

Besides adding a network-connected system console or basic VGA monitor, you can add or replace the following hardware items on your Altix 4700 series server:

- Peripheral component interface (PCI) cards into your system I/O blades.
- Disk drives in your IA blade (base I/O).

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 129.

PCI Cards

The PCI, PCI-X and PCIe based I/O sub-systems, are industry standard for connecting peripherals, storage and graphics to a processor blade. These are the primary configurable I/O system interfaces for the Altix 4700 series systems. They include:

- The IA blade (base I/O) used in the Altix 4700 system includes two half-height PCI/PCI-X slots. The two option cards install adjacent to the system disk(s) and DVD drive that are also resident in the IA blade. Note that the newer version of the IA blade that supports RAID 1 and DVD-R/W became available for order in January 2007. This enhanced version of the base I/O is generally referred to as the IA2.
- The optional three-slot PCI/PCI-X double-wide blade holds three sled-mounted PCI/PCI-X cards for easy insertion and extraction from the system.
- The optional single-wide PCI express (PCIe) blade supports two optional PCI express cards.
- The optional double-wide PCI express (PCIe) and PCI-X blade supports two optional PCI express slots and two PCI/PCI-X slots. This blade is commonly used for installing a pair of SGI 3D PCIe option cards. Note the width of the 3D cards precludes use of the PCI-X slots.

Other optional I/O blades provide additional PCI and PCI-X support for the Altix 4700 server system. Not all blades may be available with your system configuration.

Check with your SGI sales or service representative for availability. See Chapter 6, “Maintenance and Upgrade Procedures” for detailed instructions on installing or removing PCI cards in the blades. Information on installing or removing optional 3D graphics cards is included in the user’s guide shipped with the 3D graphics option.

Adding or Removing Disk Drives

The IA-blade (base I/O blade) within the IRU supports the system boot functions and contains one or two low-profile disk drives. See “Installing or Replacing a Disk Drive in the IA/IA2 Blade” on page 104 for detailed instructions on installing or removing disk drives.

Reconfigure the System Disk Drive(s) for RAID 1

System disk drive(s) ordered starting in March 2007 may come configured as either a non-RAID (jboot) or as a RAID 1 (disk mirror) system. Note that while RAID 0 (striping) is supported by the hardware, SGI recommends using either jboot or RAID 1 to ensure maximum reliability and data retention.

The type of disk configuration is ordered by the customer and configured at the SGI factory. If a requirement exists to reconfigure the system disk(s), use the following information to make the changes needed. The `lsiutil` configuration software is the only supported methodology for modifying the disks at the time this document was published. Use of other software interfaces has not been verified and cannot be supported.

Important: The example process that follows presumes there are two disks identical in size and speed installed in your Altix system IA2 (base I/O) blade. Also, note that this example of re-imaging a disk is done via network access.

The `lsiutil` file (efi version) resides at `/boot/efi`

Launch the configuration utility with the command:

```
fs0:\> lsiutil64.efi
```

You should see an interactive interface similar to the following:

```
LSI Logic MPT Configuration Utility, Version 1.xx, November 7, 200X
1 MPT Port found
Port Name      Chip Vendor/Type/Rev   MPT Rev  Firmware Rev  IOC
1.  01/00/01/0      LSI Logic SAS1068 B0   105      01100000      0
```

The SAS1068 is the chip that controls the device that the `lsiutil` utility was loaded from.

```
Select a device: [1-1 or 0 to quit] 1
```

1. Identify firmware, BIOS, and/or FCode
2. Download firmware (update the FLASH)
4. Download/erase BIOS and/or FCode (update the FLASH)
8. Scan for devices
10. Change IOC settings (interrupt coalescing)
13. Change SAS IO Unit settings
16. Display attached devices
18. Change WWID
20. Diagnostics
21. RAID actions
22. Reset bus
23. Reset target
24. Clear ACA
39. Force firmware download boot
45. Concatenate SAS firmware and NVDATA files
60. Show non-default settings
61. Restore default settings
97. Reset SAS phy
98. Reset SAS link
99. Reset port
 - e Enable expert mode in menus
 - p Enable paged mode in menus
 - w Enable logging

Main menu, select an option: [1-99 or e/p/w or 0 to quit] 8

Important: After this command is entered, the load device cannot be accessed again to open files until lsiutil has been exited and restarted.

SAS1068's links are 3.0 G, 3.0 G, down, down, down, down, down, down

B	T	L	Type	Vendor	Product	Rev	SASAddress	PhyNum
0	0	0	Disk	SGI	ST3146854SS	X422	5000c5000002cb45	0
0	1	0	Disk	SGI	ST3146854SS	X422	5000c500000121cd	1

Main menu, select an option: [1-99 or e/p/w or 0 to quit] 0

	Port Name	Chip Vendor/Type/Rev	MPT Rev	Firmware Rev	IOC
1.	01/00/01/0	LSI Logic SAS1068 B0	105	01100000	0

Select a device: [1-1 or 0 to quit] 1

1. Identify firmware, BIOS, and/or FCode
2. Download firmware (update the FLASH)

- 4. Download/erase BIOS and/or FCode (update the FLASH)
- 8. Scan for devices
- 10. Change IOC settings (interrupt coalescing)
- 13. Change SAS IO Unit settings
- 16. Display attached devices
- 18. Change WWID
- 20. Diagnostics
- 21. RAID actions
- 22. Reset bus
- 23. Reset target
- 24. Clear ACA
- 39. Force firmware download boot
- 45. Concatenate SAS firmware and NVDATA files
- 60. Show non-default settings
- 61. Restore default settings
- 97. Reset SAS phy
- 98. Reset SAS link
- 99. Reset port
 - e Enable expert mode in menus
 - p Enable paged mode in menus
 - w Enable logging

Main menu, select an option: [1-99 or e/p/w or 0 to quit] 21

- 1. Show volumes
- 2. Show physical disks
- 3. Get volume state
- 23. Replace physical disk
- 30. Create volume
- 31. Delete volume
- 32. Change volume settings
- 50. Create hot spare
- 99. Reset port
 - e Enable expert mode in menus
 - p Enable paged mode in menus
 - w Enable logging

RAID actions menu, select an option: [1-99 or e/p/w or 0 to quit] 1

0 volumes are active, 0 physical disks are active

RAID actions menu, select an option: [1-99 or e/p/w or 0 to quit] 30

B__T__L	Type	Vendor	Product	Rev	Disk Blocks	Disk MB
1. 0 0 0	Disk	SGI	ST3146854SS	X422	286749488	140014


```
2. 0 1 0 Disk      SGI      ST3146854SS      X422      286749488      140014
```

To create a volume, select two or more of the available targets

Select a target: [1-2 or RETURN to quit] 1

Select a target: [1-2 or RETURN to quit] 2

2 physical disks were created

Select volume type: [0=Mirroring, 1=Striping, default is 0]

Select volume size: [1 to 139898 MB, default is 139898]

Enable write caching: [Yes or No, default is No]

Zero the first and last blocks of the volume? [Yes or No, default is No]

Volume was created

RAID actions menu, select an option: [1-99 or e/p/w or 0 to quit] 1

1 volume is active, 2 physical disks are active

Volume 0 is Bus 0 Target 0, Type IM (Integrated Mirroring)

Volume State: degraded, enabled , resync in progress

Volume Settings: write caching disabled , auto configure

Volume Size 139898 MB, Stripe Size 0 KB, 2 Members

Primary is PhysDisk 0 (Bus 0 Target 56)

Secondary is PhysDisk 1 (Bus 0 Target 1)

RAID actions menu, select an option: [1-99 or e/p/w or 0 to quit] 99

Resetting port...

RAID actions menu, select an option: [1-99 or e/p/w or 0 to quit] 0

Main menu, select an option: [1-99 or e/p/w or 0 to quit] 0

	Port Name	Chip Vendor/Type/Rev	MPT Rev	Firmware Rev	IOC
1.	01/00/01/0	LSI Logic SAS1068 B0	105	01100000	0

Select a device: [1-1 or 0 to quit] 0

Sas(Pun0,Lun0) LSILOGICLogical Volume 3000

fs0:\> reset

.

```
.
.
EFI Boot Manager ver 1.10 [14.62]

Partition 0:
CBlades      2      Nodes      Enabled Disabled
RBlades      0      CPUs       3         0
IOBlades     1      Mem(GB)    4         0
              6         0

Please select a boot option

    EFI Shell
>>> netboot
    Boot option maintenance menu

Loading: netboot
Running LoadFile()

CLIENT MAC ADDR: 08 00 69 14 E3 0C
CLIENT IP: 137.38.82.91  MASK: 255.255.255.0  DHCP IP: 137.38.228.4
GATEWAY IP: 137.38.82.252  137.38.82.254

TSize.Running LoadFile()

Starting: netboot
ELILO
.
.
.
.
Quietly populating /dev/sda2 and /dev/sda1 with
ia64-propack51-sles10-Jan19

# reboot....
```

Note: The initial rsynchronization of the mirrored disk (depending on the size of the disk) may take up to 45 minutes or longer.

```
RAID actions menu, select an option: [1-99 or e/p/w or 0 to quit] 1

1 volume is active, 2 physical disks are active

Volume 0 is Bus 0 Target 0, Type IM (Integrated Mirroring)
```

```
Volume State: optimal, enabled
Volume Settings: write caching disabled, auto configure
Volume draws from Hot Spare Pools: 0
Volume Size 139898 MB, Stripe Size 0 KB, 2 Members
Primary is PhysDisk 1 (Bus 0 Target 56)
Secondary is PhysDisk 0 (Bus 0 Target 1)
```

For information on the two RAID 1 disks, use the `lsscsi` command as shown in the following example:

```
# lsscsi

[1:0:0:0]    disk      SGI          ST3146854SS    X422  -        /dev/sg2
[1:0:1:0]    disk      SGI          ST3146854SS    X422  -        /dev/sg3
[1:1:0:0]    disk      LSILOGIC    Logical Volume  3000  /dev/sdc /dev/sg4
```

This is an example of what you'll see on SLES10 or RHEL5 Linux operating systems. Disk [1:0:0:0] and [1:0:1:0] are the individual drives in the RAID 1, but you cannot access them through the disk driver. You access [1:1:0:0], which is `/dev/sdc` in this case. Caution must be exercised when entering Linux SCSI generic (sg) commands via the sg devices.

System Control

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

- “Levels of System Control” on page 27
- “System Controller Interaction” on page 27
- “L1 Controller” on page 28
- “Console Hardware Requirements” on page 30
- “Operating the L1” on page 31
- “L2 Operation” on page 36
- “Upgrading L1 Firmware” on page 47

The L1/L2 control system for the SGI Altix 4700 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.

Each IRU enclosure has a system controller that can communicate with other IRU system controllers when they are NUMAlinked together under a single system image. Each system controller constantly shares its information with all controllers in the system. Note that optional mass storage enclosures do *not* have a system controller.

Figure 2-1 shows an example system control network using an optional and separate (remote) workstation to monitor a single-rack Altix 4700 system.




Figure 2-1 SGI Altix 4700 L1/L2 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. Typically, an Ethernet connection to the system controller network is used. This Ethernet connection is either made from an IRU in a smaller system or from a Dense router in larger systems.

The system controller is designed into all IRUs. The level one (L1) system control hardware component is the basic administrative interface for each IRU. When an Ethernet interface is connected to the system control interface it spawns a level two (L2) application function. An Ethernet connection directly from the IRU or Dense router to a local private or public Ethernet allows the system to be administered directly from a local or remote console. Note that there is no interconnected system controller function in the optional storage modules.

The L1/L2 controllers within the system report and share status information via the NUMALink cables, thus maintaining controller configuration and topology information between all controllers in the system.

Note: Mass storage option enclosures are not specifically monitored by the system controller network. Most optional mass storage enclosures have their own internal microcontrollers for monitoring and controlling all elements of the disk array. See the owner's guide for your mass storage option for more information on this topic.

For information on attaching network connected older SGI systems see the *SGIconsole Hardware Connectivity Guide*, (P/N 007-4340-00x).

System Controller Interaction

In all Altix 4700 servers all the system controllers communicate with each other in the following ways:

- All enclosures within an Altix 4700 system communicate with each other through their NUMALink connections using low voltage differential signaling (LVDS).
- When connecting to the L1/ L2 host Ethernet connection on the system control board of an IRU or to a Dense router, the system controller spawns an L2 application providing L2 functionality.

L1 Controller

All IRUs and Dense routers have L1 controllers. The following subsections describe the basic features of all L1 controllers:

- “L1 Controller Functions” on page 28
- “L1 Front Panel Display” on page 29

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

The following list summarizes the control and monitoring functions that the L1 controller performs. Many of the L1 controller functions are common across both IRU and Dense routers; however, some functions are specific to the type of enclosure.

- Controls voltage margining within the IRU or Dense router
- Controls and monitors IRU and Dense router fan speeds
- Reads system identification (ID) PROMs
- Monitors voltage levels and reports failures
- Monitors and controls warning LEDs on the enclosure
- Monitors the On/Off power switch
- Monitors the reset switch and the nonmaskable interrupt (NMI) switch
- Reports the population of the PCI cards and the power levels of the PCI slots in installed I/O blades
- Powers on the PCI slots and their associated LEDs

L1 Front Panel Display

Figure 2-2 shows the L1 controller front panel on the IRU.




Figure 2-2 L1 Front Panel

The front panel display contains the following items:

- 2 x 12 character liquid crystal display (LCD). The display identifies the IRU, shows system status, warns of required service, and identifies a failed component.
- Power (On/Off) button (insert paper clip to actuate) and power on LED.
- Service required LED.
- Failure LED.
- Reset switch (insert paper clip to actuate).
- Non-maskable interrupt (NMI) switch (insert paper clip to actuate).

Note: The reset and NMI switch functions are available on the IRUs only.

Ethernet Switch

Use of an Ethernet switch is the preferred method of interconnecting large systems with multiple L2s and remote support hardware to multiple systems. An optional Ethernet switch provides multiple Ethernet connectors. Figure 2-3 shows example connections between the Ethernet switch, a network-connected workstation console, an IRU and other local SGI systems.




Figure 2-3 Ethernet Switch System Controller Block Diagram (Example)

Console Hardware Requirements

The console type and how these console types are connected to the Altix 4700 servers is determined by what console option is chosen.

If you have an Altix 4700 server and wish to use a “dumb terminal”, you can connect the terminal via a serial cable to the (DB-9) console port connector on the system control board of the IRU.

The terminal should 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)

Note that the serial console is generally connected to the first (bottom) IRU in any single rack configuration.

You can also connect to the RJ45 Ethernet connector on the system control board of the IRU and place it on a network. This starts an L2 on the IRU's system controller. You can use the `telnet` command to connect.

To verify that the L2 is running, at the L1 prompt enter the following:

```
001c01-L2> l2
L2 Controller is running
001c01-L2>
```

If the L2 is not running; enter:

```
001c01-L2> ! init 4
```

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

```
001c01-L1> l2 ip <ip address> <netmask> <broadcast>
001c01-L2> reboot_l1
```

These console connections enable you to view the status and error messages generated by the system 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, see “L2 Operation” on page 36.

For more information on connecting a console to an Altix 4700 series server, see “System Controller Network” on page 3. For more information on monitoring your server, see “Monitoring Your Server” on page 13.

Operating the L1

Each IRU and Dense router in the Altix 4700 system has an updated and enhanced system control implementation. This updated system controller provides both L1 and L2 functionality. The system controller utilizes an embedded version of the Linux operating system. L1 functionality is provided by an application that is always running on the system controller. When the enclosure is connected to a LAN via the L2 host connector, the system controller spawns an application that provides L2 functionality.

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 Altix 4700 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 enclosures is a complete and fully functional system controller. All the blades are interconnected by NUMALink and each shares its system control information with all other system controllers.

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 L1 port is connected directly to the console system (laptop, PC, etc.).

L1 Mode

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

```
001c01-L1>
```

Common operations are discussed in the following sections:

- “Viewing System Configuration (from an IRU’s Perspective)” on page 33
- “Command Targeting” on page 33
- “Viewing Information, Warnings, and Error Messages” on page 42

Viewing System Configuration (from an IRU's Perspective)

An L1 has limited knowledge of the system topology, depending on the system's configuration. Typically, an L1 has information only about L1s that are directly NUMALink connected.

In large configurations with more than one L1, the L1 may have knowledge of only a portion of the L1s in the system. These configurations require the use of the L2, see "L2 Operation" on page 36 for further details.

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

```
001c01-L1> config
:0 001c01 LOC
001c01-L1>
```

This example is a system with one IRU. The *<number>* that follows the colon (0, 1, and 2, from top to bottom in this example), refers to which local port the IRU is connected to or accessed through. The local (LOC) IRU is the IRU that is processing the command.

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

```
021c01-L1> config
:0 021c01 LOC
:2 021c11 L1H
:8 021r41 L0G
:5 022r41 R3G
021c01-L1>
```

Command Targeting

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

```
001c01-L1> * version
001c01:
L1 0.7.37 (Image A), Built 11/24/2005 14:59:42 [2MB image]
001c11:
L1 0.7.37 (Image A), Built 11/24/2005 14:59:42 [2MB image]
001c21:
L1 0.7.37 (Image A), Built 11/24/2005 14:59:42 [2MB image]
001c01-L1>
```

Commands can be targeted to other L1s by preceding the command with a rack and slot:

```
001c01-L1> 1.11 version
```

The command above issues a version command to IRU in rack 001, U position 11.

Some commands can be targeted to a specific blade within an IRU. Precede the command with the blade designator:

```
001c01-L1> b1 power down
```

The command above issues a power down command to the blade in blade slot 1 of the IRU in rack 001, U position 01.

Console Mode from L1

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

```
001c01-L1> Ctrl+D
entering console mode 001c01 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
001c01-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 001c01 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
001c01-L1> l1
L1 command processor engaged, <CTRL-D> for console mode.
001c01-L1>
```

L1 Console Selection

If the system contains more than one IRU and a serial connection is utilized for the console, the serial cable must be connected to the IRU that is located in the lowest rack and slot position.

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

```
001c01-L1> select  
console input: 001c01 console0  
console output: not filtered
```

The following are common subchannels associated with console communications:

- Subchannel 0A specifies Blade 0, CPU A.
- Subchannel 0C specifies Blade 0 CPU C.
- Subchannel 1A specifies Blade 1, CPU A.
- Subchannel 1C specifies Blade 1, CPU C.
- Subchannel 2A specifies Blade 2, CPU A.
- Subchannel 2C specifies Blade 2, CPU C.
- Subchannel 3A specifies Blade 3, CPU A.
- Subchannel 3C specifies Blade 3, CPU C.
- Subchannel console0 Blade 0 console subchannel.
- Subchannel console1 Blade 1 console subchannel.

The output from the `select` command:

```
“console input: 001c01 console0”
```

shows that the system controller will send input to IRU 001c01 blade 0 and the subchannel to be used is the console subchannel.

During the boot process, there is a window of time in which all processors may be 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 processor chosen to receive console input. You can turn filtering on and off with the `select filter` command.

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

```
001c01-L1>
```

```
entering console mode 001c01 console, <CTRL-T> to escape to L1  
no response from 001c01 junk bus console UART:UART_TIMEOUT
```

When this time-out condition occurs, either the IRU is hung or the subchannel is incorrect. An IRU is identified by its rack, type, and slot (001c01).

Viewing Information, Warnings, and Error Messages

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

```
001c01 ERROR: invalid arguments for 'ver' command, try "help ver"
```

The general format includes a IRU 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 Operation

As mentioned in “System Controller Interaction” on page 27 the system controller in an Altix 4700 system can provide L2 functionality.

An Ethernet cable can be plugged into the RJ45 connector on the IRU enclosure or on the Dense router. Connecting the IRU or Dense router to an active LAN via the L2 host connector will cause the system controller to spawn an L2. This connection provides network access to the system controller through the L2.

Configuring an L2's IP Address

This section refers to setting the IP address on the IRU enclosure when using an Ethernet connection.

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

Connect a serial cable to the serial console port on the target IRU and get the L1 prompt.

To see if the L2 is running (it will be if the LAN cable is plugged in and the cable is connected to an active LAN).

At the L1 prompt type:

```
001c01-L1> l2
L2 Controller is running.
001c01-L1>
```

If the L2 is not running type:

```
001c01-L1>! init 4
```

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

Verify the L2 is running again as above.

To set the IP address on the L2 type:

```
001c01-L1> l2 ip a.b.c.d 255.255.255.0 a.b.c.255
```

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

```
001c01-L1> l2 serial
```

To set the L2 system serial number:

```
001c01-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)

```
001c01-L1> l2 msys
```

If `msys` is off, turn it on:

```
001c01-L1> l2 msys on
```

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

```
001c01-L1> reboot_11
```

Once this is done for all target IRUs and Dense routers, connect them to the network (using an optional Ethernet switch if necessary).

The `rackid` on the L2 cannot be set with the L2 “`rackid`” command, instead it will be inherited from the local L1. As an example: the L2 running on the system controller in 1r41 will have a rack id of 141 (rack * 100 + slot of the local L1).

Once the L2 is running, you can telnet to the L2 using an optional LAN connected console system.

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

```
olympic-101-L2>
```

Common operations are discussed in the subsections that follow.

Viewing System Configuration

You can use the L2's `config` command to view the current system configuration from an IRU level:

```
olympic-101-L2> config
L2 127.0.0.1: - 001 (LOCAL)
L1 127.0.0.1:0:0 - 001c31
L1 127.0.0.1:0:1 - 001c21
L1 127.0.0.1:0:2 - 001c11
L1 127.0.0.1:0:3 - 001c01
L2>
```

As shown above, `config` produces a list of IRUs and their locations in the system and the system controller address of each IRU and Dense router. This is similar to the output from using the `config` command on the L1 with the addition of the L2 IP address, L1 connection, and L1 index.

The structure of the IRU and Dense router address is as follows:

a.b.c.d:x:y

where:

a.b.c.d is the IP address of the L2. (In the example above, the IP address is 127.0.0.1.)

x connection number is only 0 for Altix 4700

y is the L1 index, as follows:

0 is the local IRU

A number greater than 0 indicates that it is attached directly to or indirectly to the local IRU. An IRU is identified by its rack, type, and slot (001c01). The structure of the IRU location is as follows:

rrrbss.p

where:

rrr is the rack number.

b is the enclosure type.

ss is the slot location of the enclosure.

p is the partition of the enclosure (not present if the system is not partitioned).

In the example shown above, 001c01 is an IRU in rack 001 and slot position 01.

Setting Command Targeting

If a command is not understood by the L2 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 IRU 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 IRU 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 IRU or Dense router.

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 IRUs in a partition, enter the following:

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

Individual IRU and Dense routers can also be targeted with a short <rack>.<slot> prefix. As in 1.11 <command>.

To target individual blades in an IRU use the following syntax:

```
olympic-101-L2> 1.11 b1 power down
```

Executing the above command will power down the blade in blade slot 1 of the IRU in rack 001 U position 11.

Default Destination

When the L2 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 IRUs are found in the default destination. If you enter a command not understood by the L2, the command is sent to these IRUs.

Note: In the current implementation, if you add an IRU 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
001c01 (127.0.0.1:0:2)
001c01 (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 IRUs 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 provides sequencing that is necessary for a command to function correctly.

When L1 and L2 commands are similar, you can ensure that an L1 command is entered for the IRUs 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:

```
001c01 ERROR: invalid arguments for 'ver' command, try "help ver"
```

The general format includes an IRU 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

You can power on and power off the system with the `power` command. This command is interpreted by the L2, because the IRUs 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 IRU, specify a current destination:

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

To power on or power off all IRUs 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 data in memory will be lost.

Console Mode from the L2

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 (001c01 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 (001c01 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 communicates with the IRU set with the `select` command to be the system console or global master. All input from the console is directed to that IRU. You can set and view the system console with the `select` command.

The L2 chooses an IRU as the default console in the following order of priority:

- The IRU in the lowest numbered rack and slot, which has previously produced console output.
- The IRU 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)

    001c01-L2 detected

current system console

console input: 001c01 CPU 0A
console output: not filtered
```

The following are ten common subchannels associated with console communications:

- Subchannel 0A specifies Blade 0, CPU A.
- Subchannel 0C specifies Blade 0 CPU C.
- Subchannel 1A specifies Blade 1, CPU A.
- Subchannel 1C specifies Blade 1, CPU C.
- Subchannel 2A specifies Blade 2, CPU A.
- Subchannel 2C specifies Blade 2, CPU C.
- Subchannel 3A specifies Blade 3, CPU A.
- Subchannel 3C specifies Blade 3, CPU C.
- Subchannel console0 Blade 0 console subchannel.
- Subchannel console1 Blade 1 console subchannel.

The select command output: “console input: 001c01 console0” shows that the L2 will send console input to IRU 001c01 blade 0 and the console subchannel will be used.

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

```
L2> select 1.1
console input: 001c01 console
console output: no filtered
console detection: L2 detected
```

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

```
L2> select subchannel 1A
console input: 001c01 console CPU1A
console output: not filtered
```

During the boot process on a system with multiple CPUs, there is a window of time in which the CPUs are all producing output. This can result in a somewhat jumbled output from the L2. However, you can filter console output so that the L2 will show output from only the IRU 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 an IRU chosen to receive console input but that is not responding, a time-out condition results:

```
L2> Ctrl+D
entering console mode 001c01 CPU1A, <CTRL_T> to escape to L2

no response from 001c01 Junk bus CPU1A system not responding
no response from 001c01 Junk bus CPU1A system not responding
```

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

L1 Mode From L2

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 11 :

```
L2> r 2 s 1 l1
```

An alternate method is:

```
L2> 2.1 l1
```

```
enterling L1 mode 001c01, <CTRL-T> to escape to L2  
001c01-L1>
```

To return to L2 mode, press Ctrl+T:

```
001c01-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 returns to L1 mode:

```
re-entering L1 mode 001c01, <CTRL-T> to escape to L2  
001c01-L1>
```

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

```
001c01-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 an Ethernet connection (LAN console option) is required to execute the commands described in this section. See “Console Hardware Requirements” on page 30 and Figure 2-1 on page 26 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 IRU or by using the L1 command `reboot_l1`.

Typically, you will upgrade the firmware through the network connection from the LAN console workstation 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 IRUs 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 IRUs by replacing *all* with a rack and slot number:

```
$> /usr/cpu/firmware/sysco/flashsc --12 10.1.1.1  
/usr/cpu/firmware/sysco/11.bin 1.19
```

This updates only the IRU in rack 1, slot 19.

System Overview

This chapter provides an overview of the physical and architectural aspects of your SGI Altix 4700 series system. The major components of the Altix 4700 series systems are described and illustrated.

The Altix 4700 series is a family of multiprocessor distributed shared memory (DSM) computer systems that initially scale from 16 to 1,024 Intel 64-bit processor cores (8 to 512 sockets) as a cache-coherent single system image (SSI). Future releases will scale to larger processor counts for single system image (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 4700 series system in your lab or server room. Each 42U SGI rack holds from one to four 10U high enclosures that support up to ten compute/memory and I/O sub modules known as “blades.” These blades are single printed circuit boards (PCBs) with ASICS, processors, memory components and I/O chipsets mounted on a mechanical carrier. The blades slide directly in and out of the Altix 4700 IRU enclosures.

This chapter consists of the following sections:

- “System Models” on page 50
- “System Architecture” on page 52
- “System Features” on page 55
- “System Components” on page 59

Figure 3-1 shows the front views of a multiple-rack Altix 4700 system.




Figure 3-1 SGI Altix 4700 Systems

System Models

The basic enclosure within the Altix 4700 system is the 10U high (17.5 inch or 44.45 cm) “individual rack unit” (IRU). The IRU enclosure contains up to ten single-wide blades or two double-wide blades and eight single-wide blades. Each IRU comes with four high-speed routers. The routers connect to the installed blades via a backplane. Each router also has two ports that are brought out to external NUMalink 4 connectors. The 42U rack for this server houses all IRU enclosures, option modules, and other components; up to 64 processor sockets (128 processor cores) in a single rack. The Altix 4700 server system can expand up to 1,024 Intel 64-bit processor cores; a minimum of one IA blade (base I/O) is required for every 1,024 processor cores.

Figure 3-2 shows an example configuration of a single-rack Altix 4700 server.




Figure 3-2 SGI Altix 4700 Server System (Single Rack)

The system requires a minimum of one 42U tall rack with one single-phase power distribution unit (PDU) per IRU installed in the rack. Each single-phase PDU has 5 outlets (four of which are required to support the four power supplies in each IRU).

The three-phase PDU has 18 outlets (16 connections are required to support the four IRUs that may be installed in the rack).

You can also add additional PCI expansion blades or RAID and non-RAID disk storage to your server system.

Figure 3-3 shows an individual blade, IRU and Rack.

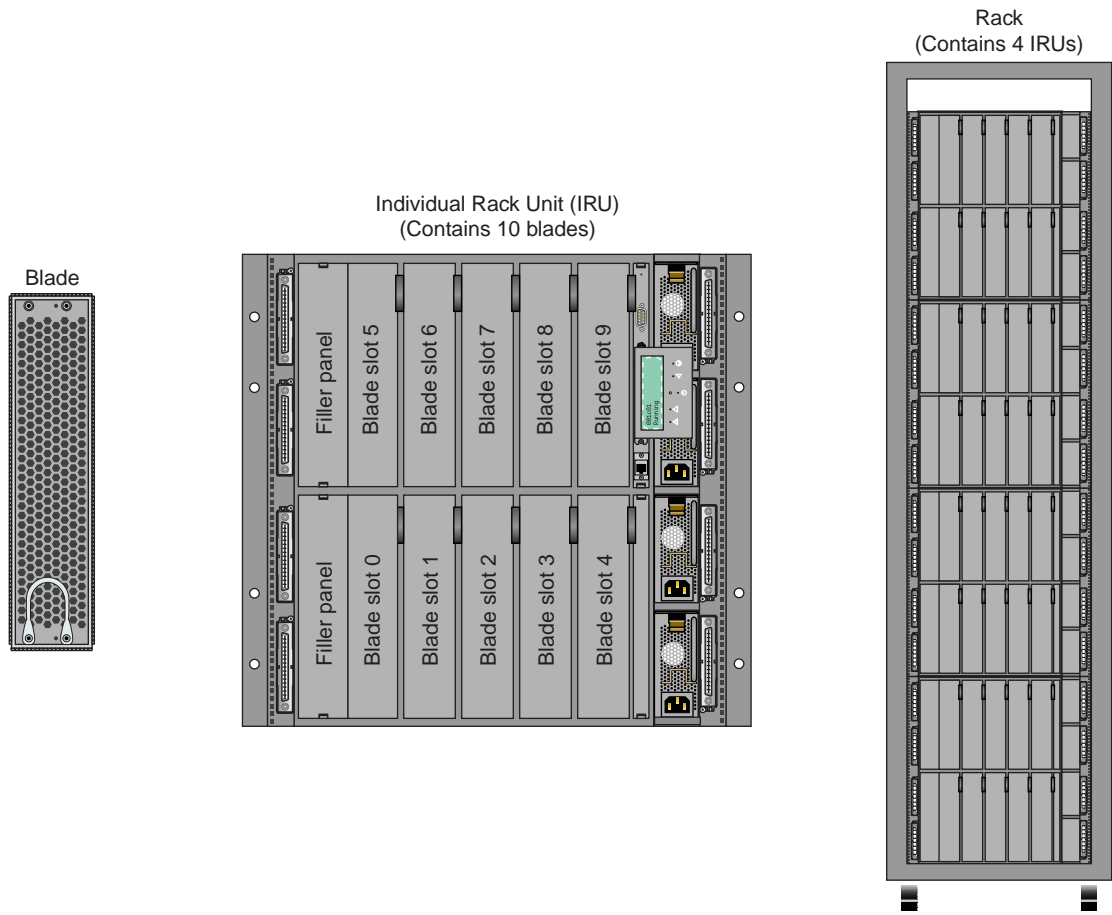


Figure 3-3 Blade, IRU and Rack Components

System Architecture

The Altix 4700 computer system is based on a distributed shared memory (DSM) architecture. The system uses a global-address-space, cache-coherent multiprocessor that scales up to sixty four Intel 64-bit processors in a single rack. Because it is modular, the DSM combines the advantages of lower entry cost with the ability to scale processors, memory, and I/O independently to a maximum of 512 processor sockets (1,024 processor cores) on a single-system image (SSI).

Larger SSI configurations may be offered in the future, contact your SGI sales or service representative for information.

The system architecture for the Altix 4700 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 four router switches in each 10U IRU enclosure.

The basic expansion building block for the NUMAlink interconnect is the processor node; each processor node consists of a Super-Hub (SHub) ASIC and one or two 64-bit processors with three levels of on-chip secondary caches. The 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 processor and memory node blade technology. This specialized ASIC acts as a crossbar between the processors, local SDRAM memory, and the network interface. The SHub ASIC memory interface enables any processor in the system to access the memory of all processors 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. Using the router ASICs with a highly specialized backplane or NUMAlink 4 cables provides a high-bandwidth, extremely low-latency interconnect between all processor, I/O, and other option blades within the system.

Figure 3-4 on page 54 shows a functional block diagram of the Altix 4700 series system IRU processor blades, NUMAlink interface, and other major components.

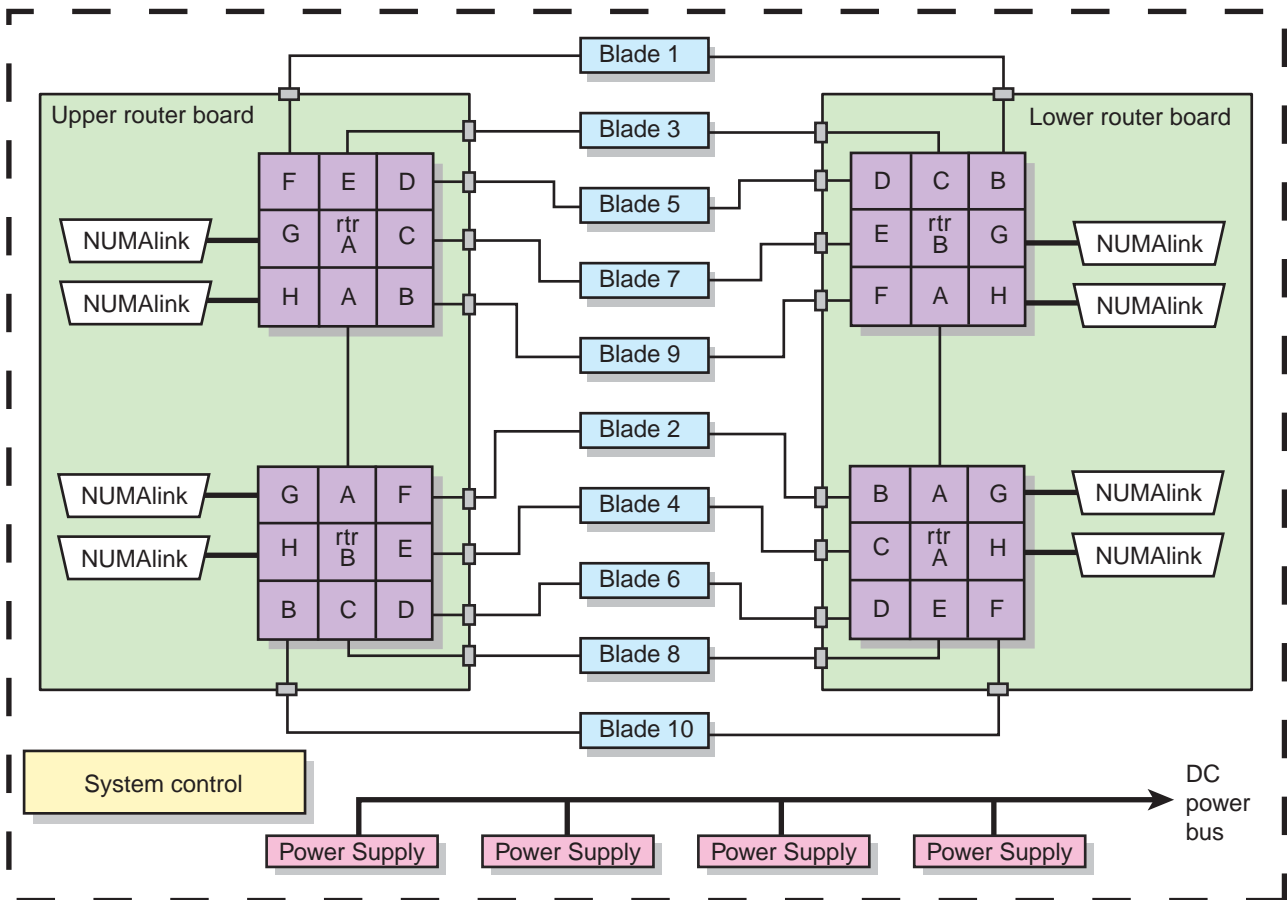


Figure 3-4 Functional Block Diagram of the Individual Rack Unit (IRU)

System Features

The main features of the Altix 4700 series server systems are introduced in the following sections:

- “Modularity and Scalability” on page 55
- “Distributed Shared Memory (DSM)” on page 55
- “Distributed Shared I/O” on page 56
- “Reliability, Availability, and Serviceability (RAS)” on page 57

Modularity and Scalability

The Altix 4700 series systems are modular systems. The components are primarily housed in building blocks referred to as individual rack units (IRUs). Additional option mass storage may be added to the rack along with additional IRUs. You can add different types of blade options to a system IRU 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 blades that create the basic functionality (compute/memory, I/O, and power) into custom 19-inch racks. The air-cooled IRU enclosure system has redundant, hot-swap fans and redundant, hot-swap power supplies at the IRU level.

Distributed Shared Memory (DSM)

In the Altix 4700 series server, memory is physically distributed both within and among the IRU enclosures (compute/memory/I/O blades); however, it is accessible to and shared by all NUMAlinked devices within the single-system image. This is to say that all NUMAlinked components sharing a single Linux operating system, operate and share the memory “fabric” of the system.

Note the following sub-types of memory within a system:

- If a processor accesses memory that is connected to the same SHub ASIC on a compute node blade, the memory is referred to as the node’s *local memory*.
- If processors access memory located in other blade nodes within the IRU, (or other NUMAlinked IRUs) the memory is referred to as *remote memory*.
- The total memory within the NUMAlinked 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 blade nodes within the IRUs (each base I/O blade node has two NUMAlink ports) and are accessible by all compute nodes within the SSI 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 4700 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/memory blade 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 4700 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 blade to reference its locally installed memory than to reference remote memory.

Reliability, Availability, and Serviceability (RAS)

The Altix 4700 server series components have the following features to increase the reliability, availability, and serviceability (RAS) of the systems.

- **Power and cooling:**
 - IRU power supplies are redundant and can be hot-swapped under most circumstances. Note that this might not be possible in a “fully loaded” system. If all the blade positions are filled, be sure to consult with a service technician before removing a power supply while the system is running.
 - IRUs have overcurrent protection at the blade and power supply level.
 - Fans are redundant and can be hot-swapped.
 - Fans run at multiple speeds in the IRUs. Speed increases automatically when temperature increases or when a single fan fails.
- **System monitoring:**
 - System controllers monitor the internal power and temperature of the IRUs, and can automatically shut down an enclosure 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 IRU and each blade/node installed has failure LEDs that indicate the failed part; LEDs are readable at the front of the IRU or 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 have a local field-replaceable unit (FRU) analyzer.
 - All system faults are logged in files.
 - Memory can be scrubbed using error checking code (ECC) when a single-bit error occurs.

System Components

The Altix 4700 series system features the following major components:

- **42U rack.** This is a custom rack used for both the compute and I/O rack in the Altix 4700 system. Up to 4 IRUs can be installed in each rack. There is also a 2U space reserved at the top for an optional router.
- **Individual Rack Unit (IRU).** This enclosure contains the compute/memory blades, IA2 blade (base I/O), standard routers and optional I/O blades for the Altix 4700. The enclosure is 10U high. Figure 3-5 on page 60 shows the Altix 4700 IRU system components.
- **Compute/Memory blade.** Holds up to two IA-64 processor sockets and 4, 8 or 12 memory DIMMs.
- **Memory-only blade.** This blade acts as a memory expansion node with no processor compute circuitry included on the blade. This blade holds 4, 8, or 12 memory DIMMs.
- **Double-wide PCI-X expansion blade.** Supports three PCI/PCI-X 133 MHz 64-bit option cards. This three-slot blade features card carriers that allow you to slide PCI/PCI-X boards directly in and out of the unit.
- **PCIe expansion blade.** The single wide PCI blade supports two PCI Express option cards.
- **Double-wide PCIe/PCI-X expansion blade.** This blade supports two PCI Express option cards and two PCI/PCI-X option cards.

Note: PCIe options may be limited, check with your SGI sales or support representative.

- **IA2 blade (Base I/O blade).** Double-wide I/O blade that supports all base system I/O functions including one or two disk drives with optional RAID 0 or 1, a DVD-R/W drive, two low-profile PCI card slots, two ethernet connectors, one SAS/SATAII port, and four USB ports. The enhanced IA blade (also known as the IA2 blade) began shipping as original equipment in most systems ordered after January 2007. You must have the IA2 version of the blade and SLES 10 or later to use RAID 1 mirroring on your system disk pair.

Note: While the system I/O blade is capable of RAID 0 support, SGI does not recommend the end user configure it in this way. RAID 0 offers no fault tolerance to the system disks, and a decrease in overall system reliability. In a RAID 0 configuration, failure of either system disk will result in data being lost on both disks, resulting in system failure. The Altix 4700 ships with RAID 1 functionality (disk mirroring) configured if the option is ordered.




Figure 3-5 Altix 4700 IRU 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 IRUs occupy multiple standard units, IRU locations within a rack are identified by the bottom unit (U) in which the IRU resides. For example, in a 42U rack, an IRU positioned in U01 through U10 is identified as U01.

Rack Numbering

Each rack is numbered with a three-digit number sequentially beginning with 001. A rack contains IRU enclosures, optional mass storage enclosures, optional router bricks and potentially other options. In a single compute rack system, the rack number is always 001.

Optional System Components

Availability of optional components for the SGI 4700 systems may vary based on new product introductions or end-of-life components. Some options are listed in this manual, others may be introduced after this document goes to production status. Check with your SGI sales or support representative for the most current information on available product options not discussed in this manual.

Rack Information

This chapter describes the physical characteristics of the tall (42U) racks in the following sections:

- “Overview” on page 63
- “Tall Altix Rack (42U)” on page 64
- “Technical Specifications” on page 66

Overview

At the time this document was published only the tall (42U) Altix rack shown in Figure 4-1 was approved for use with the Altix 4700 systems.

Tall Altix Rack (42U)

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

- **Front and rear door.** The front door is opened by grasping the wide end of the triangle-shaped door piece and pulling outward. It uses a key lock for security purposes that should open all the front doors in a multi-rack system. Note that the front door and rear door locks are keyed differently.

The rear door has a push-button key lock to prevent unauthorized access to the system. The rear doors have a master key that locks and unlocks all rear doors in a system made up of multiple racks. You cannot use the rear door key to secure the front door lock.

- **Cable entry/exit area.** Cable access openings are located in the front floor and top of the rack. Cables are only attached to the front of the IRUs; therefore, cable management occurs primarily in the front and top of the rack. In the case of Dense router modules, cables are attached and routed to the top of the rack. NUMAlink cables pass through the top of the rack. I/O and power cables normally pass through the bottom of the rack.
- **Rack structural features.** The rack is mounted on four casters; the two rear casters swivel. There are four leveling pads available at the base of the rack. The base of the rack also has attachment points to support an optional ground strap, and/or seismic tie-downs.
- **Power distribution units within the rack.** A maximum of four single-phase power distribution units (PDUs) are needed per rack, depending on the number of IRUs. Each IRU requires one 5-outlet single-phase PDU. A three-phase power distribution unit has 18 outlet connections, of which up to 16 may be needed to power four IRUs. Each three-phase rack system houses a single PDU.




Figure 4-1 Tall (42U) Altix Rack and Front Lock

Technical Specifications

Table 4-1 lists the technical specifications of the Altix 4700 tall rack.

Table 4-1 Tall Altix Rack Technical Specifications

Characteristic	Specification
Height	79.5 in. (201.9 cm)
Width	25.8 in. (65.5 cm)
Depth	43.45 in. (110.4 cm)
Weight (full)	1500 lb (682 kg)

Optional Dense Router Module

This chapter describes the function and physical components of the optional Dense router module in the following sections:

- “Overview” on page 67
- “External Components” on page 69
- “Technical Specifications” on page 72

Overview

The Dense router module features four internal router boards and functions as a high-speed switch to route network packets from one individual rack unit (IRU) to another throughout the NUMALink interconnect fabric of the server system. The router module is normally only used when your system IRUs occupy more than one rack. The key components within the Dense router module are the router chips, these are SGI custom-designed ASICs. Each of the four router boards inside the module uses a router ASIC. The router chip is an eight-port crossbar that connects any input-link channel to any of the seven other possible output-link channels supported by each router printed circuit board, see Figure 5-1 for a conceptual diagram. The Dense router module contains four individual router boards. Each router board supports eight NUMALink channels.

The Dense router module has the following features:

- 32 NUMALink channels
- One L1 controller and one Ethernet port for system controller support
- One 9-pin RS-232 type console port
- Four hot-pluggable cooling fans




Figure 5-1 Dense Router Functional Block Diagram

External Components

This section describes the external components that are located on the front and rear panels of the Dense router module. Note that the router is normally installed at the top of the rack.

The Dense router requires a power cable connection to a power module (also installed at the top of the rack).

Figure 5-2 shows an example front and rear view of a Dense router.



Figure 5-2 Dense Router Module Front and Rear Panels

Front Panel Components

The Dense router module contains the following front panel items (as shown in Figure 5-3 on page 70):

- **Sixteen NUMalink 4 connection ports.** These ports are used with systems that require more than one rack and utilize NUMalink interconnects beyond those provided on the individual IRUs.
- **Power connector port.** This port is the power input for the Dense router module. It must be connected to the 1350 W redundant power supply module located in the top section of the rack. The power supply provides 12-VDC and 3.3-VDC to the Dense router.

- **Two LEDs:**
 - On/Off LED. This LED illuminates green when the router module internal components are on and turns off when they are off.
 - Failure LED. This LED illuminates red to indicate a Dense router failure.
- **Fans.** Four 40mm fans provide N+1 redundant cooling.

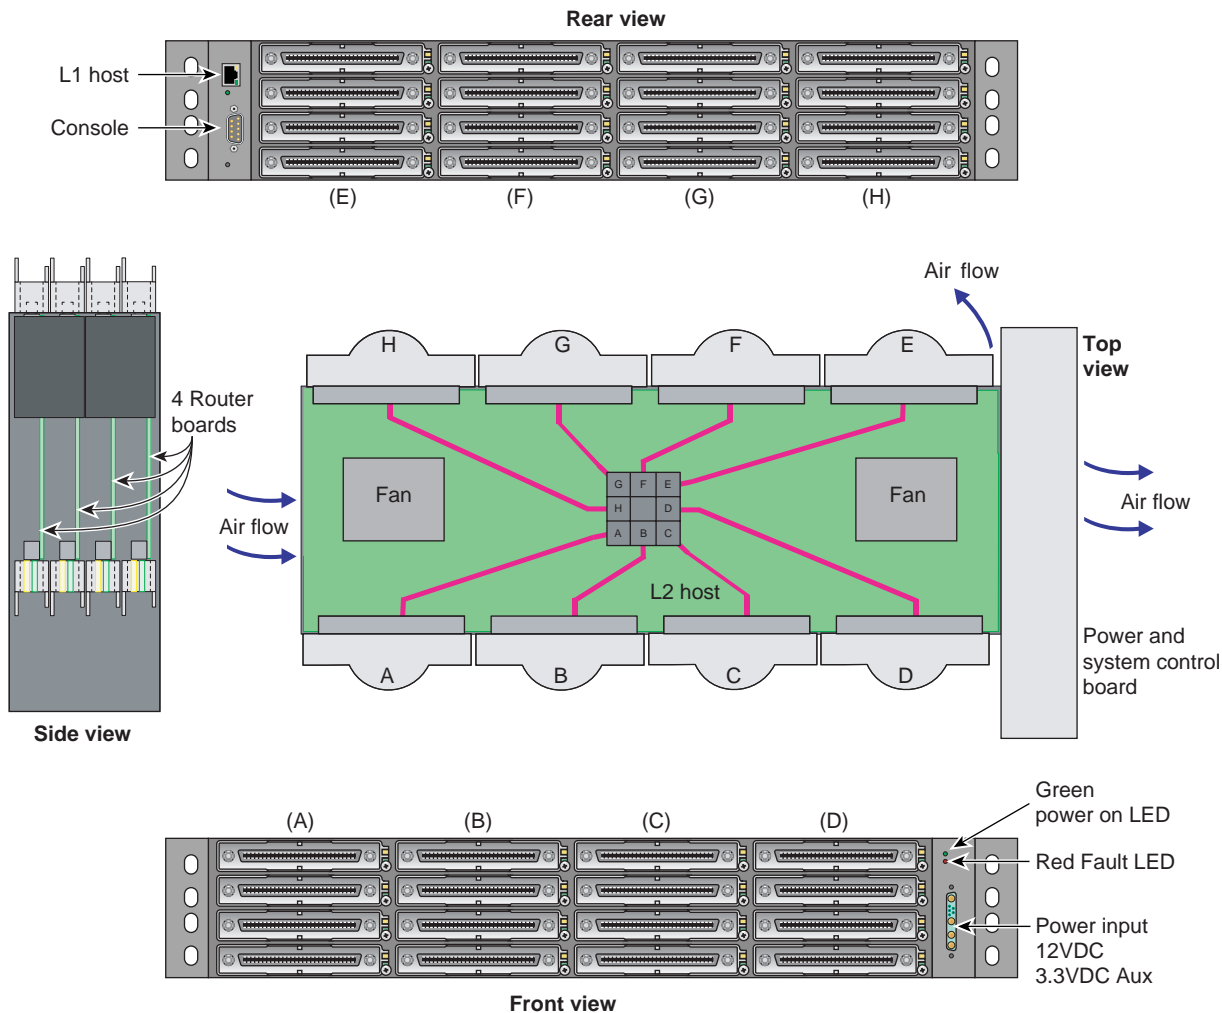


Figure 5-3 Dense Router Functionality

Rear Panel Components

The Dense router module has the following rear panel items (see Figure 5-3 on page 70):

- **L2 host connector.** The RJ-45 connector on the router provides an Ethernet connection from the Dense router's system controller to the end-user's local area network (LAN), or directly to a system console, (PC or laptop).
- **9-pin console connector.** This serial RS-232 style connection provides a console connection point for a "dumb terminal". You can connect the terminal via a serial cable to the (DB-9) console port connector on the back of the Dense router. The terminal should 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)

Note that the serial console is normally connected to the first (bottom) IRU in any single rack configuration.

- **Link connector LEDs.** Each NUMAlink connector has two LEDs, as follows:
 - The yellow LED illuminates to indicate that both the router module's port and the port to which it is connected are powered on.
 - The green LED illuminates when a connecting link has been established between the router module and the port to which it is connected.

Technical Specifications

Table 5-1 lists the technical specifications of the four-board router module.

Table 5-1 Dense Router Technical Specifications

Characteristic	Specification
Height	(2U) 3.35 in. (8.5 cm)
Width	17.38 in. (44.1 cm)
Depth	7.13 in. (18.1 cm)
Weight	14.5 lb (6.6 kg)
Input power	12VDC/3.3VDC (~ 160 W)

Table 5-2 lists the specifications of the Dense router ports.

Table 5-2 Dense Router Port Specifications

Port	Quantity	Peak Transfer Rate
NUMAlink	32	3.2 GB/s each direction
L2 host (RJ45)	1	12 Mbits/s
DB-9 serial console	1	240 Kbits/s

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 73
- “Installing or Removing Internal Parts” on page 75
- “Adding or Replacing PCI/PCI-X Cards” on page 94
- “Installing or Replacing a Disk Drive in the IA/IA2 Blade” on page 104

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 74
- “Returning the System to Operation” on page 74
- “Installing or Removing Internal Parts” on page 75

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. Follow standard procedures for gracefully halting the operating system.
2. At your system console, switch from console mode to system control mode by entering the following command:

```
$> Ctrl+T
```

The system will respond with the system control prompt:

```
L2>
```

3. From the system control prompt (L2>), power off the system with the following command:
L2> **power down**
4. Locate the power distribution unit(s) (PDUs) in the front 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 PDU circuit breaker switches to the “on” position.
2. At your system console, enter the following command:
L2> **power up**
3. Verify that the LEDs on the L1/L2 controller(s) 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 7, 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 6-1. Components not listed must be installed or removed by a qualified SGI field engineer.

Table 6-1 lists the customer-replaceable components and the page on which you can find the instructions for installing or removing the component.

Table 6-1 Customer-replaceable Components and Maintenance Procedures

Component	Procedure
L1 controller panel	“Remove and Replace the System Control Display Panel” on page 76
IRU power supplies	“Removing or Adding IRU Power Supplies” on page 79
IRU fans (blowers)	“Removing and Replacing IRU Fans (Blowers)” on page 80
Replacing IRU blades	“Removing and Replacing a Compute/Memory Blade” on page 84
Add or replace DIMMs	“Memory DIMM Group Placement and Guidelines” on page 87
Replace the DVD drive	“Remove and Replace the DVD Drive” on page 90
PCI, PCI-X, PCIe cards	“Adding or Replacing PCI/PCI-X Cards” on page 94
2D Gfx PCI card in base I/O	“Add or Remove a PCI Card in Base IA Blades” on page 102
Base I/O blade disk drives	“Installing or Replacing a Disk Drive in the IA/IA2 Blade” on page 104

Replacing IRU Components

While the individual rack units (IRUs) are not considered end-user replaceable, a number of their internal components can be removed and replaced. These include:

- L1 controller panel
- IRU power supplies
- IRU cooling fans (also called blowers)
- IRU blades and their components

Remove and Replace the System Control Display Panel

To remove and replace either of the top two power supplies in an IRU, you must first remove the system control display panel. Use the following steps to remove the system control display:

1. Using a T10 torx driver, loosen the lower screw that holds the display panel to the IRU chassis, see Figure 6-1 on page 77.
2. Lift the display assembly up about 1/3-inch (84 mm) and pull it outward from the IRU chassis.
3. Don't unplug the ribbon cable unless you are replacing the system control panel.
4. Let the control panel gently dangle and go to the section "Removing or Adding IRU Power Supplies" on page 79.
5. Replace the applicable power supply and then use the following set of steps to replace the system control panel.

Replace the system control panel using the following steps:

1. Position the controller panel assembly slightly above the retention screw and carefully slid it back into place, see Figure 6-2 on page 78. Be careful not to snag the ribbon cable while replacing the panel.
2. Use the T10 torx driver to tighten the retention screw that holds the panel to the IRU chassis.
3. Restart the IRU or system as applicable.




Figure 6-1 Removing the L1 Controller Panel




Figure 6-2 Replacing the L1 Controller Panel

Removing or Adding IRU Power Supplies

Under most circumstances a single power supply in an IRU can be replaced without shutting down the IRU or the complete system. In the case of a fully configured (loaded) IRU, this may not be possible. In a fully loaded system, power supply failure could result in shutdown of the IRU due to the power loss. Note that when replacing either of the top two supplies, it is necessary to remove the system controller panel (see “Remove and Replace the System Control Display Panel” on page 76 before proceeding).

Use the following steps to replace a power supply in the IRU:

1. Disconnect the power cord from the IRU power supply.
2. Press down on the retention latch at the top of the power supply, see Figure 6-3.
3. Using the power supply handle, pull the power supply straight out of the IRU chassis.




Figure 6-3 Removing an IRU Power Supply

4. Install the new power supply with the retention latch at the top of the supply.

5. Slide the power supply into the chassis until the retention latch engages.
6. Reconnect the power cord to the power supply.




Figure 6-4 Replacing an IRU Power Supply

Removing and Replacing IRU Fans (Blowers)

The IRU fans (blowers) are located at the rear of the unit (see Figure 6-5). You will need to access the rack from the back to remove and replace a fan. The IRU's system controller issues a warning message when a fan is not running properly. This means the fan RPM level is not within tolerance. When an IRU fan fails, the following things happen:

1. The L1 display on the IRU with the failed fan shows a warning message
`ATTN: Fan (number) warning`
2. The console will show the same warning indicating the rack and IRU position
`001c01 L2> Fan (number) warning limit reached @ 0 RPM`
3. A line will be added to the L1 system controller's log file indicating the fan warning.
4. If ESP is used, a warning message will be sent to it also, (see "SGI Electronic Support" on page 115).

The system controller monitors the temperature within each IRU. If the temperature increases due to the failed fan, the remaining five fans will run at a higher RPM to compensate for the missing fan. The system will continue running until a scheduled maintenance occurs. The fan numbers for the IRU (as viewed from the rear) are shown in Figure 6-5.




Figure 6-5 IRU Fan Assembly (Blowers)

Note that a fan can be replaced while the IRU is operating. You will need a number 1 phillips-head screw driver to complete the procedure.

Use the following steps and illustrations to replace an IRU fan:

1. Undo the (captive) screw (located in the middle of the blower assembly handle). The handle has a notch for the screw access, see Figure 6-6.
2. Grasp the blower assembly handle and pull the assembly straight out.




Figure 6-6 Removing a Fan From the IRU

3. Slide a new blower assembly completely into the open slot, see Figure 6-7.

4. Tighten the blower assembly screw to secure the new fan.




Figure 6-7 Replacing an IRU Fan

Removing or Replacing a System Blade From the IRU

The blades installed in your SGI Altix 4700 can generally be removed from or installed in an IRU without powering down the entire system. However, the operating system should be halted and the power to the IRU should be shut down in most cases. Note that the IA blade holding your system disk(s) should never be removed without backing up and gracefully powering down the unit.

Removing and Replacing a Compute/Memory Blade

Use the following information and illustrations to remove and replace a compute/memory blade from the host IRU:

1. Open the front door of the rack and access the IRU blade slated for replacement.
2. Move the blade's retention tab gently to the left to disengage the locking tab. Be careful not to use too much force as the tab is breakable.
3. Pull the blade's retention tab outward until it is fully extended, see Figure 6-8.
4. Slide the blade outward by pulling the bottom ring gently out.
5. Remove the blade completely from the IRU and place it on an anti-static work surface or in anti-static packaging.
6. Slide the new blade carefully into the IRU slot until fully seated.
7. Carefully push the retention tab back into place so that the blade is locked down, see Figure 6-9 for an example.
8. Restart the IRU or system.




Figure 6-8 Removing a Compute/Memory Blade From the IRU




Figure 6-9 Replacing a Compute/Memory Blade in the IRU

Memory DIMM Group Placement and Guidelines

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

These twelve DIMM slots are organized into three groups of four DIMMs each, as shown in Figure 6-10.

DIMMs are installed one per DIMM slot, and must be installed in groups of four. Figure 6-11 shows an example of the physical DIMM locations within the processor blade assembly.

Follow these guidelines when installing DIMM groups:

- Memory is increased or decreased in four-DIMM group increments only. Each of the four DIMMs that make up a DIMM group must be the same memory size and speed; however, different blades can differ in memory size.
- The first four DIMMs (group 0) must be in place for the blade's memory to operate properly.
- Mixing blades with different capacity DIMM groups within an IRU is permitted.
- When installing 4 GB DIMMs, check with your SGI service representative. Space limitations may require that a maximum of eight 4 GB DIMMs be installed on each blade.

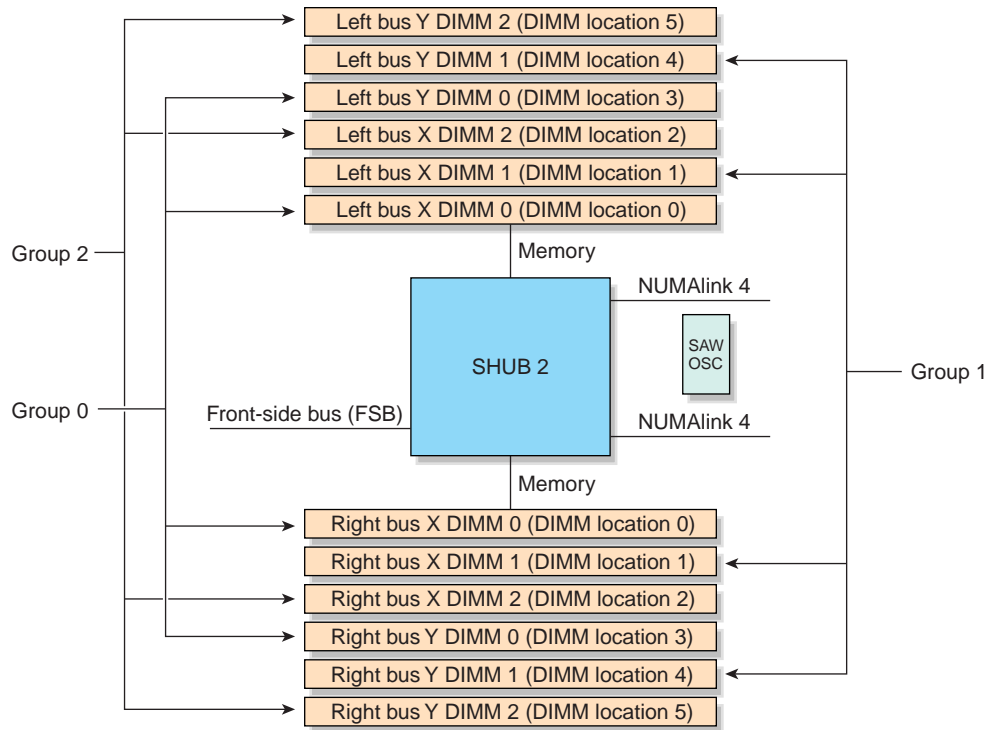


Figure 6-10 DIMM Group Locations Diagram




Figure 6-11 DIMM Memory Slot Location Example

Remove and Replace the DVD Drive

Use the following steps to remove the DVD drive from an IA blade:

1. Power off the IRU, use instructions in “Powering the System On and Off” on page 5.
2. Extract the IA blade you are going to install the new DVD drive in. See the information in “Replacing IRU Components” on page 76.
3. After removing the IA blade from the IRU chassis, place it on a stable flat surface with the left side facing upward (see Figure 6-12).
4. Use a Phillips type screwdriver to remove the DVD access panel screw, (see Figure 6-12).
5. Lift the DVD access panel clear of the IA blade and set it aside.
6. Unscrew and remove the DVD drive assembly retention screw.
7. Gently but firmly disengage the DVD power and data cable connectors from the old DVD drive and slide the DVD unit out of the IA blade housing.
8. Slide the new DVD drive assembly into the IA blade housing.
9. Connect the DVD power and data cable connectors to the replacement drive.
10. Screw in the DVD drive assembly retention screw.
11. Replace the DVD access panel and tighten the retention screw.
12. Return the IA blade to its IRU slot.
13. Power-up the IRU and/or system.




Figure 6-12 DVD Removal from IA Blade




Figure 6-13 DVD Drive Replacement in the IA Blade

Overview of PCI Express Operation

This section provides a brief overview of the PCI Express (PCIe) technology that will be available as an option with your system. PCI Express has both compatibility and differences with older PCI/PCI-X technology. Check with your SGI sales or service representative for more detail on PCI Express board options available with the SGI Altix 4700.

PCI Express is compatible with PCI/PCI-X in the following ways:

- Compatible software layers
- Compatible device driver models
- Same basic board form factors
- PCIe controlled devices appear the same as PCI/PCI-X devices to most software

PCI Express technology is different from PCI/PCI-X in the following ways:

- PCI Express uses a point-to-point serial interface vs. a shared parallel bus interface used in older PCI/PCI-X technology
- PCIe hardware connectors are not compatible with PCI/PCI-X (see Figure 6-14)
- Potential sustained throughput of x16 PCI Express is approximately four times that of the fastest PCI-X throughputs



Figure 6-14 Comparison of PCI/PCI-X Connector with PCI Express Connectors

PCI Express technology uses two pairs of wires for each transmit and receive connection (4 wires total). These four wires are generally referred to as a lane or x1 connection. This is also called “by

1”. SGI Altix 4700 PCIe technology uses a x16 connector (64 wires) or “by 16” in PCI Express card slots. This technology will support PCIe boards that use connectors up to x16 in size. Table 6-2 shows this concept.

Table 6-2 SGI PCIe Support Levels

SGI x16 PCIe Connector	
x1 PCIe cards	Supported
x2 PCIe cards	Supported
x4 PCIe cards	Supported
x8 PCIe cards	Supported
x16 PCIe cards	Supported
x32 PCIe cards	Not supported

Adding or Replacing PCI/PCI-X Cards



Warning: Before installing, operating, or servicing any part of this product, read the “Safety Information” on page 129.

This section provides instructions for adding or replacing a PCI/PCI-X card in the I/O blade or PCI expansion blades installed in your system’s Individual Rack Units (IRUs). To maximize the operating efficiency of your cards, be sure to read all the introductory matter before beginning the installation.

There are different types of blades used in the IRU that hold PCI/PCI-X and PCIe boards. Note the following guidelines:

- The IA/IA2 blade (base I/O blade) containing the system disk(s) and DVD will **only** accommodate half-height PCI/PCI-X cards.
- The three-slot double-wide PCI blade supports **only** PCI/PCI-X cards that mount and install with an SGI sliding mounting bracket.



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

Installing PCI Cards in the Three-Slot Blade

Use the following instructions and illustrations to install or replace a PCI/PCI-X card in the three-slot blade:

1. Follow the instructions in “Replacing IRU Components” on page 76 and extract the three-slot blade from the IRU.
2. To extract the card carrier from a slot in the blade, pull the handle on the selected card carrier straight out, as shown in Figure 6-15.
3. If you are replacing a card, gently remove the existing card from the carrier.
4. If you are adding a card, extract the carrier metal filler plate by pushing down on it, as shown in Figure 6-16. This filler plate covers the area where your card connectors will protrude out to the rear panel of the blade.




Figure 6-15 Removing a Card Carrier From the Three-Slot Blade




Figure 6-16 Extracting the Carrier Metal Filler Plate

Note: If you are adding or replacing a half-height PCI card, skip ahead 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 6-17).

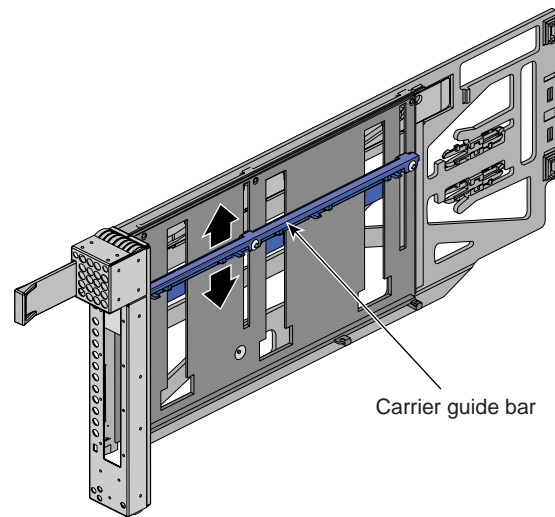


Figure 6-17 Adjusting the Carrier Guide Bar

- c. Mount the card in the carrier, (see Figure 6-18) and ensure 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. Skip step 6 and proceed to step 7.




Figure 6-18 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 6-19.

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 6-19).




Figure 6-19 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 6-20).




Figure 6-20 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 6-21.
- 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 6-21 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 6-22.




Figure 6-22 Installing a Card in a Slot

8. Push in horizontally on the carrier seating bar, as shown in Figure 6-23, to seat the card securely in the slot.




Figure 6-23 Seating the Card in the Slot

9. Install the blade into the IRU.
10. Power on the IRU or restart the system.
11. Run the `ls -pci` PCI hardware inventory command to verify the installation. This command lists PCI hardware that the operating system discovered during the boot operation.

Add or Remove a PCI Card in Base IA Blades

To add or replace a PCI card in the IA/IA2 blade (base I/O), follow these steps:

1. Power off the IRU by following the power-off instructions in the “Powering the System On and Off” on page 5 if you are removing the IA blade.
2. Extract the blade you are going to install the new PCI card in. See the information in “Replacing IRU Components” on page 76 if you have not already removed the blade. Note that in the case of extracting an IA blade that holds the system disk, you must shut down the operating system.
3. After removing the IA blade from the IRU chassis, place it on a stable flat surface with left side facing upward (see Figure 6-24).
4. Using a Tx20 Torx driver, remove the front screw securing the PCI card access door.
5. Using a Tx20 Torx driver, remove the top PCI card access door screw.
6. Remove the PCI card access door.
7. Add/replace/remove a PCI card(s) following standard ESD rules. Make sure any newly installed PCI card(s) are firmly seated in the slot and secure.




Figure 6-24 Removing the Card Access Panel From the IA Blade




Figure 6-25 Replacing the IA Blade PCI Access Door

8. Replace the PCI card access door by installing it into the slotted opening and sliding it forward, as shown in Figure 6-25.
9. Reinstall the top PCI card access door screw.
10. Reinstall the front PCI card access door screw.
11. Reinstall the IA blade in the IRU chassis.
12. Return the system to service.
13. Boot your operating system software. (See your software operation guide if you need instructions to boot your operating system.)
14. Run the `ls -pci` PCI hardware inventory command to verify the installation. This command lists PCI hardware that the operating system discovered during the boot operation.

Installing or Replacing a Disk Drive in the IA/IA2 Blade

The IA/IA2 blade (base I/O) in an IRU has two system disk drive bays as shown in Figure 6-26. This section describes how to install or remove the drives. The IA2 base I/O blade is an enhanced version of the IA that supports RAID 1, DVD-R/W, and DVI/VGA graphics. IA2 blades are available on most systems ordered starting January 2007.




Figure 6-26 Location of Disk Drive Bays

Removing a System Disk Drive

Before you remove a disk drive, you must power off the IRU or system according to the instructions in the “Powering the System On and Off” on page 5.

To remove a disk drive, follow these steps:

1. Open the front of the rack and access the two disk drive bays, as shown in panel 1 of Figure 6-27.
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 blade’s backplane connector.
3. Carefully slide the disk drive out of the blade and place it on an ESD-safe surface. Do not use the handle to pull the disk drive out of the bay.

After you remove the disk drive, go to the next section to install a new one, or, power on the IRU according to the instructions in the “Powering the System On and Off” on page 5.

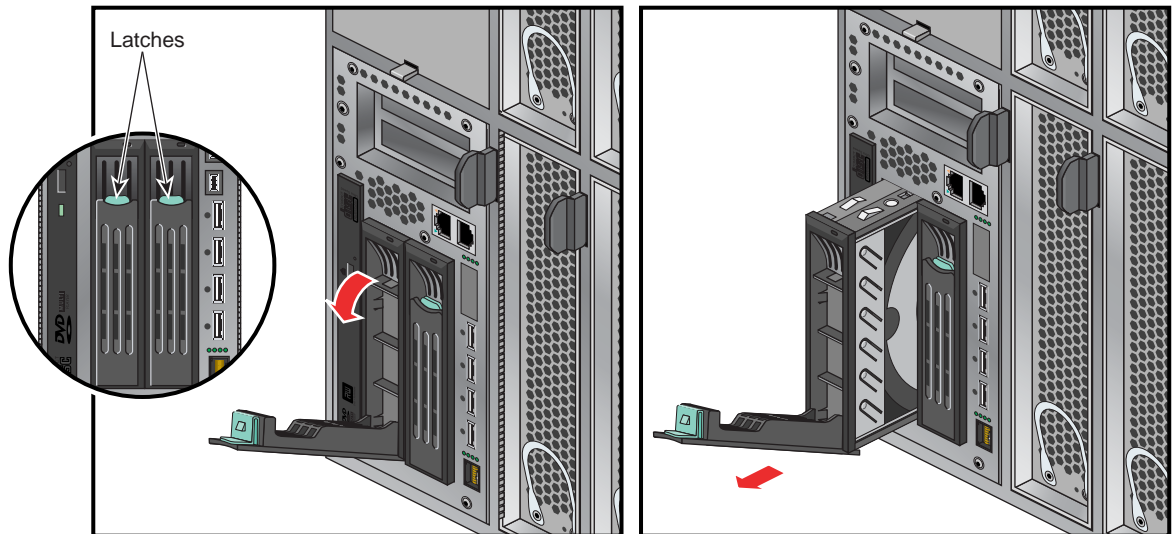


Figure 6-27 Removing a Disk Drive

Installing a System Disk Drive

Before you install a disk drive, you must power off the IRU that houses the IA blade (base I/O) where the disk(s) are installed, see the instructions in the “Powering the System On and Off” on page 5.

Note that you should always back up all files before replacement of the disk. Be sure to replace any Altix 4700 hard disk with an SGI approved drive.



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 6-28 and follow these steps:

1. Access the two disk drive bays (see panel 1 in Figure 6-28).

Note: If you are installing only one disk drive, install it in the left-side housing nearest the DVD drive.

2. Position the drive assembly so that it engages the bay guide rails, and then gently push the drive into the blade. 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.

After you finish installing the disk drive, power on the IRU or system according to the instructions in the “Powering the System On and Off” on page 5.




Figure 6-28 Replacing a Disk Drive

Troubleshooting and Diagnostics

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

- “Troubleshooting Chart” on page 110
- “L1 Controller Error Messages” on page 111
- “SGI Electronic Support” on page 115

Troubleshooting Chart

Table 7-1 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 115.

Table 7-1 Troubleshooting Chart

Problem Description	Recommended Action
The system will not power on.	<p>Ensure that the power cords of the IRU are seated properly in the power receptacles.</p> <p>Ensure that the PDU circuit breakers are on and properly connected to the wall source.</p> <p>If the power cord is plugged in and the circuit breaker is on, contact your SSE.</p>
An individual IRU will not power on.	<p>Ensure the power cables of the IRU are plugged in.</p> <p>View the L1 display; see Table 7-2 if an error message is present.</p> <p>If the L1 controller is not running, contact your SSE.</p>
The system will not boot the operating system.	Contact your SSE.
The Service Required LED illuminates on an IRU.	View the L1 display of the failing IRU; see Table 7-2 for a description of the error message.
The Failure LED illuminates on an IRU.	View the L1 display of the failing IRU; see Table 7-2 for a description of the error message.
The green or yellow LED of a NUMAlink port is not illuminated.	Ensure that the NUMAlink cable is seated properly on both ends.
The PWR LED of a populated PCI slot is not illuminated.	Reseat the PCI card.
The Fault LED of a populated PCI slot is illuminated (on).	Reseat the PCI card. If the fault LED remains on, replace the PCI card.
The amber LED of a disk drive is on.	Replace the disk drive.

L1 Controller Error Messages

Table 7-2 lists error messages that the L1 controller generates and displays on the L1 display. This display is located on the front of the IRUs.

Note: In Table 7-2, 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 7-2 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 IRU.
ATTN: <power VRM description> low fault limit reached @ x.xxV	30-second power-off sequence for the IRU.
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 7-2 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 IRU's air inlet has exceeded 30° C.
ATTN: <temp sensor description> critical temperature reached @ xxC xxF	The ambient temperature at the IRU's air inlet has exceeded 35 °C.
ATTN: <temp sensor description> fault temperature reached @ xxC xxF	The ambient temperature at the IRU'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 IRU's air inlet has exceeded 27 °C.
ATTN: <temp sensor description> critical temperature reached @ xxC xxF	The ambient temperature at the IRU's air inlet has exceeded 31 °C.
ATTN: <temp sensor description> fault temperature reached @ xxC xxF	The ambient temperature at the IRU's air inlet has exceeded 35 °C.
Temperature stable message:	
ATTN: <temp sensor description> stabilized @ xxC/xxF	The ambient temperature at the IRU'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.
IRU appears to have been powered down	The L1 controller has registered a fault and has shut down.

LED Status Indicators

There are a number of LEDs on the front of the IRUs that can help you detect, identify and potentially correct functional interruptions in the system. The following subsections describe these LEDs and ways to use them to understand potential problem areas.

IRU Power Supply LEDs

Each power supply installed in an IRU has a single bi-color (green/amber) status LED. The LED will either light green or amber (yellow), or flash green or yellow to indicate the status of the individual supply. See Table 7-3 for a complete list.

Table 7-3 Power Supply LED States

Power supply status	Green LED	Amber LED
No AC power to the supply	Off	Off
Power supply has failed	Off	On
Power supply problem warning	Off	Blinking
AC available to supply (standby) but IRU is off	Blinking	Off
Power supply on (IRU on)	On	Off

IRU NUMAlink Router Port LEDs

Each IRU supports a total of eight NUMAlink connectors (located on the front of the unit). Each of these connectors has two status LEDs (one green and one amber).

- The amber LED illuminates to indicate that both the NUMAlink module and the module to which it is connected are powered on.
- The green LED illuminates when a link has been established between the NUMAlink module and the module to which it is connected.

If both LEDs are dark, check the connections at both ends of the NUMAlink cable to ensure they are firmly seated. Check the power-on status of both units the cable is connected with.

Compute/Memory Blade LEDs

Each compute/memory blade installed in an IRU has a total of eight LED indicators arranged in two rows of four and visible behind the perforated sheetmetal of the blade:

- One green LED shows power-on complete status for the blade.
- One red LED shows power failure or bad voltage status within the blade.
- Two green NUMalink indicators show NI0 and NI1 connection status between the blade and the router board within the IRU. Constant green is a good connection.
- Four amber (heartbeat LEDs) indicate compute activity (the LEDs light up according to the number, activity and type of processors installed in the blade). If the IRU is fully powered on and booted and none of the amber LEDs are lit there is most likely a problem with the compute/memory blade. Try reseating the blade in the slot. Confirm the two green NUMalink status LEDs are on. If there is no LED activity on the blade, it must be replaced.




Figure 7-1 Compute Blade Status LED Locations

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 7-2 shows the sequence of events that occurs if you use all of the SGI Electronic Support capabilities.

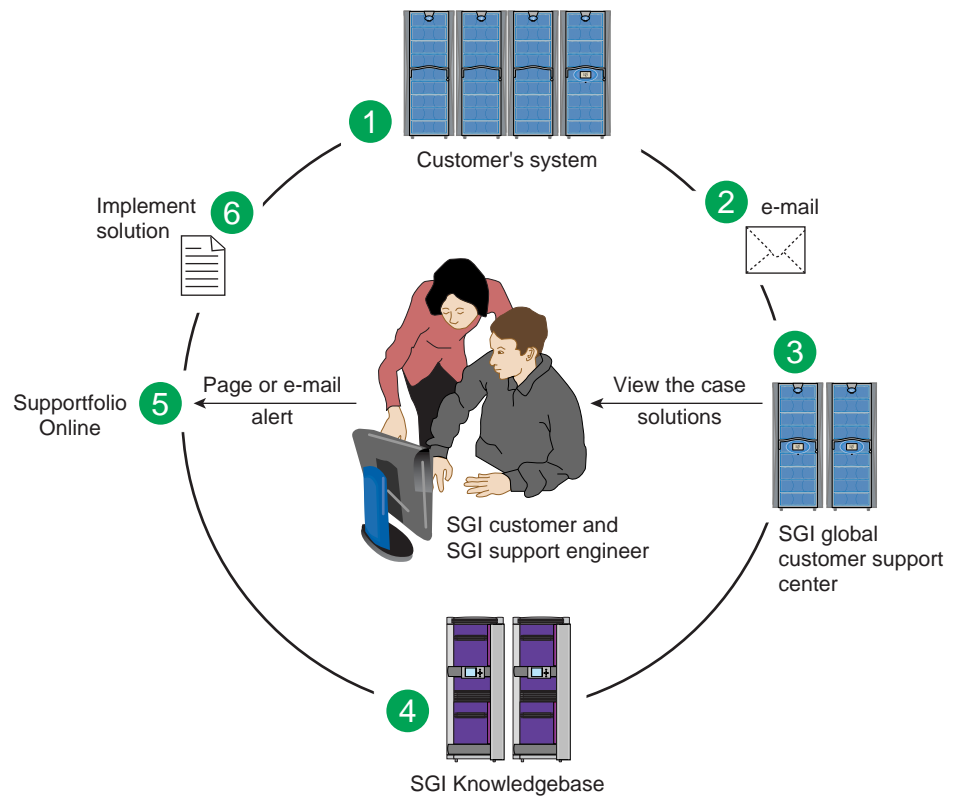


Figure 7-2 Full Support Sequence

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 119
- “Physical Specifications” on page 120
- “Environmental Specifications” on page 121
- “Power Specifications” on page 122
- “I/O Port Specifications” on page 123

System-level Specifications

Table A-1 summarizes the Altix 4700 system configuration ranges. Note that while each compute/memory board can house either one or two processor sockets; each socket can potentially hold one or two processor “cores”.

Table A-1 Altix 4700 System Configuration Ranges

Category	Minimum	Maximum
Processors	8 processor cores (4 blades) ^a	1024 processor cores
Individual Rack Units (IRUs)	1 per rack	4 per rack
Compute/memory blade DIMM capacity	4 DIMMs per blade	12 DIMMs per blade
System main memory DIMMs	16 per system image	3072 per system image
Number of base I/O blades	One per IRU	8 per system image

- a. Single-socket compute blades support one or two cores each.
Dual-socket blades support two or four cores per blade.

Physical Specifications

Table A-2 shows the physical specifications of the Altix 4700 system.

Table A-2 Altix 4700 Physical Specifications

Feature	Specification
Dimensions for a single tall rack, including doors and side panels	Height: 79.5 in. (201.9 cm) Width: 25.8 in. (65.5 cm) Depth: 43.45 in. (110.4 cm)
Shipping dimensions	Height: 81.25 in. (206.4 cm) Width: 42 in. (106.7 cm) Depth: 51 in. (129.5 cm)
Weight of full rack (maximum)	1500 lb (682 kg)
Shipping weight (maximum)	1600 lb (727 kg)
Access requirements	
Front	48 in. (121.9 cm)
Rear	48 in. (121.9 cm)
Side	None
Individual Rack Unit (IRU) enclosure specifications	Dimensions: 17.4 in high x 17.5 in wide x 32.5 in deep (with blower) (44.2 cm high x 44.45 cm wide x 82.5 cm deep) Weight: ~263 lbs. (119.5 kg)

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) 8% to 95% non-operating (no condensation)
Heat dissipation Altix 4700 (rack)	70.29 kBTU/hr maximum (based on 20.6 kW)
Cooling requirement	Ambient air or optional water cooling
Air flow: intake (front), exhaust (rear)	Approximately 1500 CFM (normal operation)
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	20.6 kW per rack
Hold-up time	20 ms
Three-phase power requirements	
Voltage	US/Japan 208V (180-254 VAC) International 400V (312-440 VAC)
Frequency	50-60 Hz
Power	20.6 kW per rack
Total harmonic distortion	Less than 10% 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 124
- “Serial Ports” on page 125
- “External SAS/SATA Port Connector” on page 127
- “USB Type A Connector” on page 128

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 IRU and the optional Dense router modules have 9-pin serial interface connectors. These ports are for console interface and 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 Connector

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

External SAS/SATA Port Connector

Figure A-3 shows the connector pin locations for the external Serial Attached SCSI (SAS) and Serial Advanced Technology Attachment (SATA) connector. This combined technology connector is used to support optional external storage. Table A-7 lists the pin assignments for this SAS/SATA connector.




Figure A-3 Pin Number Locations for External SAS/SATA Port

Table A-7 SAS/SATA Pin Assignments

Pin Number	Signal Name	Pin Number	Signal Name
1	Gnd	8	Gnd
2	2tx_h	9	3tx_h
3	2tx_l	10	3tx_l
4	Gnd	11	Gnd
5	2rx_l	12	3rx_l
6	2rx_h	13	3rx_h
7	Gnd	14	Gnd

USB Type A Connector

Figure A-4 shows the USB type A connector provided on the base I/O that supports general USB applications and optional keyboard and mouse configurations.

Table A-8 lists the pin assignments for the USB type A connector.




Figure A-4 Pin Number Locations for USB Type A Connector

Table A-8 Pin Assignments for USB Type A Connector

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

Safety Information and Regulatory Specifications

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

- “Safety Information” on page 129
- “Regulatory Specifications” on page 131

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 131
- “CE Notice and Manufacturer’s Declaration of Conformity” on page 131
- “Electromagnetic Emissions” on page 132
- “Shielded Cables” on page 134
- “Electrostatic Discharge” on page 134
- “Laser Compliance Statements” on page 135
- “Lithium Battery Statements” on page 136

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

警告使用者：

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

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 lå tuijota såteeseenstirra ej in i strålen.



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



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

Lithium Battery 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 suosittelemaan tyyppiin. Hävitä käytetty paristo valmistajan ohjeiden mukaisesti.



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

Index

A

- Altix 4700 server
 - illustration, 51
- Altix 4700 servers
 - monitoring locations, 13

B

- battery statements, 136
- block diagram
 - system, 54
- brick slot numbering, 61

C

- cache coherency, 56
- CMN number, 131
- command targeting
 - L1 mode, 33
 - L2 mode, 40
- connector pin assignment
 - SAS/SATA, 127
 - USB type A, 128
- console
 - L1 mode power on, 9
 - powering on bricks, 9
- console communications, 44
 - subchannels, 44

- console mode
 - from L1 mode, 34
 - from L2 emulator mode, 43
- console selection, 44
- controller
 - L1, 27
- customer service, xix

D

- Dense router
 - features, 67
 - front panel components, 69
 - overview, 67
 - port specifications, 72
 - rear panel components, 71
 - technical specifications, 72
- disk drive installation, 106
- documentation
 - available via the World Wide Web, xvii
 - conventions, xix

E

- embedded support partner (ESP), 117
 - system support, 12
- environmental specifications, 121
- error messages, 111
- ESP (embedded support partner)
 - system support, 12

F

- failed fan, 80
- fan numbers for the IRU, 81
- fan RPM level, 80
- front panel display
 - L1 controller, 29

I

- information, warning, and error messages
 - viewing in L2 mode, 36, 42
- install
 - disk drive, 106

K

- Knowledgebase, 117

L

- L1 console selection, 35
- L1 controller, 13
 - description, 28
 - error messages, 111
 - front panel display, 29
 - front panel display items, 29
 - functions, 28
- L1 firmware upgrade, 47
- L1 mode
 - command targeting, 33
 - from L1 software, 32
 - from L2 emulation mode, 45
 - powering on at console, 9
 - to console mode, 34
 - viewing information, warning, and error messages, 36
 - viewing system configuration, 33

L1 software

- modes of operation, 31
- to L1 mode, 32

L2 emulation mode

- to L1 mode, 45
- viewing system configuration, 38

L2 emulator mode

- powering off, 42
- powering on, 42
- resetting system, 42
- setting command targeting, 40
- to console mode, 43
- viewing information, warning, and error messages, 42

laser compliance statements, 135**lithium battery warning statements, 2, 136****M**

- modes of operation
 - L1 software, 31
- monitoring
 - embedded support partner, 12
 - IRUs with L1 controller, 13
 - server, 13

N

- non-uniform memory access (NUMA), 57
- numbering
 - bricks in a rack, 61
 - racks, 61

P

- PCI and PCI-X cards
 - installation procedure
 - PCI and PCI-X card installation, 94

physical specifications
 Altix 4700 Physical Specifications, 120

pinouts
 Ethernet connector, 124
 serial connector, 125

power specifications, 122

powering off

L2 mode, 42

powering on

at system console, 9

in L1 mode, 9

L2 emulator mode, 42

preparation, 6

product support, xix

R

RAS features, 57

replace a fan, 80

Replacing, 76

resetting system

L2 mode, 42

S

SAS/SATA

external connector pin assignments, 127

server

monitored at L1 controller, 13

monitoring locations, 13

SGI Knowledgebase. See Knowledgebase

subchannels

associated with console communications, 44

support

embedded support partner (ESP), 12

Supportfolio Online, 117

system architecture, 52

system block diagram, 54

system components

Altix 4700 server, 60

list of, 59

system controller, 25

system features, 55

system overview, 49

T

tall rack

features, 64

technical specifications

Dense router, 72

system level, 119

technical support, xix

to monitor IRUs, 13

troubleshooting

L1 controller error messages, 111

problems and recommended actions, 110

U

upgrading L1 firmware, 47

USB type A

connector pin assignments, 128

V

viewing system configuration

L1 mode, 33

W

warning messages

viewing in L2 mode, 36, 42