

SGI® Altix® UV 1000 System User's Guide

Document Number 007-5663-002

COPYRIGHT

© 2010 SGI. 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 SGI.

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, 46600 Landing Parkway, Fremont, CA 94538.

TRADEMARKS AND ATTRIBUTIONS

Silicon Graphics, SGI, the SGI logo, Altix, NUMAlink and NUMAflex are trademarks or registered trademarks of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries worldwide.

Intel, Itanium and Xeon 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.

Infiniband is a trademark of the InfiniBand Trade Association.

Linux is a registered trademark of Linus Torvalds.

Red Hat and all Red Hat-based trademarks are trademarks or registered trademarks of Red Hat, Inc. in the United States and other countries.

SUSE LINUX is a registered trademark of Novell Inc.

Windows is a registered trademark of Microsoft Corporation in the United States and other countries.

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

Record of Revision

Version	Description
001	June 2010 First Release
002	October 2010 Addition of PCIe maintenance information and updates to system drawings.

Contents

List of Figures	ix
List of Tables	xi
Audience	xiii
Important Information	xiii
Chapter Descriptions	xiv
Related Publications	xv
Conventions	xvii
Product Support	xviii
Reader Comments	xviii
1. Operation Procedures	1
Precautions	1
ESD Precaution	1
Safety Precautions	2
System Control Network Overview	3
System Controller Access	4
Connecting to the UV System Control Network	5
Communicating with the System	5
The Command Line Interface	5
SMN Specific CLI Commands	6
Example SMN <code>uvcli</code> Commands	6
Additional CLI Commands Used With the SMN	7
The SGI Management Center Graphical User Interface	8
Powering the System On and Off	8
Preparing to Power On	9
Powering-On and Off From the SGI Management Center Interface	13

Powering On and Off from the Command Line Interface	14
Power On an Altix UV System From the SMN Command Line	14
Command Options for Power On	15
Optional Power On From the CMC Command Line	15
Bootting Directly From a CMC	15
Serial Console Hardware Requirements.	15
Establishing a Serial Connection to the CMC on UV 1000.	16
Establishing CMC IP Hardware Connections	16
Using DHCP to Establish an IP Address	17
Using a Static IP Address	18
Power On the System Using the CMC Network Connection	18
Optional Power On Using the SMC to Connect to the CMC	19
Monitoring Power On	20
Power off an Altix UV System	21
Additional CLI Power Command Options	21
Using Embedded Support Partner (ESP).	23
System Control Interface Options	24
Optional Components	25
PCIe Cards	25
2. System Control	27
Levels of System Control	27
System Management Node (SMN) Overview	28
CMC Overview	29
BMC Overview	31
System Controller Interaction	31
IRU Controllers	32
Chassis Management Controller Functions	32
1U Console Option	32
Flat Panel Rackmount Console Option Features	33
3. System Overview	35
System Models	37
System Architecture	39

System Features	41
Modularity and Scalability	41
Distributed Shared Memory (DSM)	41
Distributed Shared I/O	43
Chassis Management Controller (CMC)	43
ccNUMA Architecture	43
Cache Coherency	43
Non-uniform Memory Access (NUMA)	44
Reliability, Availability, and Serviceability (RAS)	44
System Components	46
Bay (Unit) Numbering	48
Rack Numbering	48
Optional System Components	48
4. Rack Information	49
Overview	49
Altix UV 1000 Series Rack (42U)	50
Altix System Rack Technical Specifications	54
5. PCIe and Disk Add or Replace Procedures	55
Maintenance Precautions and Procedures	55
Preparing the System for Maintenance or Upgrade	56
Returning the System to Operation	56
Overview of PCI Express (PCIe) Operation	57
Adding or Replacing PCIe Cards in the Expansion Enclosure	58
Installing Cards in the 1U PCIe Expansion Chassis	59
Installing PCIe Cards in a Riser Blade	62
Installing or Replacing a Drive in the Dual Disk Riser Blade	66
6. Troubleshooting and Diagnostics	69
Troubleshooting Chart	70
LED Status Indicators	71
IRU Power Supply LEDs	71
Compute/Memory Blade LEDs	71
SGI Electronic Support	74

A.	Technical Specifications and Pinouts	. 77
	System-level Specifications	. 77
	Physical Specifications	. 78
	Environmental Specifications	. 79
	Power Specifications	. 80
	I/O Port Specifications	. 81
	Ethernet Port	. 82
	Serial Ports	. 83
	USB Type A Connector	. 85
B.	Safety Information and Regulatory Specifications	. 87
	Safety Information	. 87
	Regulatory Specifications	. 89
	CMN Number	. 89
	CE Notice and Manufacturer’s Declaration of Conformity	. 89
	Electromagnetic Emissions	. 90
	FCC Notice (USA Only)	. 90
	Industry Canada Notice (Canada Only)	. 91
	VCCI Notice (Japan Only)	. 91
	Korean Class A Regulatory Notice	. 91
	Shielded Cables	. 92
	Electrostatic Discharge	. 92
	Laser Compliance Statements	. 93
	Lithium Battery Statement	. 94

List of Figures

Figure 1-1	System Management Node Rear Connections	4
Figure 1-2	UV CMC Connections	4
Figure 1-3	IRU Power Supply Cable Location Example	9
Figure 1-4	Single-Phase 2-Outlet PDU Example	10
Figure 1-5	Single-Phase 8-Outlet PDU	11
Figure 1-6	Three-Phase PDU Examples	12
Figure 2-1	System Management Node Front and Rear Panels	29
Figure 2-2	SGI Altix UV 1000 LAN-attached System Control Network	30
Figure 2-3	Optional 1U Rackmount Console	33
Figure 2-4	System Management Node (SMN) Direct Video Connection Ports	34
Figure 3-1	SGI Altix UV 1000 System Example	36
Figure 3-2	SGI Altix UV 1000 IRU and Rack	38
Figure 3-3	Functional Block Diagram of the Individual Rack Unit	40
Figure 3-4	Blade Node Block Diagram	42
Figure 3-5	Altix UV 1000 IRU System Components Example	47
Figure 3-6	BaseIO Riser Enabled Blade Front Panel Example	48
Figure 4-1	Altix UV 1000 Series Rack Example	51
Figure 4-2	Front Lock on Tall (42U) Altix Rack	52
Figure 4-3	Optional Water-Chilled Cooling Units on Rear of Altix Rack	53
Figure 5-1	Comparison of PCI/PCI-X Connector with PCI Express Connectors	57
Figure 5-2	The PCIe Expansion Enclosure	60
Figure 5-3	Card Slot Locations	61
Figure 5-4	Removing a PCIe Enabled Riser Blade	62
Figure 5-5	PCIe Card Access in Riser Blade	63
Figure 5-6	Install/Replace a New PCIe Card in the Riser Blade.	64
Figure 5-7	Replace the Blade After PCIe Installation	65
Figure 5-8	Location of Disk Drive Bays	67

Figure 5-9	Removing/Replacing a Disk Drive 68
Figure 6-1	UV Compute Blade Status LED Locations 73
Figure 6-2	Full Support Sequence Example 74
Figure A-1	Ethernet Port 82
Figure A-2	Serial Port Connector 83
Figure A-3	Pin Number Locations for USB Type A Connector 85
Figure B-1	VCCI Notice (Japan Only) 91
Figure B-2	Korean Class A Regulatory Notice 91

List of Tables

Table 4-1	Tall Altix Rack Technical Specifications	54
Table 5-1	SGI Altix UV 1000 PCIe Support Levels	58
Table 5-2	PCIe Expansion Slot Bandwidth Support Levels	59
Table 6-1	Troubleshooting Chart	70
Table 6-2	Power Supply LED States	71
Table 6-3	Altix UV Blade LED Status Information	72
Table A-1	Altix UV 1000 System Configuration Ranges	77
Table A-2	Altix UV 1000 Physical Specifications	78
Table A-3	Environmental Specifications	79
Table A-4	Power Specifications	80
Table A-5	Ethernet Pinouts	82
Table A-6	Serial Port Pinout.	84
Table A-7	Pin Assignments for USB Type A Connector	85

About This Guide

This guide provides an overview of the architecture, general operation and descriptions of the major components that compose the SGI® Altix® UV 1000 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 UV 1000 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 UV 1000 system, with the exception of the following items that you can perform yourself:

- Using your system console controller to enter commands and perform system functions such as powering on and powering off, as described in this guide.
- Adding and replacing PCIe cards, as described in this guide.
- Adding and replacing disk drives in dual-disk enabled riser blades.
- Using the On/Off switch and other switches on the rack PDUs.
- 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 overall system control network interface and provides basic instructions for operating the controllers.
- Chapter 3, “System Overview,” provides technical overview information needed to understand the basic functional architecture of the SGI Altix UV 1000 systems.
- Chapter 4, “Rack Information,” describes the rack sizes and general features.
- Chapter 5, “PCIe and Disk Add or Replace Procedures,” provides instructions for installing or removing the customer-replaceable components of your system.
- Chapter 6, “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 UV 1000 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 UV 1000 series system:

- *SGI UV CMC Controller Software User's Guide*
(P/N 007-5636-00x)

This guide describes how to use the system console controller commands to monitor and manage your SGI Altix UV 1000 or UV 100 system via line commands. Coverage of control includes descriptions of the interface and usage of the commands. These commands are primarily used when a system management node is not present in the system. Note that it does **not** cover controller command information for the SGI Altix UV 10.

- *SGI Altix UV System Management Node Administrator's Guide*
(P/N 007-5694-00x)

This guide covers the system management node (SMN) for SGI Altix UV 1000 and SGI Altix UV 100 series systems. It describes the software and hardware components used with the SMN as well as providing an overview of the UV system control network. System network addressing is covered and a chapter on how to use KVM to enable remote console access from the system management node is included.

- *SGI Management Center Quick Start Guide*
(P/N 007-5672-00x)

This document may be helpful to users or administrators of SGI Altix UV systems using the SGI Management Center interface. The guide provides introductory information on configuration, operation and monitoring of your Altix UV system using the management center software.

- *SGI Management Center System Administrator's Guide*
(P/N 007-5642-00x)

This guide is intended for system administrators who work with the SGI Management Center software GUI to manage and control Altix UV 100 or 1000 systems. Depending on your system configuration and implementation, this guide may be optional. The manual is written with the assumption the user has a good working knowledge of Linux.

- *SGI UV Systems Linux Configuration and Operations Guide*
(P/N 007-5629-00x)

This guide discusses the operational management of SGI ProPack for Linux when used on the SGI Altix UV 100 or UV 1000 family of servers. It explains how to perform general system configuration and operation under Linux.

- *SGI UV Systems Installation Guide*
(P/N 007-5675-00x)

This guide covers software installation on UV 100 or UV 1000 systems and their SMNs.

- *Linux Application Tuning Guide for SGI X86-64 Based Systems*
(P/N 007-5646-00x)

This guide includes a chapter that covers advanced tuning strategies for applications running on Altix UV 100 or UV 1000 systems as well as other SGI X86 based systems.

- 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 shipped with Linux 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. References in the documentation to these pages include the name of the command and the section number in which the command is found. For example, to display a man page, type the request on a command line:

```
man commandx
```

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. International customers can visit <http://www.sgi.com/support/>. Click on the “Support Centers” link under the “Online Support” heading for information on how to contact your nearest SGI customer support center.

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
46600 Landing Parkway
Fremont, California 94538

SGI values your comments and will respond to them promptly.

Operation Procedures

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

- “Precautions” on page 1
- “System Control Network Overview” on page 3
- “Powering the System On and Off” on page 8
- “Using Embedded Support Partner (ESP)” on page 23
- “System Control Interface Options” on page 24
- “Optional Components” on page 25

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 a grounding wrist strap when you handle any 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 87.



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 Control Network Overview

All Altix UV 1000 system individual rack units (IRUs) use an embedded chassis management controller (CMC). The CMC communicates with both the blade-level board management controllers (BMCs) and the system management node (SMN), which runs the SGI Management Center software. In concert with the SGI Management Center software, they are generically known as the system control network.

The *SGI Management Center System Administrator's Guide* (P/N 007-5642-00x) provides information on using the GUI to administer your Altix UV 1000 system.

The Altix UV 1000 system control network provides control and monitoring functionality for each compute blade, power supply, and fan assembly in each IRU enclosure in the system.

The SGI Management Center is an application that provides control over multiple IRUs, and communication to other UV systems. Remote administration requires that the SMN be connected by an Ethernet connection to a private or public Local Area Network (LAN).

The CMC network provides the following functionality:

- Powering the entire system on and off.
- Powering individual IRUs on and off.
- Power on/off individual blades in an IRU.
- Monitoring the environmental state of the system.
- Partitioning the system.
- Enter controller commands to monitor or change particular system functions within a particular IRU. See the *SGI UV CMC Controller Software User's Guide* (P/N 007-5636-00x) for a complete list of command line interface (CLI) commands.
- Provides access to the system OS console allowing you to run diagnostics and boot the system.

System Controller Access

Access to the UV system controller network is accomplished by the following connection methods:

- A LAN connection to the system management node (running the SGI Management Center software application). This can also be done using an optional VGA-connected console, see Figure 1-1.
- A direct ethernet connection to the **SBK** port (see Figure 1-2) on a CMC (also see the note below).
- A serial connection to the “Console” port on the CMC (see note below).

Note: In systems with less than four racks, a connection to any CMC is supported. In systems with more than four racks, use a CMC that is also used to interconnect the four rack building block groups - this CMC will have a connection to the SBK connector.




Figure 1-1 System Management Node Rear Connections




Figure 1-2 UV CMC Connections

Connecting to the UV System Control Network

The ethernet connection is the preferred method of accessing the system console.

Administrators can perform one of the following options for connectivity:

- If the SMN is plugged into the customer LAN, connect to the SMN (SSH w/ X11 Forwarding) and start the SGI Management Center remotely.
- An in-rack system console can be directly connected to the system management node via VGA and PS2. You can then log into the SMN and perform system administration either through CLI commands or via the SGI Management Center interface. Note that the CMC ethernet port, (labeled **SBK**) requires connecting to a network with a DHCP server if the SMN node is not used. The CMC is factory set to DHCP mode and thus has no fixed IP address and cannot be accessed until an IP address is established.
- A serial connection is used to communicate directly with the CMC. This connection is typically used for service purposes or for system controller and system console access in small systems where an ethernet connection or in-rack system console is not used or available.

Communicating with the System

The two primary ways to communicate with and administer the UV 1000 system are through SGI Management Center interface or the UV command line interface (CLI).

The Command Line Interface

The UV command line interface is accessible by logging into either a system maintenance node (SMN) or chassis management controller (CMC).

Log in as `root`, (default password is `root`) when logging into the CMC.

Login as `sysco`, when logging into the SMN.

Once a connection to the SMN or CMC is established, various system control commands can be entered. See “Powering On and Off from the Command Line Interface” on page 14 for specific examples of using the CLI commands.

SMN Specific CLI Commands

The following CLI command options are available specifically for the SMN:

- h|--help This help message.
- hh|--help This help message + CLI help message.
- q|--quiet No diagnostic message.
- s|--system Select UV system. If only one system is present, this one is selected.
Otherwise, this option is mandatory.
- S|--show depth Show nodes at depth >= 1 using optional supplied pattern.
Default pattern=*
- t|--target One target in one of the two following formats:
 - a. rack[/slot[/blade]]
 - b. r{1..}{s{i}}{1..2}{{b{n}}{0..15}}

Note: This format is NOT for `uvcli` only.

Examples: r1i02 = rack 1, slot 2

r2i1b4 = rack 2, slot 1, blade 4

Select the target from the CLI command itself, or, if not available, using the `-t` option.

Example SMN `uvcli` Commands

The following are examples of `uvcli` commands:

- `uvcli --help` This help.
- `uvcli -- leds --help` Help on leds command.
- `uvcli leds r1i1b4` Show leds on BMC located at rack 1, slot1, blade 4.


```
uvcli -t 1/1 leds    Show leds on all BMCs in rack 1, slot 1.
uvcli -- leds -v r1i1  Same as previous command but more verbose.
uvcli -S 1          Show all system serial numbers.
uvcli -S 1 '*/part*' Show all system partitions.
```

Additional CLI Commands Used With the SMN

The following list of available CLI commands are specifically for the SMN:

```
auth    authenticate SSN/APPWT change
bios    perform bios actions
bmc     access BMC shell
cmc     access CMC shell
config  show system configuration
console access system consoles
help    list available commands
hel     access hardware error logs
hwcfg   access hardware configuration variable
leds    display system LED values
log     display system controller logs
power   access power control/status
```

Type '<cmd> --help' for help on individual commands.

The SGI Management Center Graphical User Interface

The SGI Management Center interface is a server monitoring and management system. The SGI Management Center provides status metrics on operational aspects for each node in a system. The interface can also be customized to meet the specific needs of individual systems.

The *SGI Management Center System Administrator's Guide* (P/N 007-5642-00x) provides information on using the interface to monitor and maintain your Altix UV 1000 system. Also, see Chapter 2 in this guide for additional reference information on the SGI Management Center interface.

Powering the System On and Off

This section explains how to power on and power off individual rack units, or your entire Altix UV 1000 system, as follows:

- “Preparing to Power On” on page 9
- “Power On an Altix UV System From the SMN Command Line” on page 14
- “Power off an Altix UV System” on page 21

Using a system controller connection, you can power on and power off individual blades, IRUs or the entire system.

If you are using an SGI Management Center interface, you can monitor and manage your server from a remote location. For details, see the documentation for the power management tool you are using in concert with the SGI Management Center.

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 23.

Preparing to Power On

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

1. Check to ensure that the power connector on the cable between the rack's power distribution units (PDUs) and the wall power-plug receptacles are securely plugged in.
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, see the example in Figure 1-3. Setting the circuit breakers on the PDUs to the "On" position will apply power to the IRUs and will start the CMCs in the IRUs. Note that the CMC in each IRU stays powered on as long as there is power coming into the unit. Turn off the PDU breaker switch on each of the PDUs that supply voltage to the IRUs power supplies if you want to remove all power from the unit.




Figure 1-3 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 the following three figures) are turned on to provide power to the server when the system is powered on.

Figure 1-4 shows an example of a single-phase 2-plug PDU that can be used with the Altix UV 1000 system. This is the PDU that is used to distribute power to the IRUs when the system is configured with single-phase power.




Figure 1-4 Single-Phase 2-Outlet PDU Example

Figure 1-5 shows an example of an eight-plug single-phase PDU that can be used in the Altix UV 1000 rack system. This unit is used to support auxiliary equipment in the rack.




Figure 1-5 Single-Phase 8-Outlet PDU

Figure 1-6 shows examples of the three-phase PDUs that can be used in the SGI Altix UV 1000 system. These PDUs are used to distribute power to the IRUs when the system is configured with three-phase power




Figure 1-6 Three-Phase PDU Examples

Powering-On and Off From the SGI Management Center Interface

Commands issued from the SGI Management Center interface are typically sent to all enclosures and blades in the system (up to a maximum 4096 compute cores) depending on set parameters. SGI Management Center services are started and stopped from scripts that exist in */etc/init.d*

SGI Management Center, is commonly installed in */opt/sgi/sgimc*, and is controlled by one of these services—this allows you to manage SGI Management Center services using standard Linux tools such as `chkconfig` and `service`.

If your SGI Management Center interface is not already running, or you are bringing it up for the first time, use the following steps:

1. Power on the server running the SGI Management Center interface.
2. Open an ssh or other terminal session command line console to the SMN using a remote workstation or local VGA terminal.
3. Use the information in the section “Preparing to Power On” on page 9 to ensure that all system components are supplied with power and ready for bring up.
4. Log in to the SMN as `root` (the default password is **sgisgi**).
5. On the command line, enter **mgrclient** and press **Enter**.
The SGI Management Center Login dialog box is displayed.
6. Enter a user name (**root** by default) and password (**root** by default) and click **OK**.
The SGI Management Center interface is displayed.
7. The power on (green button) and power off (red button) are located in the middle of the SGI Management Center GUI’s Tool Bar - icons which provide quick access to common tasks and features.

See the *SGI Management Center System Administrator’s Guide* for more information.

Powering On and Off from the Command Line Interface

The Altix UV 1000 command line interface is accessible by logging into either the system management node (SMN) as `root` or the CMC as `root`.

Instructions issued at the command line interface of a local console prompt typically only affect the local partition or a part of the system. Depending on the directory level you are logged in at, you may power up an entire partition (SSI), a single rack, or a single IRU enclosure. In CLI command console mode, you can obtain only limited information about the overall system configuration. An SMN has information about the IRUs in its SSI. Each IRU has information about its internal blades, and also (if other IRUs are attached via NUMALink to the IRU) information about those IRUs.

Power On an Altix UV System From the SMN Command Line

1. Login to the SMN as `root`, via a terminal window similar to the following:

The default password for logging in to the SMN as `root` is `sgisgi`.

```
# ssh -X root@uv-system-smn
```

```
root@system-smn:~/hw>
```

Once a connection to the SMN is established, the SMN prompt is presented and various system control commands can be entered.

2. To see a list of available commands enter the following:

```
root@uv-system-smn:~/hw>ls /sysco/bin
```

3. Change the working directory to `sysco`, similar to the following:

```
root@uv-system-smn:~/hw>cd /sysco
```

In the following example the system is powered on without monitoring the progress or status of the power-on process. When a power command is issued, it checks to see if the individual rack units (IRUs) are powered on; if not on, the power command powers up the IRUs and then the blades in the IRU are powered on.

4. Enter the `power on` command, similar to the following:

```
sysco@uv-system-smn:~/hw>power on
```

The system will take time to fully power up (depending on size and options).

Command Options for Power On

The following example command options can be used with either the SMN or CMC CLI:

usage: **power [-vcow] on|up [TARGET]...**turns power on

-v, --verbose	verbose output
-c, --clear	clear EFI variables (system and partition targets only)
-o, --override	override partition check
-w, --watch	watch boot progress

To monitor the power-on sequence during boot, see the section “Monitoring Power On” on page 20, the `-uvpower` option must be included with the command to power on.

Optional Power On From the CMC Command Line

Because every Altix UV 1000 system comes with a system management node there should be few reasons for powering on the system from a CMC. Use the following information if you have a need to power on from a CMC rather than the SMN CLI or the SGI Management Center GUI. If the SMN is not available you can still boot the system directly by using the CMC, see “Booting Directly From a CMC”.

Booting Directly From a CMC

If a system management node (SMN) is not available, it is possible to power on and administer your system directly from the CMC. When available, the optional SMN should always be the primary interface to the system.

The console type and how these console types are connected to the Altix UV 1000 systems is determined by what console option is chosen. Establish either a serial connection and/or network/Ethernet connection to the CMC.

Serial Console Hardware Requirements

The console type and how these console types are connected to the Altix UV 1000 servers is determined by what console option is chosen. If you have an Altix UV 1000 server and wish to use a serially-connected “dumb terminal”, you can connect the terminal via a serial cable to the

(DB-9) RS-232-style console port connector on the CMC. The terminal should be set to the following functional modes:

- Baud rate of 115,200
- 8 data bits
- One stop bit, no parity
- No hardware flow control (RTS/CTS)

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

Establishing a Serial Connection to the CMC on UV 1000

If you have an Altix UV 1000 system and wish to use a serially-connected "dumb terminal", you can connect the terminal via a serial cable to the (DB-9) RS-232-style console port connector on the CMC board of the IRU.

1. The terminal should be set to the operational modes described in the previous subsection.

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

2. On the system management node (SMN) port, the CMC is configured to request an IP address via dynamic host configuration protocol (DHCP).
3. If your system does not have an SMN, the CMC address cannot be directly obtained by DHCP and will have to be assigned, see the following subsections for more information.

Establishing CMC IP Hardware Connections

For IP address configuration, there are two options: DHCP or static IP. The following subsections provide information on the setup and use of both.

Note: Both options require the use of the CMC's serial port, refer to Figure 1-2 on page 4.

Network (LAN RJ-45) connections to the Altix UV 1000 CMC are always made via the **SBK** port.

For DHCP, you must determine the IP address that the CMC has been assigned; for a static IP, you must also configure the CMC to use the desired static IP address.

To use the serial port connection, you must attach and properly configure an RS-232 cable to the CMC's "CONSOLE" port. Configure the serial port as described in "Serial Console Hardware Requirements" on page 15.

When the serial port session is established, the console will show a CMC login, and the user can login to the CMC as user "root" with password "root".

Using DHCP to Establish an IP Address

To obtain and use a DHCP generated IP address, plug the CMC's external network port (**SBK**) into a network that provides IP addresses via DHCP, the CMC can then acquire an IP address.

To determine the IP address assigned to the CMC, you must first establish a connection to the CMC serial port (as indicated in the section "Serial Console Hardware Requirements" on page 15), and run the command "ifconfig eth0". This will report the IP address that the CMC is configured to use.

Running the CMC with DHCP is not recommended as the preferred option for UV 1000 systems. The nature of DHCP makes it difficult to determine the IP address of the CMC, and it is possible for that IP address to change over time, depending on the DHCP configuration usage. The exception would be a configuration where the system administrator is using DHCP to assign a "permanent" IP address to the CMC.

To switch from a static IP back to DHCP, the configuration file `/etc/sysconfig/ifcfg-eth0` on the CMC must be modified (see additional instructions in the "Using a Static IP Address" section). The file must contain the following line to enable use of DHCP:

```
BOOTPROTO=dhcp
```

Using a Static IP Address

To configure the CMC to use a static IP address, the user/administrator must edit the configuration file `/etc/sysconfig/ifcfg-eth0` on the CMC. The user can use the "vi" command (i.e. "vi /etc/sysconfig/ifcfg-eth0") to modify the file.

The configuration file should be modified to contain these lines:

```
BOOTPROTO=static
IPADDR=<IP address to use>
NETMASK=<netmask>
GATEWAY=<network gateway IP address>
HOSTNAME=<hostname to use>
```

Note that the "GATEWAY" and "HOSTNAME" lines are optional.

After modifying the file, save and write it using the vi command `:"w!"`, and then exit vi using `:"q"`. Then reboot the CMC (using the "reboot" command); after it reboots, it will be configured with the specified IP address.

Power On the System Using the CMC Network Connection

You can use a network connection to power on your UV system as described in the following steps:

1. Establish a Network/Ethernet connection (as detailed in the previous subsections). CMCs have their rack and "U" position set at the factory. The CMC will have an IP address, similar to the following:

```
SBK 172.17.<rack>.<slot>
```

2. You can use the IP address of the CMC to login, as follows:

```
ssh root@<IP-ADDRESS>
```

Typically, the default password for the CMC set out of the SGI factory is `root`. The default password for logging in as `sysco` on the SMN is `sgisgi`.

The following example shows the CMC prompt:

```
SGI Chassis Manager Controller, Firmware Rev. 0.x.xx
CMC:r1i1c>
```

This refers to rack 1, IRU 1, CMC.

3. Power up your Altix UV system using the power on command, as follows:

```
CMC:r1i1c> power on
```

The system will take time to fully power up (depending on size and options). Larger systems take longer to fully power on. Information on booting Linux from the shell prompt is included at the end of the subsection (“Monitoring Power On” on page 20).

Optional Power On Using the SMC to Connect to the CMC

Typically, the default password for the CMC set out of the SGI factory is **root**.

Use the following steps to establish a network connection from the SMN to the CMC and power on the system using the CMC prompt and the command line interface:

1. Establish a network connection to the CMC by using the `ssh` command from the SMN to connect to the CMC, similar to the following example:

Note: This is only valid if your PC or workstation that is connected to the CMC (via the network connection) has its `/etc/hosts` file setup to include the CMCs.

```
ssh root@hostname-cmc
```

The following example shows the CMC prompt:

```
SGI Chassis Manager Controller, Firmware Rev. x.x.xx
```

```
CMC:r1i1c>
```

This refers to rack 1, IRU 1, CMC.

2. Power up your Altix UV system using the power-on command, as follows:

```
CMC:r1i1c> power on
```

Note that the larger a system is, the more time it will take to power up completely. Information on booting Linux from the shell prompt is included at the end of the subsection (“Monitoring Power On” on page 20).

Monitoring Power On

Open a separate window on your PC or workstation and establish another connection to the SMN or CMC and use the `uvcon` command to open a system console and monitor the system boot process. Use the following steps:

```
CMC:r1i1c> uvcon
uvcon: attempting connection to localhost...
uvcon: connection to SMN/CMC (localhost) established.
uvcon: requesting baseio console access at r001i01b00...
uvcon: tty mode enabled, use 'CTRL-]' 'q' to exit
uvcon: console access established
uvcon: CMC <--> BASEIO connection active
*****
***** START OF CACHED CONSOLE OUTPUT *****
*****
***** [20100512.143541] BMC r001i01b10: Cold Reset via NL
broadcast reset
***** [20100512.143541] BMC r001i01b07: Cold Reset via NL
broadcast reset
***** [20100512.143540] BMC r001i01b08: Cold Reset via NL
broadcast reset
***** [20100512.143540] BMC r001i01b12: Cold Reset via NL
broadcast reset
***** [20100512.143541] BMC r001i01b14: Cold Reset via NL
broadcast reset
***** [20100512.143541] BMC r001i01b04: Cold Reset via NL....
```

Note: Use CTRL-] q to exit the console.

Depending upon the size of your system, it can take 5 to 10 minutes for the Altix UV system to boot to the EFI shell. When the `shell>` prompt appears, enter `fs0`, as follows:

```
shell> fs0
```

At the `fs0` prompt, enter the Linux boot loader information, as follows:

```
fs0> /efi/suse/elilo
```

The ELILO Linux Boot loader is called and various SGI configuration scripts are run and the SUSE Linux Enterprise Server 11 Service Pack *x* installation program appears.

Power off an Altix UV System

To power down the Altix UV system, use the `power off` command, as follows:

```
CMC:r1i1c> power off
==== r001i01c (PRI) ====
```

You can also use the `power status` command, to check the power status of your system

```
CMC:r1i1c> power status
==== r001i01c (PRI) ====
```

```
on: 0, off: 32, unknown: 0, disabled: 0
```

The following command options can be used with the `power off|down` command:

```
usage: power [-vo] off|down [TARGET]...turns power off
-v, --verbose          verbose output
-o, --override         override partition check
```

Additional CLI Power Command Options

The following are examples of command options related to power status of the system IRUs. These commands and arguments can be used with either the SMN or CMC CLI.

```
usage: power [-vchow] reset [TARGET]...toggle reset
-v, --verbose          verbose output
-c, --clear            clear EFI variables (system and partition
targets only)
-h, --hold            hold reset high
-o, --override         override partition check
-w, --watch           watch boot progress
```

usage: **power [-v] ioreset** [TARGET]...toggle I/O reset
-v, --verbose verbose output

usage: **power [-vhow] cycle** [TARGET]...cycle power off on
-v, --verbose verbose output
-h, --hold hold reset high
-o, --override override partition check
-w, --watch watch boot progress

usage: **power [-v10ud] [status]** [TARGET]...show power status
-v, --verbose verbose output
-1, --on show only blades with on status
-0, --off show only blades with off status
-u, --unknown show only blades with unknown status
-d, --disabled show only blades with disabled status

usage: **power [-ov] nmi|debug** [TARGET]...issue NMI
-o, --override override partition check
-v, --verbose verbose output

usage: **power [-v] margin [high|low|norm|<value>]** [TARGET]...power
margin control
high|low|norm|<value> margin state
-v, --verbose verbose output

usage: **power --help**
--help display this help and exit

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 6.

System Control Interface Options

You can monitor and interact with your Altix UV 1000 server from the following sources:

- Using the SGI 1U rackmount console option you can connect directly to the system management node (SMN) for basic monitoring and administration of the Altix system. See “1U Console Option” in Chapter 2 for more information; SLES 11 or later is required.
- A PC or workstation on the local area network can connect to the SMN’s external ethernet port and set up remote console sessions or display GUI objects from the SGI Management Center interface.
- A serial console display can be plugged into the CMC at the rear of IRU 001. You can also monitor IRU information and system operational status from other IRUs that are connected to IRU 001.

These console connections enable you to view the status and error messages generated by the chassis management controllers in your Altix UV 1000 rack. For example, you can monitor error messages that warn of power or temperature values that are out of tolerance. See the section “1U Console Option” in Chapter 2, for additional information.

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 UV 1000 series server:

- Peripheral component interface (PCIe) cards into the optional PCIe expansion chassis.
- PCIe cards into the blade-mounted PCIe riser card.
- Disk drives in your dual disk drive riser card equipped compute blade.

PCIe Cards

The PCIe based I/O sub-systems, are industry standard for connecting peripherals, storage, and graphics to a processor blade. The following are the primary configurable I/O system interfaces for the Altix UV 1000 series systems:

- The optional two-slot internal PCIe riser card is a compute blade-installed riser card that supports one x8 and one x16 PCIe Gen2 card.
- The optional external PCIe riser card is a compute blade-installed riser card that supports two x16 PCI express Gen2 ports. These ports can be used to connect to an optional I/O expansion chassis that supports multiple PCIe cards. Each x16 connector on the riser card can support one I/O expansion chassis.

Important: PCIe cards installed in a two-slot internal PCIe riser card are *not* hot swappable or hot pluggable. The compute blade using the PCIe riser must be powered down and removed from the system before installation or removal of a PCIe card. Also see “Installing Cards in the 1U PCIe Expansion Chassis” on page 59 for more information.

Not all blades or PCIe cards may be available with your system configuration. Check with your SGI sales or service representative for availability. See Chapter 5, “PCIe and Disk Add or Replace Procedures” for detailed instructions on installing or removing PCIe cards or UV 1000 system disk drives.

System Control

This chapter describes the general interaction and functions of the overall Altix UV 1000 system control. System control parameters depend somewhat on the overall size and complexity of the SGI Altix UV 1000 but will generally include the following three areas:

- The system management node (SMN) which runs the SGI Management Center software
- The chassis management controllers (CMC) boards - one per IRU
- The individual blade-based board management controllers (BMC) - report to the CMCs

Note: While it is possible to operate and administer a very basic (single-rack) SGI Altix UV 1000 system without using an SMN and SGI Management Center, this is an exception rather than rule.

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 LAN connection to the system controller network is used. This Ethernet connection is made from a remote PC/workstation connected to the system management node (SMN). The SMN is a separate stand-alone server installed in the SGI Altix UV 1000 rack. The SMN acts as a gateway and buffer between the UV system control network and any other public or private local area networks.

Important: The SGI Altix UV system control network is a private, closed network. It should not be reconfigured in any way to change it from the standard SGI Altix UV factory installation. It should not be directly connected to any other network. The Altix UV system control network is not designed for and does not accommodate additional network traffic, routing, address naming (other than its own schema), or DHCP controls (other than its own configuration). The Altix UV system control network also is not security hardened, nor is it tolerant of heavy network traffic, and is vulnerable to Denial of Service attacks.

System Management Node (SMN) Overview

An Ethernet connection directly from the SMN (Figure 2-1 on page 29) to a local private or public Ethernet allows the system to be administered directly from a local or remote console via the SGI Management Center interface (most often installed on the SMN). Note that there is no direct inter-connected system controller function in the optional expansion PCIe modules.

The system controller network is designed into all IRUs. Controllers within the system report and share status information via the CMC Ethernet interconnect cables. This maintains controller configuration and topology information between all controllers in an SSI. Figure 2-2 on page 30 shows an example system control network using an optional and separate (remote) workstation to monitor a single-rack SGI Altix UV 1000 system. It is also possible to connect an optional PC or console directly to the SMN, see Figure 2-4 on page 34.

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 administering network connected SGI systems using the SGI Management Center, see the *SGI Management Center System Administrator's Guide* (P/N 007-5642-00x).




Figure 2-1 System Management Node Front and Rear Panels

CMC Overview

The CMC system for the SGI Altix UV 1000 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.

Viewing the system from the rear, the CMC blade is on the right side of the IRU. The CMC accepts direction from the SMN and supports powering-up and powering-down individual or groups of compute blades and environmental monitoring of all units within the IRU. The CMC sends operational requests to the Baseboard Management Controller (BMC) on each compute/memory blade installed. The CMC provides data collected from the compute nodes within the IRU to the system management node upon request.

CMCs can communicate with the blade BMCs and other IRU CMCs when they are linked together under a single system image (SSI); also called a partition. Each CMC shares its information with the SMN and other CMCs within the SSI. Note that the system management node (server), optional mass storage units and PCIe expansion enclosures do *not* have a CMC installed.




Figure 2-2 SGI Altix UV 1000 LAN-attached System Control Network

BMC Overview

Each compute blade in an IRU has a baseboard management controller (BMC). The BMC is a built-in specialized microcontroller hardware component that monitors and reports on the functional “health” status of the blade. The BMC provides a key functional element in the overall Intelligent Platform Management Interface (IPMI) architecture.

The BMC acts as an interface to the higher levels of system control such as the IRU’s CMC board and the higher level control system used in the system management node. The BMC can report any on-board sensor information that it has regarding temperatures, power status, operating system condition and other functional parameters that may be reported by the blade. When any of the preset limits fall out of bounds, the information will be reported by the BMC and an administrator can take some corrective action. This could entail a node shutdown, reset (NMI) or power cycling of the individual blade.

The individual blade BMCs do not have information on the status of other blades within the IRU. This function is handled by the CMCs and the system management node. Note that blades equipped with an optional BaseIO riser board have a dedicated BMC Ethernet port.

System Controller Interaction

In all Altix UV 1000 servers all the system controller types (SMNs, CMCs and BMCs) communicate with each other in the following ways:

- System control commands and communications are passed between the SMN and CMCs via a private dedicated gigabit Ethernet. The CMCs communicate directly with the BMC in each installed blade by way of the IRU’s internal backplane.
- All CMCs can communicate with each other via an Ethernet “ring” configuration network established within an SSI.
- In larger configurations the system control communication path includes a private, dedicated Ethernet switch that allows communication between SSI environments.

IRU Controllers

All IRUs have a chassis management controller (CMC) board installed. The following subsection describes the basic features of the controllers:

Note: For additional information on controller commands, see the *SGI UV CMC Controller Software User's Guide* (P/N 007-5636-00x).

Chassis Management Controller Functions

The following list summarizes the control and monitoring functions that the CMC performs. Many of the controller functions are common across all IRUs, however, some functions are specific to the type of enclosure.

- Monitors individual blade status via blade BMCs
- Controls and monitors IRU fan speeds
- Reads system identification (ID) PROMs
- Monitors voltage levels and reports failures
- Monitors and controls warning LEDs
- Monitors the On/Off power process
- Provides the ability to create multiple system partitions
- Provides the ability to flash system BIOS

1U Console Option

The SGI optional 1U console (Figure 2-3 on page 33) is a rackmountable unit that includes a built-in keyboard/touchpad. It uses a 17-inch (43-cm) LCD flat panel display of up to 1280 x 1024 pixels.




Figure 2-3 Optional 1U Rackmount Console

Flat Panel Rackmount Console Option Features

The 1U flat panel console option has the following listed features:

1. **Slide Release** - Move this tab sideways to slide the console out. It locks the drawer closed when the console is not in use and prevents it from accidentally sliding open.
2. **Handle** - Used to push and pull the module in and out of the rack.
3. **LCD Display Controls** - The LCD controls include On/Off buttons and buttons to control the position and picture settings of the LCD display.

4. **Power LED** - Illuminates blue when the unit is receiving power.

The 1U console attaches to the system management node server using PS/2 and HD15M connectors or to an optional KVM switch (not provided by SGI). See Figure 2-4 for the SMN video connection points. The 1U console is basically a “dumb” VGA terminal, it cannot be used as a workstation or loaded with any system administration program.

The 27-pound (12.27-kg) console automatically goes into sleep mode when the cover is closed.




Figure 2-4 System Management Node (SMN) Direct Video Connection Ports

System Overview

This chapter provides an overview of the physical and architectural aspects of your SGI Altix UV 1000 series system. The major components of the Altix UV 1000 series systems are described and illustrated.

The Altix UV 1000 series is a family of multiprocessor distributed shared memory (DSM) computer systems that initially scale from 32 to 4,096 Intel processor cores as a cache-coherent single system image (SSI). Future releases may 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 lower entry-level cost with global scalability in processors, memory, and I/O. You can install and operate the Altix UV 1000 series system in your lab or server room. Each 42U SGI rack holds one or two 18-U high enclosures that support up to 16 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 UV 1000 IRU enclosures.

This chapter consists of the following sections:

- “System Models” on page 37
- “System Architecture” on page 39
- “System Features” on page 41
- “System Components” on page 46

Figure 3-1 shows the front view of a single-rack Altix UV 1000 system.




Figure 3-1 SGI Altix UV 1000 System Example

System Models

The basic enclosure within the Altix UV 1000 system is the 18U high “individual rack unit” (IRU). The IRU enclosure contains up to 16 single-wide blades connected to each other via a backplane. Each IRU has ports that are brought out to external NUMALink 5 connectors. The 42U rack for this server houses all IRU enclosures, option modules, and other components; up to 64 processor sockets (512 processor cores) in a single rack. The Altix UV 1000 server system can expand up to 4,096 Intel processor cores per SSI; a minimum of one BaseIO equipped blade is required for every 4,096 processor cores. Higher core counts in an SSI may be available in future releases, check with your SGI sales or service representative for current information.

Figure 3-2 shows an example of how IRU placement is done in a single-rack Altix UV 1000 server.

The system requires a minimum of one 42U tall rack with four single-phase power distribution units (PDUs) per IRU installed in the rack. Each single-phase PDU has two outlets (five are required to support the eight power supplies in an IRU and two power connections to the SMN).

The three-phase PDU has 9 outlets (8 connections are required to support each IRU installed in a rack).

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




Figure 3-2 SGI Altix UV 1000 IRU and Rack

System Architecture

The Altix UV 1000 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 512 processor cores 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 cores in each of 4 racks (2,048 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 UV 1000 system is a fifth-generation NUMAflex DSM architecture known as NUMALink 5. In the NUMALink 5 architecture, all processors and memory can be tied together into a single logical system. This combination of processors, memory, and internal switches constitute the interconnect fabric called NUMALink within each 18U IRU enclosure.

The basic expansion building block for the NUMALink interconnect is the processor node; each processor node consists of a Hub ASIC and two four-core, six-core, or eight-core processors with on-chip secondary caches. The Intel processors are connected to the Hub ASIC via quick path interconnects.

The Hub ASIC is the heart of the processor and memory node blade technology. This specialized ASIC acts as a crossbar between the processors and the network interface. The Hub ASIC enables any processor in the SSI to access the memory of all processors in the SSI.

Figure 3-3 on page 40 shows a functional block diagram of the Altix UV 1000 series system IRU.

A two-port channel extender blade is used in routerless system topologies to reduce the number of external NUMALink cables required to interconnect a system. The two-port channel extender blade installs in the NUMALink network slots 0 through 3 in a single IRU system. NUMALink network slots 0 and 1 are reserved for 8-port channel extender blades used in multi-IRU systems.

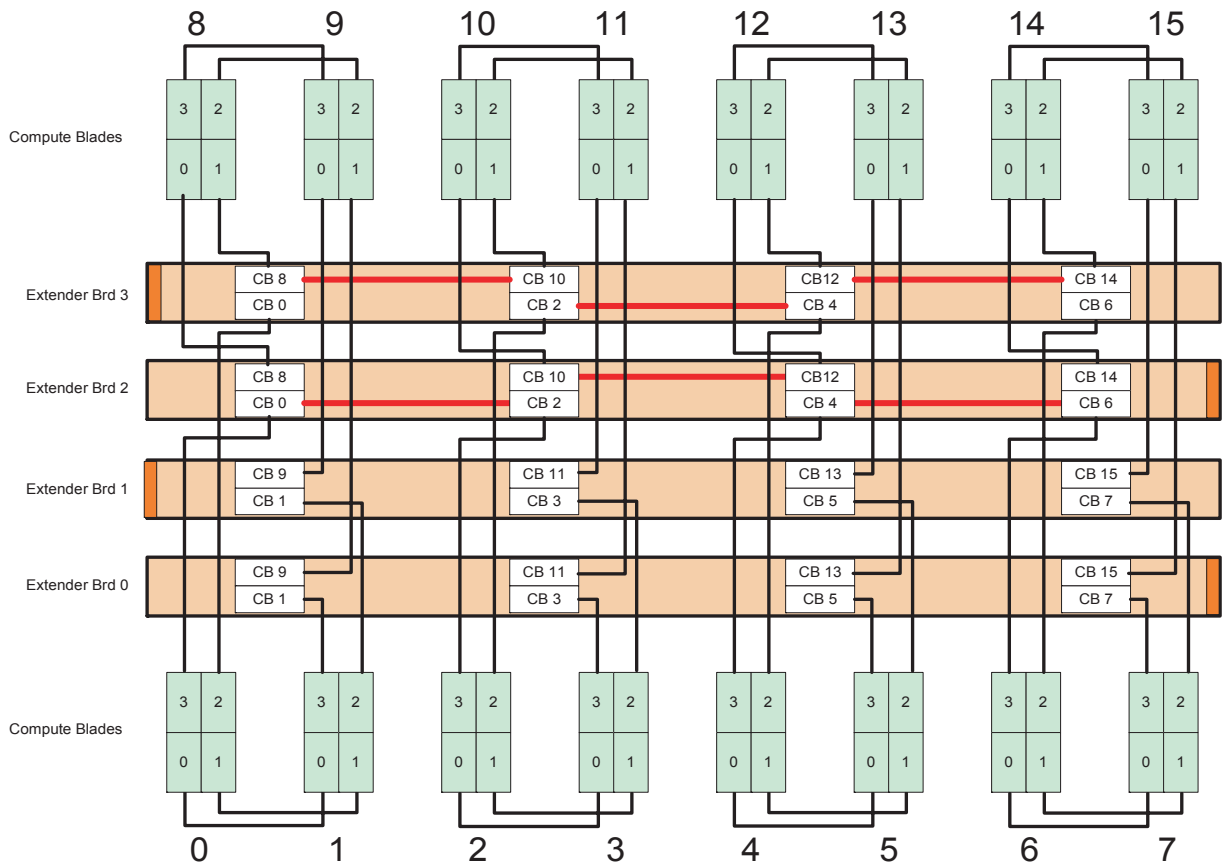


Figure 3-3 Functional Block Diagram of the Individual Rack Unit

System Features

The main features of the Altix UV 1000 series server systems are discussed in the following sections:

- “Modularity and Scalability” on page 41
- “Distributed Shared Memory (DSM)” on page 41
- “Chassis Management Controller (CMC)” on page 43
- “Distributed Shared I/O” on page 43
- “Reliability, Availability, and Serviceability (RAS)” on page 44

Modularity and Scalability

The Altix UV 1000 series systems are modular systems. The components are primarily housed in building blocks referred to as individual rack units (IRUs). Additional optional 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. The air-cooled IRU enclosure system has redundant, hot-swap fans and redundant, hot-swap power supplies.

Distributed Shared Memory (DSM)

In the Altix UV 1000 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 (SSI). This is to say that all NUMAlinked components sharing a single Linux operating system, operate and share the memory “fabric” of the system. 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. Note the following sub-types of memory within a system:

- If a processor accesses memory that it is connected to on a compute node blade, the memory is referred to as the node’s *local memory*. Figure 3-4 on page 42 shows a conceptual block diagram of the blade’s memory, compute and I/O pathways.
- 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*.

The following riser cards plug into the mezzanine connector:

1. Base I/O card
2. Boot drive
3. Integrated PCIE GEN2 (supports two PCIe cards)
4. I/O expansion (supports two PCIe card expansion channels).




Figure 3-4 Blade Node Block Diagram

Distributed Shared I/O

Like DSM, I/O devices are distributed among the blade nodes within the IRUs. Each BaseIO riser card equipped blade node is accessible by all compute nodes within the SSI (partition) through the NUMAlink interconnect fabric.

Chassis Management Controller (CMC)

Each IRU has a chassis management controller (CMC) located directly below the upper set of cooling fans in the rear of the IRU. The chassis manager supports powering up and down of the compute blades and environmental monitoring of all units within the IRU.

One GigE port from each compute blade connects to the CMC blade via the internal IRU backplane. A second GigE port from each even number blade slot is also connected to the CMC. This second port is used to support an optional BaseIO riser card in the even numbered slots.

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 UV 1000 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 may 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 UV 1000 server series uses 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 UV 1000 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.
 - All main memory has Intel Single Device Data Correction, to detect and correct 8 contiguous bits failing in a memory device. Additionally, the main memory can detect and correct any two-bit errors coming from two memory devices (8 bits or more apart).
 - All high speed links including Intel Quick Path Interconnect (QPI), Intel Scalable Memory Interconnect (SMI), and PCIe have CRC check and retry.
 - The NUMALink interconnect network is protected by cyclic redundancy check (CRC).

- Each blade/node installed has status LEDs that indicate the blade's operational condition; LEDs are readable at the front of the IRU.
- 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 UV 1000 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 UV 1000 system. Up to two IRUs can be installed in each rack. There is also space reserved for a system management node and other optional 19-inch rackmounted components.
- **Individual Rack Unit (IRU).** This enclosure contains eight power supplies, 2-16 compute/memory blades, BaseIO and other optional riser enabled blades for the Altix UV 1000. The enclosure is 18U high. Figure 3-5 on page 47 shows the Altix UV 1000 IRU system components.
- **Compute blade.** Holds two processor sockets and 8 or 16 memory DIMMs. Each compute blade can be ordered with a riser card that enables the blade to support various I/O options.
- **BaseIO enabled compute blade.** I/O riser enabled blade that supports all base system I/O functions including two ethernet connectors, one SAS port, one BMC ethernet port and three USB ports.

Note: While the BaseIO 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 shutdown. The Altix UV 1000 ships with RAID 1 functionality (disk mirroring) configured if the option is ordered.

- **Dual disk enabled compute blade.** This riser enabled blade supports two hard disk drives that act as the system disks for the SSI. This blade must be installed adjacent to and physically connected with the BaseIO enabled compute blade. Jbod, RAID 0 and RAID 1 are supported. Note that you must have the BaseIO riser blade optionally enabled to use RAID 1 mirroring on your system disk pair.
- **Two-Slot Internal PCIe enabled compute blade.** The internal PCIe riser based compute blade supports two internally installed PCI Express option cards.
- **External PCIe enabled compute blade.** This riser enabled board must be used in conjunction with a PCIe expansion enclosure. A x16 adapter card connects from the blade to the expansion enclosure, supporting up to four PCIe option cards.

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




Figure 3-5 Altix UV 1000 IRU System Components Example




Figure 3-6 BaseIO Riser Enabled Blade Front Panel 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 U18 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, 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 UV 1000 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 current information on available product options not discussed in this manual.

Rack Information

This chapter describes the physical characteristics of the tall (42U) UV 1000 racks in the following sections:

- “Overview” on page 49
- “Altix UV 1000 Series Rack (42U)” on page 50
- “Altix System Rack Technical Specifications” on page 54

Overview

At the time this document was published only the tall (42U) Altix rack (shown in Figure 4-2) was available from the SGI factory for use with the Altix UV 1000 systems. Other racks may be available to house the system IRUs, check with your SGI sales or service representative for information.

Altix UV 1000 Series Rack (42U)

The tall rack (shown in Figure 4-1 on page 51) has the following features and components:

- **Front and rear door.** The front door is opened by grasping the outer end of the rectangular-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 (see Figure 4-2 on page 52).

Note: The front door and rear door locks are keyed differently. The optional water-chilled rear door (see Figure 4-3 on page 53) does not use a lock.

The standard 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. Multiple cables are attached to the front of the IRUs; therefore, a significant part of the cable management occurs in the front part of the rack. The stand-alone system management nodes have cables that attach at the rear of the rack. Rear cable connections will also be required for optional storage modules installed in the same rack with the IRU(s). Optional inter-rack communication cables can pass through the top of the rack. These are necessary whenever the system consists of multiple racks. 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 in the rack.** Ten outlets are required for a single IRU system as follows:
 - Allow eight outlets for the first IRU
 - Two outlets for a maintenance node (server)
 - Allow eight more outlets for each additional IRU in the system

Eight two-outlet single-phase power distribution units (PDUs) are needed for a base rack system with two IRUs.

In addition an 8 outlet single-phase PDU is required for the system management node and other optional equipment. Each three-phase power distribution unit has 9 outlet connections; 18 outlets are required to power two IRUs.




Figure 4-1 Altix UV 1000 Series Rack Example




Figure 4-2 Front Lock on Tall (42U) Altix Rack




Figure 4-3 Optional Water-Chilled Cooling Units on Rear of Altix Rack

Altix System Rack Technical Specifications

Table 4-1 lists the technical specifications of the Altix UV 1000 series tall rack.

Table 4-1 Tall Altix Rack Technical Specifications

Characteristic	Specification
Height	79.5 in. (201.9 cm)
Width	31.3 in. (79.5 cm)
Depth	45.8 in. (116.3 cm)
Single-rack shipping weight (approximate)	2,381 lbs. (1,082 kg) air cooled 2,581 (1,173 kg) water assist cooling
Single-rack system weight (approximate)	2,300 lbs. (1,045 kg) air cooled 2,500 lbs (1,136 kg) water assist cooling
Voltage range	North America/International
Nominal	200-240 VAC /230 VAC
Tolerance range	180-264 VAC
Frequency	North America/International
Nominal	60 Hz /50 Hz
Tolerance range	47-63 Hz
Phase required	Single-phase or 3-phase
Power requirements (max)	34.57 kVA (33.88 kW) approximate
Hold time	16 ms
Power cable	8 ft. (2.4 m) pluggable cords

PCIe and Disk Add or Replace Procedures

This chapter provides information about installing and removing PCIe cards and system disk drives from your SGI system, as follows:

- “Maintenance Precautions and Procedures” on page 55
- “Adding or Replacing PCIe Cards in the Expansion Enclosure” on page 58
- “Installing or Replacing a Drive in the Dual Disk Riser Blade” on page 66

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 56
- “Returning the System to Operation” on page 56



Warning: To avoid problems that could void your warranty, your SGI or other approved system support engineer (SSE) should perform all the setup, addition, or replacement of parts, cabling, and service of your SGI Altix UV 1000 system, with the exception of the following items that you can perform yourself:

- Using your system console or network access workstation to enter commands and perform system functions such as powering on and powering off, as described in this guide.
- Installing, removing or replacing cards in the optional PCIe equipped riser blade.
- Installing, removing or replacing PCIe cards in the optional 1U PCIe expansion chassis.
- Adding and replacing disk drives used with your system and using the ESI/ops panel (operating panel) on optional mass storage.

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. Go to the section “Powering the System On and Off” in Chapter 1 if you are not familiar with power down procedures.
3. After the system is powered off, locate the power distribution unit(s) (PDUs) in the front of the rack and turn off the circuit breaker switches on each PDU.

Note: Powering the system off is not a requirement when replacing a RAID 1 system disk. Addition of a non-RAID disk can be accomplished while the system is powered on, but the disk is not automatically recognized by system software.

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. Power up the system. If you are not familiar with the proper power down procedure, review the section “Powering the System On and Off” in Chapter 1.
3. Verify that the LEDs on the system power supplies and system blades turn on and illuminate green which indicates that the power-on procedure is proceeding properly.

If your system does not boot correctly, see “Troubleshooting Chart” in Chapter 6, for troubleshooting procedures.

Overview of PCI Express (PCIe) Operation

This section provides a brief overview of the PCI Express (PCIe) technology 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 specific PCI Express board options available with the SGI Altix UV 1000.

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 5-1)
- Potential sustained throughput of x16 PCI Express is approximately four times that of the fastest PCI-X throughputs



Figure 5-1 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 - also called “by 1”.

SGI Altix UV 1000 PCIe technology is available up to 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 5-1 shows this concept.

For information on which slots in the PCIe expansion chassis support what lane levels, see Table 5-2 on page 59.

Table 5-1 SGI Altix UV 1000 PCIe Support Levels

SGI x16 PCIe Connectors	Support levels in optional chassis
x1 PCIe cards	Supported in all four slots
x2 PCIe cards	Supported in all four slots
x4 PCIe cards	Supported in all four slots
x8 PCIe cards	Supported in two slots
x16 PCIe cards	1 slot supported
x32 PCIe cards	Not supported

Adding or Replacing PCIe Cards in the Expansion Enclosure



Warning: Before installing, operating, or servicing any part of this product, read the “Safety Information” on page 87.

This section provides instructions for adding or replacing a PCIe card in a PCIe expansion enclosure installed in your system. To maximize the operating efficiency of your cards, be sure to read all the introductory matter before beginning the installation.



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

Installing Cards in the 1U PCIe Expansion Chassis

The PCIe expansion chassis functions in a similar manner to a computer chassis that supports PCIe slots. Always follow the manufacturer’s instructions or restrictions for installing their card.

Important: Replacement (swapping) of a PCIe card in the 1U chassis may be done while the system is powered on. Addition of a new card while the system is running requires a reboot to initiate recognition and functionality. Removal (without replacement) of an existing PCIe card may cause system error messages. When installing PCIe cards, ensure that the input current rating specified on the AC input label is not exceeded.

The EB4-1U-SGI chassis provides space for up to four (4) PCIe cards in the following “lane” bandwidth configurations:

Table 5-2 PCIe Expansion Slot Bandwidth Support Levels

PCIe expansion enclosure slot #	PCIe connector level supported by slot	PCIe slot number location in board “carriage”
Slot 1	Up to x16	Bottom-left side
Slot 2	Up to x4	Top-left side
Slot 3	Up to x8	Top-right side
Slot 4	Up to x4	Bottom-right side

Note: Before installing the PCIe expansion cards, be sure to remove each respective slot cover and use its screw to secure your expansion card in place.

1. Working from the front of the expansion chassis, locate the two “thumb screws” that hold the PCIe board “carriage” in the expansion chassis.
2. Turn the two thumb screws counter-clockwise until they disengage from the 1U chassis.
3. Pull the T-shaped board “carriage” out of the chassis until the slots are clear of the unit.
4. Select an available slot based on the lane support your PCIe card requires, see Table 5-2.
5. Remove the metal slot cover from the selected slot and retain its screw.

6. Fit the PCIe card into the slot connector with the connector(s) extending out the front of the bracket, then secure the board with the screw that previously held the metal slot cover.
7. Push the PCIe board “carriage” back into the enclosure until it is seated and twist the retaining thumb screws clockwise (right) until fully secure.

Important: After installation, be sure to power on the PCIe expansion enclosure before re-booting your system.




Figure 5-2 The PCIe Expansion Enclosure




Figure 5-3 Card Slot Locations

Installing PCIe Cards in a Riser Blade

If your system is equipped with an optional PCIe riser blade, you can install or replace PCIe cards (up to x16) using the following steps and illustrations:

1. Obtain a T10 torx driver to perform the necessary disassembly of the blade. It is recommended that you use an anti-static wrist strap during this procedure.
2. Shut down the system or the individual blade you will be upgrading and open the front door of the rack.
3. Grasp the front pull bars and depress the black release button on the blade and pull it out of the IRU, see Figure 5-4. Place it on a flat static-safe working surface.




Figure 5-4 Removing a PCIe Enabled Riser Blade

4. Unscrew and remove the two T10 torx-head screws located on either side of the rear part of the blade's sheetmetal housing, see Figure 5-5 on page 63. Set the screws aside for the re-installation process.
5. Slide the blade's sheetmetal cover back slightly and then lift it away from the unit.
6. Near the front of the blade you will remove a single T10 torx-head screw from the C-shaped bracket that holds both of the optional PCIe cards in place on the blade's riser board. Set the screw aside for the reinstallation process.




Figure 5-5 PCIe Card Access in Riser Blade

7. If you are replacing a PCIe card, gently remove the existing board from the PCIe connector on the riser and place in a static-safe bag or container.
8. Install a new PCIe card by aligning the board's interface connector with the riser's and inserting until the card is flush to the connector, see Figure 5-6 on page 64. The PCIe connector(s) such as graphics, Fibre Channel, etc. should protrude from the blade's front.




Figure 5-6 Install/Replace a New PCIe Card in the Riser Blade

9. Re-secure the C-shaped bracket that holds the PCIe option cards in place on the blade's riser.
10. Replace the blade's sheetmetal top cover and securing screws.
11. Slide the blade back into the IRU (see Figure 5-7 on page 65) seat it fully and restart the system or blade.




Figure 5-7 Replace the Blade After PCIe Installation

Installing or Replacing a Drive in the Dual Disk Riser Blade

The dual-disk drive riser card permits a blade to house two system disk drive bays as shown in Figure 5-8 on page 67. This section describes how to install or remove the drives. The blade supports RAID 1, RAID 0 and jbod disk drives.

Note: A RAID 1 drive may be replaced while the system is operating. Removal of a jbod or RAID 0 drive while the system is operating will cause generation of system errors and possible loss of data.

Use the following steps and illustrations to add or replace a disk drive in the Altix UV 1000 dual disk drive riser blade.

To remove a disk drive:

1. Press in and down on the red button until the handle is released, see Figure 5-9 on page 68.
2. Pull the handle upward and out until the locking mechanism is cleared.
3. Grasp the disk drive by the side and extract it from the disk riser blade.

Installation of a disk drive into the riser blade is the opposite of extraction, use the following steps:

1. Install a drive (pre-mounted on sled) into the dual disk drive riser blade (see Figure 5-8 on page 67) with the red button at the bottom.
2. Slide the disk and sled into the riser until the actuating teeth can grab the riser plate.

Important: Use care - do not insert the drive too far into the disk riser blade.

3. Push the drive handle down until the handle clicks into place (this completes the insertion of the disk drive).




Figure 5-8 Location of Disk Drive Bays




Figure 5-9 Removing/Replacing a Disk Drive

Troubleshooting and Diagnostics

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

- “Troubleshooting Chart” on page 70
- “LED Status Indicators” on page 71
- “SGI Electronic Support” on page 74

Troubleshooting Chart

Table 6-1 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 74.

Table 6-1 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>Confirm the PDU(s) supporting the IRU are on.</p>
The system will not boot the operating system.	Contact your SSE.
An amber LED illuminates on a blade.	See Table 6-2 on page 71 for a description of the status message.
The amber (yellow) status LED of an IRU power supply is lit or the LED is not lit at all. See Table 6-2 on page 71.	Ensure the power cable to the supply is firmly connected at both ends and that the PDU is turned to on. Check and confirm the supply is fully plugged into it's slot. If the green LED does not light, contact your support engineer.
The PWR LED of a populated PCIe slot is not illuminated.	Reseat the PCI card.
The Fault LED of a populated PCIe slot is illuminated (on).	Reseat the card. If the fault LED remains on, replace the card.
The amber LED of a disk drive is on.	Replace the disk drive.

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 bi-color 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 6-2 for a complete list.

Table 6-2 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

Compute/Memory Blade LEDs

Each compute/memory blade installed in an IRU has a total of 19 LED indicators visible behind the perforated sheetmetal of the blade. Note that the heartbeat LEDs may not be available or functional on all versions of the compute blades:

At the bottom end (or left side) of the blade:

- There are 8 amber heartbeat LEDs
- One green 12V power good LED
- One blue unit identifier (UI) LED
- One green 3.3V auxiliary power LED

At the top end (or right side) of the blade:

- There are four green NumaLink status LEDs
- An auxiliary power good green LED
- System power good green LED
- Ethernet 0 green LED
- Ethernet 1 green LED

If the blade is properly seated and the system is powered on and there is no LED activity showing on the blade, it must be replaced. Table 6-3 lists the blade's LED status information, Figure 6-1 on page 73 shows the locations of the blade LEDs.

Table 6-3 Altix UV Blade LED Status Information

LED #	LED color	LED status information/message
01	Amber	CPU 0 heartbeat
02	Amber	CPU 1 heartbeat
03	Amber	CPU 0 heartbeat
04	Amber	CPU 1 heartbeat
05	Amber	CPU 0 heartbeat
06	Amber	CPU 1 heartbeat
07	Amber	CPU 0 heartbeat
08	Amber	CPU 1 heartbeat
09	Green	12V power good to blade
10	Blue	Unit (blade) Identifier - blade is selected when lighted
11	Green	3.3V auxiliary power is present on the blade
12	Green	NUMALink 0 status on blade
13	Green	Auxiliary power to blade is good
14	Green	NUMALink 01 status on blade
15	Green	System power to blade is good

Table 6-3 (continued) Altix UV Blade LED Status Information

LED #	LED color	LED status information/message
16	Green	NUMALink 02 status on blade
17	Green	Ethernet 0 activity on blade
18	Green	NUMALink 03 status on blade
19	Green	Ethernet 1 activity on blade




Figure 6-1 UV 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 6-2 shows the sequence of events that occurs if you use all of the SGI Electronic Support capabilities.

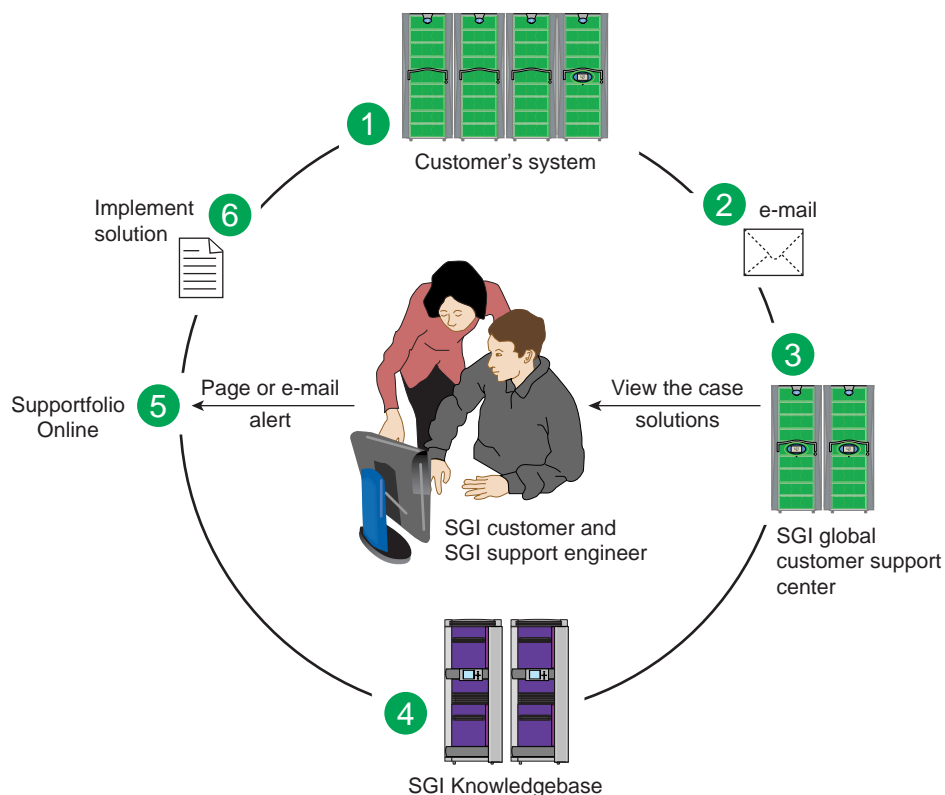


Figure 6-2 Full Support Sequence Example

The sequence of events can be described as follows:

1. Embedded Support Partner (ESP) software 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>

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 77
- “Physical Specifications” on page 78
- “Environmental Specifications” on page 79
- “Power Specifications” on page 80
- “I/O Port Specifications” on page 81

System-level Specifications

Table A-1 summarizes the Altix UV 1000 system configuration ranges. Note that while each compute/memory board can house two processor sockets; each socket can potentially hold four, six or eight processor “cores”.

Table A-1 Altix UV 1000 System Configuration Ranges

Category	Minimum	Maximum
Processors	8 processor cores (2 blades) ^a	4,096 processor cores
Individual Rack Units (IRUs)	1 per rack	2 per rack
Blades per IRU	2 per IRU	16 per IRU
Compute/memory blade DIMM capacity	8 DIMMs per blade	16 DIMMs per blade
CMC units	1 per IRU	2 per IRU
Number of BaseIO riser enabled blades	One per SSI	One per SSI

a. Dual-socket blades support eight to 16 cores per blade.

Physical Specifications

Table A-2 shows the physical specifications of the Altix UV 1000 system.

Table A-2 Altix UV 1000 Physical Specifications

Feature	Specification
Dimensions for a single 24-inch wide tall rack, including doors and side panels	Height: 79.5 in. (201.9 cm) Width: 31.3 in. (79.5 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)
Single-rack shipping weight (approximate)	2,381 lbs. (1,082 kg) air cooled 2,581 lbs. (1,173 kg) water assist cooling
Single-rack system weight (approximate)	2,300 lbs. (1,045 kg) air cooled 2,500 lbs. (1,136 kg) water assist cooling
Access requirements	
Front	48 in. (121.9 cm)
Rear	48 in. (121.9 cm)
Side	None
18U-high Individual Rack Unit (IRU) enclosure specifications	Dimensions: 30.5 in high x 24 in (flange width) wide x 29 in deep 77.5 cm high x 70 cm wide x 73.6 cm deep Weight: ~610 lbs. (277 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 full Altix UV 1000 (rack)	115.6 kBTU/hr maximum (based on 33.88 kW) approximate
Cooling requirement	Ambient air or optional water cooling
Air flow: intake (front), exhaust (rear)	Approximately 2,800 CFM (normal operation)
Maximum altitude	10,000 ft. (3,049 m) operating 40,000 ft. (12,195 m) non-operating

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	33.88 kW per rack
Hold-up time	16 ms
Three-phase power requirements	
Voltage	US/Japan 208V (180-254 VAC) International 400V (312-440 VAC)
Frequency	50-60 Hz
Power	33.88 kW per compute 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 82
- “Serial Ports” on page 83
- “USB Type A Connector” on page 85

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 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.

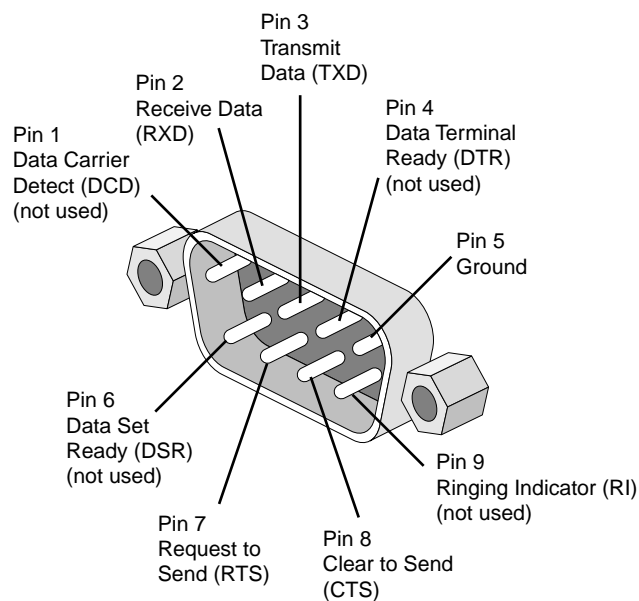


Figure A-2 Serial Port Connector

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

USB Type A Connector

Figure A-3 shows the USB type A connector provided on the BaseIO riser blade that supports general USB applications and optional keyboard and mouse configurations.

Table A-7 lists the pin assignments for the USB type A connector.




Figure A-3 Pin Number Locations for USB Type A Connector

Table A-7 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 87
- “Regulatory Specifications” on page 89

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 89
- “CE Notice and Manufacturer’s Declaration of Conformity” on page 89
- “Electromagnetic Emissions” on page 90
- “Shielded Cables” on page 92
- “Electrostatic Discharge” on page 92
- “Laser Compliance Statements” on page 93
- “Lithium Battery Statement” on page 94

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)

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-2 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 när denna del är öppnad ålå tuijota såteeseenstirra ej in i strålen.



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



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

Lithium Battery Statement



Warning: There is danger of explosion if a lithium battery is incorrectly replaced in this product. Replace a battery only with the same or equivalent type recommended by the manufacturer. Dispose of used batteries according to the manufacturer's instructions.
