



SGI® UV 2000 System User Guide

Document Number 007-5832-004

COPYRIGHT

© 2012-2015 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, 900 North McCarthy Blvd. Milpitas, CA 95035.

TRADEMARKS AND ATTRIBUTIONS

Silicon Graphics, SGI, the SGI logo, 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.

LSI, MegaRAID, and MegaRAID Storage Manager are trademarks or registered trademarks of LSI Corporation.

Linux is a registered trademark of Linus Torvalds in the U.S. and other countries.

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, 2012 First Release
002	May, 2013 MIC/GPU Blade and RAID updates
003	May, 2014 Cooling fan and system drive remove and replace information
004	November, 2015 Service reference and branding updates

Contents

List of Figures	xi
List of Tables	xiii
Audience	xv
Important Information	xv
Chapter Descriptions	xvi
Related Publications	xvii
Obtaining SGI Publications	xviii
Conventions	xix
Product Support	xx
Reader Comments	xx
1. Operation Procedures	1
Precautions	1
ESD Precaution	1
Safety Precautions	2
Power Connections Overview	2
System Connections Overview	7
Connecting to the UV System Control Network	7
System Controller Access	8
Serial Console Hardware Requirements	9
Establishing a Serial Connection to the CMC on SGI UV 2000	9
Establishing CMC IP Hardware Connections	9
Using DHCP to Establish an IP Address	10
Using a Static IP Address	11

System Control Overview	12
Communicating with the System	13
The SGI Management Center Graphical User Interface	13
Powering-On and Off From the SGI Management Center Interface	13
The Command Line Interface	14
Powering On and Off from the Command Line Interface	14
Command Options for Power On	15
Power On the System From the SMN Command Line.	15
Specific CLI Commands Used With the SMN	16
Optional Power On From the CMC Command Line	16
Booting Directly From a CMC	17
Power On the System Using the CMC Network	17
Optional Power On Using the SMN to Connect to the CMC	18
Monitoring Power On	19
Power off a UV System	20
Additional CLI Power Command Options	21
Optional SGI Remote Services (SGI RS)	22
SGI Remote Services Primary Capabilities	22
SGI Remote Services Benefits	22
SGI Remote Service Operations Overview	23
Optional Components	25
PCIe Cards	25
PCIe Drive Controllers in BaseIO Blade	27
RAID PCIe Disk Controller	27
Non-RAID PCIe Disk Controller	28
2. System Control	29
Levels of System Control	29
System Management Node (SMN) Overview	30
CMC Overview	31
BMC Overview	33
System Controller Interaction	33

IRU Controllers	34
Chassis Management Controller Functions.	34
1U Console Option	34
Flat Panel Rackmount Console Option Features	35
3. System Overview	37
System Models	39
System Architecture	41
System Features	43
Modularity and Scalability	43
Distributed Shared Memory (DSM)	43
Distributed Shared I/O	45
Chassis Management Controller (CMC)	45
ccNUMA Architecture	45
Cache Coherency	45
Non-uniform Memory Access (NUMA)	46
Reliability, Availability, and Serviceability (RAS)	46
System Components	48
Optional BaseIO SSDs	50
MIC/GPU Enabled Compute Blade	51
Bay (Unit) Numbering	52
Rack Numbering	52
Optional System Components	52
4. Rack Information	53
Overview	53
SGI UV 2000 Series Rack (42U)	54
SGI UV 2000 System Rack Technical Specifications	58
5. Optional Octal Router Chassis Information	59
Overview	59
SGI UV 2000 Series NUMALink Octal Router Chassis.	60
SGI UV 2000 External NUMALink System Technical Specifications	62

6.	Add or Replace Procedures	. 63
	Maintenance Precautions and Procedures	. 63
	Preparing the System for Maintenance or Upgrade	. 64
	Returning the System to Operation	. 64
	Overview of PCI Express (PCIe) Operation	. 65
	Adding or Replacing PCIe Cards in the Expansion Enclosure	. 66
	Installing Cards in the 1U PCIe Expansion Chassis	. 67
	Removing and Replacing an IRU Enclosure Power Supply	. 70
	Replacing a System Fan (Blower)	. 72
	Replacing a System Disk	. 75
7.	Troubleshooting and Diagnostics	. 79
	Troubleshooting Chart	. 80
	LED Status Indicators	. 81
	IRU Power Supply LEDs	. 81
	Compute/Memory Blade LEDs	. 82
	Accessing Online Support Information and Services	. 83
	SGI Customer Portal	. 83
	Technical Assistance	. 83
	Other Resources	. 84
	SGI Warranty Levels	. 84
	Optional SGI Remote Services (SGI RS)	. 84
A.	Technical Specifications and Pinouts	. 85
	System-level Specifications	. 85
	Physical Specifications	. 86
	Environmental Specifications	. 87
	Power Specifications	. 88
	I/O Port Specifications	. 89
	BaseIO VGA Port Information	. 89
	Ethernet Port	. 91
	Serial Ports	. 92
	USB Type A Connector	. 94

B. Safety Information and Regulatory Specifications	95
Safety Information	95
Regulatory Specifications	97
CMN Number	97
CE Notice and Manufacturer's Declaration of Conformity	97
Electromagnetic Emissions	98
FCC Notice (USA Only)	98
Industry Canada Notice (Canada Only)	99
VCCI Notice (Japan Only)	99
Korean Class A Regulatory Notice	99
Shielded Cables	100
Electrostatic Discharge	100
Laser Compliance Statements	101
Lithium Battery Statement	102

List of Figures

Figure 1-1	IRU Power Supply Cable Location Example	3
Figure 1-2	Single-Phase 2-Outlet PDU Example	4
Figure 1-3	Single-Phase 8-Outlet PDU	5
Figure 1-4	Three-Phase PDU Examples	6
Figure 1-5	System Management Node Rear Video Connections	8
Figure 1-6	UV CMC Connection Faceplate Example	8
Figure 1-7	SGI Remote Services Process Overview.	24
Figure 1-8	PCIe Option Blade Example with Full-Height and Low-Profile Slots	26
Figure 1-9	PCIe Option Blade Example with Two Low-Profile Slots	26
Figure 1-10	BaseIO Blade and PCIe Disk Controller Example	27
Figure 2-1	System Management Node Front and Rear Panels	31
Figure 2-2	SGI UV 2000 LAN-attached System Control Network Example	32
Figure 2-3	Optional 1U Rackmount Console	35
Figure 2-4	System Management Node (SMN) Direct Video Connection Ports	36
Figure 3-1	SGI UV 2000 Single-Rack System Example	38
Figure 3-2	SGI UV 2000 IRU and Rack	40
Figure 3-3	Functional Block Diagram of the Individual Rack Unit (IRU)	42
Figure 3-4	Blade Node Block Diagram Example	44
Figure 3-5	SGI UV 2000 IRU System Components Example	49
Figure 3-6	BaseIO Riser Enabled Blade Front Panel Example	50
Figure 3-7	MIC/GPU Enabled Compute Blade Example Front View	51
Figure 4-1	SGI UV 2000 Series Rack Example	55
Figure 4-2	Front Lock on Tall (42U) Rack	56
Figure 4-3	Optional Water-Chilled Cooling Units on Rear of SGI 42U Rack	57
Figure 5-1	SGI UV 2000 Optional NUMalink ORC (Rear View)	60
Figure 5-2	SGI UV 2000 Optional ORC Chassis Example (Front View)	61
Figure 6-1	Comparison of PCI/PCI-X Connector with PCI Express Connectors	65

Figure 6-2	The PCIe Expansion Enclosure 68
Figure 6-3	Card Slot Locations 69
Figure 6-4	Removing an Enclosure Power Supply 70
Figure 6-5	Replacing an IRU Enclosure Power Supply 71
Figure 6-6	UV 2000 Rear Fan Assembly (Blowers) 73
Figure 6-7	Removing a Fan From the Rear Assembly 74
Figure 6-8	Replacing an Enclosure Fan 75
Figure 6-9	Removing the UV 2000 System Disk Example 76
Figure 6-10	Replacing the UV 2000 System Disk Example 77
Figure 7-1	UV Compute Blade Status LED Locations Example. 82
Figure A-1	VGA Port Pinouts 89
Figure A-2	Ethernet Port 91
Figure A-3	Serial Port Connector 92
Figure A-4	Pin Number Locations for USB Type A Connector 94
Figure B-1	VCCI Notice (Japan Only) 99
Figure B-2	Korean Class A Regulatory Notice 99

List of Tables

Table 4-1	Tall Rack Technical Specifications	58
Table 5-1	External NUMalink Technical Specifications	62
Table 6-1	SGI UV 2000 PCIe Support Levels	66
Table 6-2	PCIe Expansion Slot Bandwidth Support Levels	67
Table 7-1	Troubleshooting Chart	80
Table 7-2	Power Supply LED States	81
Table A-1	SGI UV 2000 System Configuration Ranges	85
Table A-2	SGI UV 2000 Physical Specifications	86
Table A-3	Environmental Specifications	87
Table A-4	Power Specifications	88
Table A-5	VGA Pin Functions	90
Table A-6	Ethernet Pinouts	91
Table A-7	Serial Port Pinout.	93
Table A-8	Pin Assignments for USB Type A Connector	94

About This Guide

This guide provides an overview of the architecture, general operation and descriptions of the major components that compose the SGI® UV 2000 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 UV 2000 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 installation or service provider should perform all the set up, addition, or replacement of parts, cabling, and service of your SGI UV 2000 system, with the exception of the following items that you can perform yourself as needed:

- 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.
- Removing and replacing the IRU power supplies or system cooling fans.
- 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 UV 2000 systems.
- Chapter 4, “Rack Information,” describes the rack sizes and general features.
- Chapter 5, “Optional Octal Router Chassis Information,” describes the optional NUMALink router technology available in SGI UV 2000 systems consisting of two or more racks. This router technology is available in an enclosure “package” known as the Octal Router Chassis.
- Chapter 6, “Add or Replace Procedures,” provides instructions for installing or removing the customer-replaceable components of your system.
- Chapter 7, “Troubleshooting and Diagnostics,” provides recommended actions if problems occur on your system.
- Appendix A, “Technical Specifications and Pinouts,” provides physical, environmental, and power specifications for your system. Also included are the pinouts for the non-proprietary connectors.
- Appendix B, “Safety Information and Regulatory Specifications,” lists regulatory information related to use of the UV 2000 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 UV 2000 series system:

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

This guide describes how to use the system console controller commands to monitor and manage your SGI UV 2000 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 UV 10, UV 20, UV 30, UV 300 or UV 300EX.

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

This guide covers the system management node (SMN) for SGI UV 2000 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 UV systems using the SGI Management Center interface. The guide provides introductory information on configuration, operation and monitoring of your 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 SGI UV 2000 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 System Software Installation and Configuration Guide*
(P/N 007-5948-00x)

In UV systems that come with pre-installed Linux software operating systems; this document describes how to re-install it when necessary. Also, this guide is a reference document for people who manage the operation of SGI UV 2000 systems. It explains how to

perform general system configuration and operation under Linux for SGI UV. For a list of manuals supporting SGI Linux releases and SGI online resources, see the SGI Performance Suite documentation.

- *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 SGI UV 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.

Obtaining SGI Publications

You can obtain SGI documentation as follows:

Use the SGI customer portal and support website at:
<http://support.sgi.com>

Click on the following:

Support by Product > *productname* > **Documentation**

If you do not find what you are looking for, you can search for a specific product name by selecting **Search Knowledgebase** and using the category **Documentation**.

SGI systems shipped with Linux include a set of Linux man pages, formatted in the standard UNIX “man page” style. You can view man pages by typing `man title` at a command line.

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.

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
- “Power Connections Overview” on page 2
- “System Control Overview” on page 12
- “Optional SGI Remote Services (SGI RS)” on page 22
- “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 95.



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.

Power Connections Overview

Prior to operation, your SGI UV 2000 system should be set up and connected by a professional installer. If you are powering on the system for the first time or want to confirm proper power connections, 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 1-1 on page 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 IRU's power supplies if you want to remove all power from the unit.

Important: In a system configuration using 2-outlet single-phase PDUs, each power supply in an IRU should be connected to a different PDU within the rack. This will ensure the maximum amperage output of a single PDU is not exceeded if a power supply fails.




Figure 1-1 IRU Power Supply Cable Location Example

3. If you plan to power on a server that includes optional mass storage enclosures, make sure that the power switch on the rear of each PSU/cooling module (one or two per storage 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 1-2 shows an example of a single-phase 2-plug PDU that can be used with the SGI UV 2000 system. This PDU can be used to distribute power to the IRUs when the system is configured with single-phase power.




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

Figure 1-3 on page 5 shows an example of an eight-plug single-phase PDU that can be used in the SGI UV 2000 rack system. This unit is used to support auxiliary equipment in the rack.




Figure 1-3 Single-Phase 8-Outlet PDU

Figure 1-4 shows examples of the three-phase PDUs that can be used in the SGI UV 2000 system. These PDUs are used to distribute power to the IRUs when the system is configured with three-phase power. The enlarged section shows an optional PDU status interface panel.




Figure 1-4 Three-Phase PDU Examples

System Connections Overview

You can monitor and interact with your SGI UV 2000 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 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. If a CMC is configured with an IP address, you can log directly into it from another PC or workstation on the local area network.

These console connections enable you to view the status and error messages generated by the chassis management controllers in your SGI UV 2000 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.

The following subsections describe the options for establishing and using communication connections to work with your SGI UV 2000 system.

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, see Figure 1-5 on page 8. 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 is factory set to DHCP mode and thus has no fixed IP address and cannot be accessed until an IP address is established. See the subsection “Using DHCP to Establish an IP Address” on page 10 for more information on this topic.

A micro-USB serial connection can be 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 in-rack system console is not used or available.

System Controller Access

Access to the SGI UV 2000 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 1-5.
- A micro-USB serial cable connection to the “Console” port (see Figure 1-6) on the CMC (see note below). See also “Serial Console Hardware Requirements” on page 9.

Note: Each IRU has two chassis management controller (CMC) slots located in the rear of the IRU directly below the cooling fans. Only **one** CMC is supported in each IRU. The CMC slot on the right is the slot that is populated.




Figure 1-5 System Management Node Rear Video Connections




Figure 1-6 UV CMC Connection Faceplate Example

Serial Console Hardware Requirements

The console type and how these console types are connected to the SGI UV 2000 servers is determined by what console option is chosen. If you have an SGI UV 2000 server and wish to use a serially-connected “dumb terminal”, you can connect the terminal via a micro-USB serial cable to the 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 SGI UV 2000

If you have an SGI UV 2000 system and wish to use a serially-connected “dumb terminal”, you can connect the terminal via a micro-USB serial cable to the 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 1-6 on page 8.

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 a micro-USB cable to the CMC's "CONSOLE" port. Configure the serial port as described in "Serial Console Hardware Requirements" on page 9.

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 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 9), 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 SGI UV 2000 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") or other editing method to modify the file.

The configuration file should be modified as follows:

```
BOOTPROTO=dhcp
```

must be commented out, and the entries:

```
BOOTPROTO=static  
IPADDR=  
NETMASK=
```

must be uncommented and set appropriately. Obtain the appropriate values for the IPADDR and NETMASK from your system administrator/IT organization.

```
GATEWAY=<network gateway IP address>  
HOSTNAME=<hostname to use>
```

Note that the GATEWAY and HOSTNAME lines are optional.

Once the changes are made, save the file and reboot the CMC (using the "reboot" command); after it reboots, it will be configured with the specified IP address.

System Control Overview

All SGI UV 2000 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 UV 2000 system control network provides control and monitoring functionality for each compute blade, power supply, and fan assembly in each individual rack unit (IRU) enclosure in the system. The IRU is a 10U-high enclosure that supplies power, cooling, network fabric switching and system control for up to eight compute blades. A single chassis management controller blade is installed at the rear of each IRU.

The *SGI Management Center System Administrator's Guide* (P/N 007-5642-00x) provides more detailed information on using the GUI to administer your SGI UV 2000 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 in concert with the SMN 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 Software User 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.

Communicating with the System

The two primary ways to communicate with and administer the SGI UV 2000 system are through SGI Management Center interface or the UV command line interface (CLI).

The SGI Management Center Graphical User Interface

The SGI Management Center interface is a server monitoring and management system. The SGI Management Center GUI 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 SGI UV 2000 system. Also, see Chapter 2 in this guide for additional reference information on the SGI Management Center interface.

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 128 blades or 4096 processor cores per SSI) depending on set parameters. SGI Management Center services are started and stopped from scripts in */etc/init.d*

SGI Management Center, is commonly installed in */opt/sgi/sgimc* (in a Linux SMN) 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`. In a Windows configured management server the default location would commonly be *c:\Program Files\SGI\SGIMC*.

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 “Power Connections Overview” on page 2 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 **root**).

5. On the command line, enter **mgrcli** 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.

The Command Line Interface

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

Note: The command line interface is virtually the same when used from either the SMN or the CMC. Using the command line interface from the SMN may require that the command be targeted to a specific UV 2000 system if the SMN is managing more than one SSI.

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.

Powering On and Off from the Command Line Interface

The SGI UV 2000 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.

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** [system-SSN]...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 19, the `-uvpower` option must be included with the command to power on.

Power On the 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>
```

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>ls /sysco/bin/help
```

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

```
root@uv-system-smn>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>power on
```

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

Specific 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.

Optional Power On From the CMC Command Line

Most SGI UV 2000 systems come with a system management node (SMN) and there should be few reasons for powering on the system from a CMC. Note that the *SGI UV CMC Software User Guide*, publication (007-5636-00x), describes how to use the chassis management controller (CMC) to monitor and manage SGI UV 2000 systems.

Note: Some basic systems are sold without an SMN and require slightly different administrative procedures. These “SMN-less” systems are restricted in regards to number of blades and router options available.

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”.

Note: The command line interface for the CMC is virtually the same as that for the SMN, with the exception that the CMC does **not** have the ability to target a system when multiple systems are supported from one SMN.

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 SGI UV 2000 systems is determined by what console option is chosen. To monitor or administer a system through the CMC network, you will need to establish a mini-USB serial connection to the CMC. See the information in the following two subsections.

Power On the System Using the CMC Network

You can use a use a direct mini-USB serial connection to the CMC to power on your UV system; note that this process is not the standard way to administer a system. Use the following steps:

1. Establish a 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:

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

2. You can use the IP address of the CMC to perform an SSH 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. x.x.xx  
CMC:r1i1c>
```

This refers to rack 1, IRU 1, CMC.

3. Power up your 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 19).

Optional Power On Using the SMN 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 SMN 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 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 19).

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 on system size and the operating system used, it can take 5 to 10 minutes for the 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 in the following examples:

```
fs0:\> \efi\SuSE\elilo (for SLES 11)
```

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.

```
fs0:\> \efi\redhat\grub (for RHEL 6)
```

The Linux Boot loader for RHEL 6 is called and various SGI configuration scripts are run and the RHEL Service Pack *x* installation program appears.

For information on booting SLES 12 or RHEL 7 (covered by *grubx64.efi*) or creating boot options, see the *SGI UV System Software and Installation Guide*, (P/N 007-5948-00x)

Important: If your SGI UV 2000 system includes a system management node (SMN), be aware that SGI does not support the SLES 12 operating system on the SMN server used with UV 2000.

Power off a UV System

Be sure all clients are logged off prior to power down. Log into the SGI UV system as the root user, and type the following command to stop the operating system:

```
# halt
```

To power down the 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 [system-SSN]...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 [system-SSN]...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 [system-SSN]...toggle I/O reset
-v, --verbose          verbose output

usage: power [-vhow] cycle [system-SSN]...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] [system-SSN]...show power status
-v, --verbose          verbose output
-l, --on               show only blades with on status
-O, --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 [system-SSN]...issue NMI
-o, --override         override partition check
-v, --verbose          verbose output

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

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

Optional SGI Remote Services (SGI RS)

The optional SGI RS system automatically detects system conditions that indicate potential future problems and then notifies the appropriate personnel. This enables you and SGI global support teams to proactively support systems and resolve issues before they develop into actual failures.

SGI Remote Services provides a secure connection to SGI Customer Support - on demand. This can ensure business continuance with SGI systems management and optimization.

SGI Remote Services Primary Capabilities

- 24x7 remote monitoring and data gathering of SGI UV customer systems
- Alerts and notification on changes, failures and potential failures
- Log files immediately available
- Configuration fingerprint
- Secure file transfer
- Optional secure remote access to customer systems

SGI Remote Services Benefits

- Improved uptime and system availability
- Proactive identification of issues before they create an outage
- Increase system stability by monitoring hardware and software version compatibility
- Reduced time to resolve support cases
- Greater operational efficiency
- Less involvement of customer staff during troubleshooting
- Faster support case resolution
- Improved productivity

Proactive potential problem identification can result in higher system availability

Automated Alerts and, in some instances, Case Opening results in faster problem resolution time and less direct involvement required by Customer Support Teams. SGI Remote Services are available for all UV systems and also other specific SGI systems.

SGI Remote Service Operations Overview

An SGI Support Services Software Agent runs on each SGI system at your location, enabling remote system monitoring and secure communication to SGI Support staff. Your basic hardware and software configuration as well as system health information is captured and stored in the Cloud. Figure 1-7 shows an example visual overview of the monitoring and response process.

Cloud intelligence automatically reviews select Event Logs around the clock (every five minutes) to identify potential failure information. If the Cloud intelligence detects a critical Event, it notifies SGI support personnel.

This monitoring requires no changes to customer systems or firewalls as long as the SGI Agent can send HTTPS messages to highly secure Cloud and Global Access Servers. It will also have no impact on customer network or system performance. All communication between SGI global support and customer systems is kept secure using Secure Socket Layer (SSL) encryption. All communication with SGI is initiated from the customer site using HTTPS protocol on port 443.

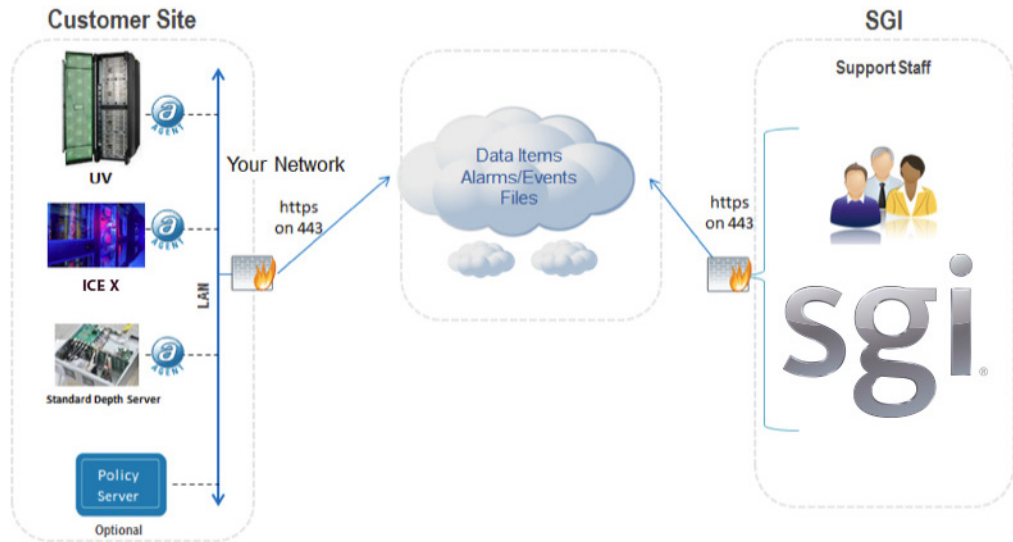


Figure 1-7 SGI Remote Services Process Overview

Optional Components

Besides adding a network-connected system console or basic VGA monitor, you can add or replace the following hardware items on your SGI UV 2000 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 SGI UV 2000 series systems:

- The optional full -height two-slot internal PCIe blade is a dual-node compute blade that supports one full-height x16 PCIe Gen3 card in the top slot and one low-profile x16 PCIe Gen3 card in the lower slot. See Figure 1-8 on page 26 for an example.
- The optional dual low-profile PCIe blade supports two PCIe x16 Gen3 cards. See Figure 1-9 on page 26 for an example.
- The optional external PCIe I/O expansion chassis supports up to four PCIe cards. The external PCIe chassis is supported by connection to a compute blade using an optional host interface card (HIC). Each x16 PCIe enabled blade host interface connector can support one I/O expansion chassis. See Chapter 6 for more details on the optional external PCIe chassis.

Important: PCIe cards installed in an optional two-slot PCIe blade 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(s). Also see “Installing Cards in the 1U PCIe Expansion Chassis” on page 67 for more PCIe related 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 6, “Add or Replace Procedures” for detailed instructions on installing or removing PCIe cards or SGI UV 2000 system disk drives.




Figure 1-8 PCIe Option Blade Example with Full-Height and Low-Profile Slots




Figure 1-9 PCIe Option Blade Example with Two Low-Profile Slots

PCIe Drive Controllers in BaseIO Blade

The SGI UV 2000 system offers a RAID or non-RAID (JBOD) PCIe-based drive controller that resides in the BaseIO blade's PCIe slot. Figure 1-10 shows an example of the system disk HBA controller location in the BaseIO blade.

Note: The PCIe drive controller (upper-right section in Figure 1-10) always hosts the system boot drives.




Figure 1-10 BaseIO Blade and PCIe Disk Controller Example

RAID PCIe Disk Controller

At the time this document was published, the optional RAID controller used in the BaseIO blade PCIe slot is an LSI MegaRAID SAS 9280-8e. This PCIe 2.0 card uses two external SAS control connectors and supports the following:

- RAID levels 0, 1, 5, 6, and 10
- Advanced array configuration and management utilities

- Support for global hot spares and dedicated hot spares
- Support for user-defined stripe sizes: 8, 16, 32, 64, 128, 256, 512, or 1024 KB

The RAID controller also supports the following advanced array configuration and management capabilities:

- Online capacity expansion to add space to an existing drive or a new drive
- No reboot necessary after expansion
- Online RAID level migration, including drive migration, roaming and load balancing
- Media scan
- User-specified rebuild rates (specifying the percentage of system resources to use from 0 percent to 100 percent)
- Nonvolatile random access memory (NVRAM) of 32 KB for storing RAID system configuration information; the MegaRAID SAS firmware is stored in flash ROM for easy upgrade.

Non-RAID PCIe Disk Controller

At publication time, the LSI 9200-8e low-profile PCIe drive controller HBA is the default non-RAID system disk controller for the SGI UV 2000. This drive controller has the following features:

- Supports SATA and SAS link rates of 1.5 Gb/s, 3.0 Gb/s, and 6.0 Gb/s
- Provides two x4 external mini-SAS connectors (SFF-8088)
- The HBA has onboard Flash memory for the firmware and BIOS
- The HBA is a 6.6-in. x 2.713-in., low-profile board
- Supports eight-lane, full-duplex PCIe 2.0 performance
- The HBA has multiple status and activity LEDs and a diagnostic UART port
- A x8 PCIe slot is required for the HBA to operate within the system

System Control

This chapter describes the general interaction and functions of the overall SGI UV 2000 system control. System control parameters depend somewhat on the overall size and complexity of the SGI UV 2000 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 UV 2000 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 UV 2000 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 UV system control network is a private, closed network. It should not be reconfigured in any way to change it from the standard SGI UV factory installation. It should not be directly connected to any other network. The 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 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 31) 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 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 interconnects. This maintains controller configuration and topology information between all controllers in an SSI. Figure 2-2 on page 32 shows an example system control network using an optional and separate (remote) workstation to monitor a single-rack SGI UV 2000 system. It is also possible to connect an optional PC or (in-rack) console directly to the SMN, see Figure 2-4 on page 36.

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 user'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 UV 2000 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. A CMC port from each chassis management controller connects to a dedicated Ethernet switch that provides a synchronous clock signal to all the CMCs in an SSI.

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 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 as well as 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 UV 2000 LAN-attached System Control Network Example

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 SGI UV 2000 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 network. 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 an SMN and multiple 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 Software User 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 35) 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 UV 2000 series system. The major components of the SGI UV 2000 series systems are described and illustrated.

The SGI UV 2000 series is a family of multiprocessor distributed shared memory (DSM) computer systems that can scale from 16 to 2,048 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 SGI UV 2000 series system in your lab or server room. Each 42U SGI rack holds one to four 10-U high enclosures that support up to eight compute/memory and I/O sub modules known as “blades.” These blades contain 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 SGI UV 2000 IRU enclosures.

This chapter consists of the following sections:

- “System Models” on page 39
- “System Architecture” on page 41
- “System Features” on page 43
- “System Components” on page 48

Figure 3-1 shows the front view of a single-rack SGI UV 2000 system.




Figure 3-1 SGI UV 2000 Single-Rack System Example

System Models

The basic enclosure within the SGI UV 2000 system is the 10U high “individual rack unit” (IRU). The IRU enclosure contains up to eight compute blades connected to each other via a backplane. Each IRU has ports that are brought out to external NUMALink 6 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 SGI UV 2000 server system requires a minimum of one BaseIO equipped blade for every 2,048 processor cores. Higher core counts in an SSI may be available in future releases, check with your SGI sales or service representative for the most current information.

Note: Special systems operated without a system management node (SMN) must have an optional external DVD drive available to connect to the BaseIO blade.

Figure 3-2 shows an example of how IRU placement is done in a single-rack SGI UV 2000 server.

The system requires a minimum of one 42U tall rack with three single-phase power distribution unit (PDU) plugs per IRU installed in the rack. Three outlets are required to support each power shelf. There are three power supplies per power shelf and two power connections are required for an SMN.

You can also add additional PCIe expansion enclosures or RAID and non-RAID disk storage to your server system. Power outlet needs for these options should be calculated in advance of determining the number of outlets needed for the overall system.




Figure 3-2 SGI UV 2000 IRU and Rack

System Architecture

The SGI UV 2000 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 processor count, memory, and I/O independently in each rack. Note that a maximum of 2,048 cores are supported on a single-system image (SSI). Larger SSI configurations may be offered in the future, contact your SGI sales or service representative for additional information.

The system architecture for the SGI UV 2000 system is a sixth-generation NUMAflex DSM architecture known as NUMALink 6 or NL6. In the NUMALink 6 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 and between each 10U IRU enclosure.

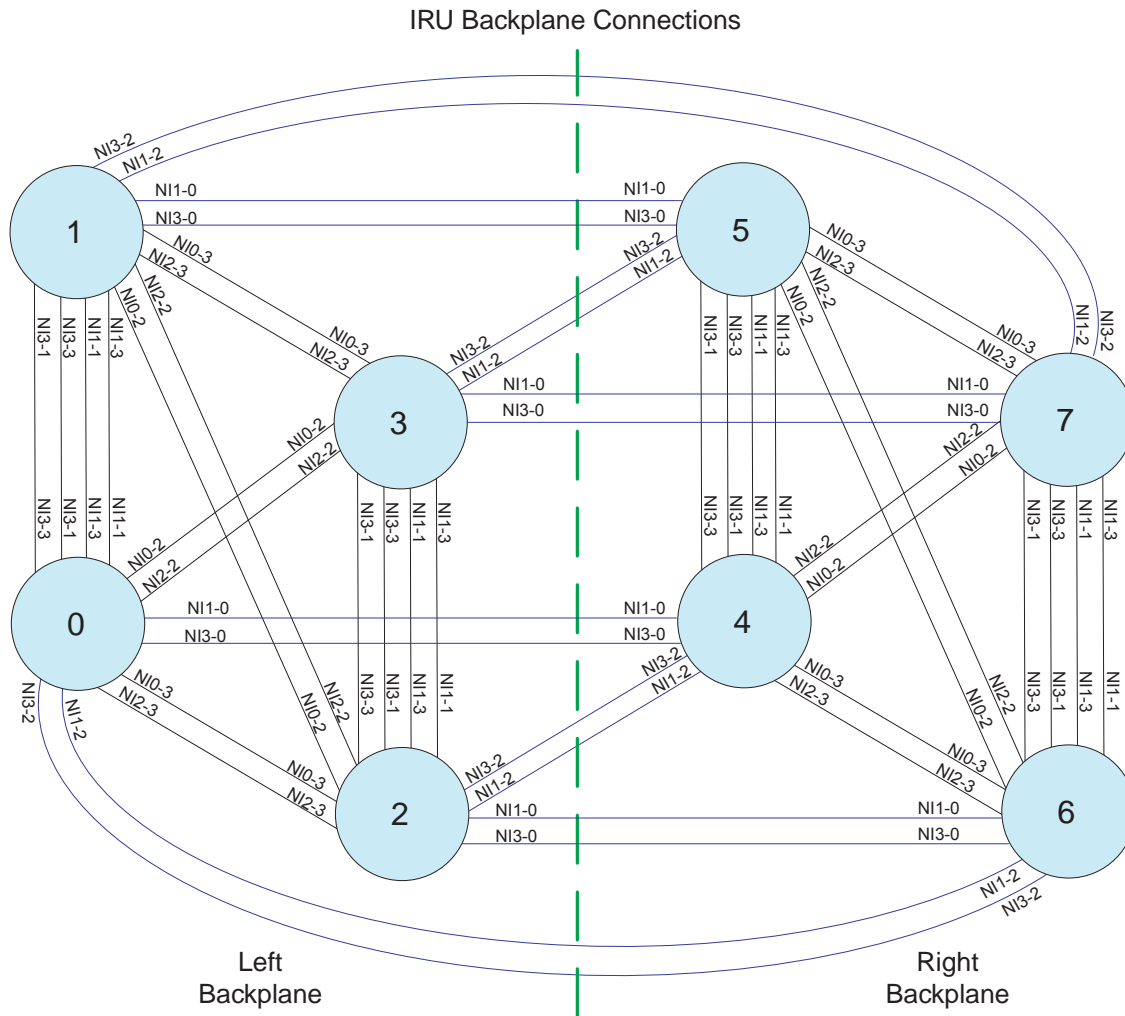
The basic expansion building block for the NUMALink interconnect is the processor node; each processor node consists of a dual-Hub ASIC (also known as a HARP) and two eight-core processors with on-chip secondary caches. The Intel processors are connected to the dual-Hub ASIC via quick path interconnects (QPIs). Each dual-HUB ASIC is also connected to the system's NUMALink interconnect fabric through one of sixteen NL6 ports.

The dual-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 42 shows a functional block diagram of the SGI UV 2000 series system IRU.

System configurations of up to eight IRUs can be constructed without the use of external routers. Routerless systems can have any number of blades up to a maximum of 64. Routerless system topologies reduce the number of external NUMALink cables required to interconnect a system.

External optional routers are needed to support multi-rack systems with more than four IRUs, see Chapter 5, “Optional Octal Router Chassis Information” for more information.



Note: this drawing only shows the cabling between the compute blades within the IRU. External cabling (cabling that exits the IRU) is not shown.

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

System Features

The main features of the SGI UV 2000 series server systems are discussed in the following sections:

- “Modularity and Scalability” on page 43
- “Distributed Shared Memory (DSM)” on page 43
- “Chassis Management Controller (CMC)” on page 45
- “Distributed Shared I/O” on page 45
- “Reliability, Availability, and Serviceability (RAS)” on page 46

Modularity and Scalability

The SGI UV 2000 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, MIC/GPU capability 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 SGI UV 2000 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 is directly connected to its resident socket, the memory is referred to as *local memory*. Figure 3-4 on page 44 shows a conceptual block diagram of the blade’s memory, compute and I/O pathways.
- If a processor needs to access memory located in another socket, or on another blade 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*.




Figure 3-4 Blade Node Block Diagram Example

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 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 blade slot is also connected to the CMC. This connection is used to support a BaseIO riser card. Only one BaseIO is supported in an SSI. The BaseIO must be the first blade (lowest) in the SSI.

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 SGI UV 2000 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 SGI UV 2000 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 SGI UV 2000 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 SGI Remote Services (SGI RS) program which automatically detects system conditions that indicate potential future problems and then notifies the appropriate personnel. This enables you and the SGI global support team to provide pro-active and rapid issue resolution. See the section “Optional SGI Remote Services (SGI RS)” in Chapter 1 for more information.
- 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 SGI UV 2000 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 SGI UV 2000 system. Up to four 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 three power supplies, 2-8 compute/memory blades, BaseIO and other optional riser enabled blades for the SGI UV 2000. The enclosure is 10U high. Figure 3-5 on page 49 shows the SGI UV 2000 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 BMC Ethernet port and three USB ports. System disks are always controlled by a PCIe disk controller installed in the BaseIO blade's PCIe slot. Figure 3-6 on page 50 shows a front-view example of the BaseIO blade.

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. The SGI UV 2000 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 normally 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. Either two half-height or one half-height and one full-height cards are supported.
- **MIC/GPU PCIe enabled compute blade.** This blade supports one optional MIC or GPU card in the upper slot via PCIe interface to the bottom node board. Option cards are limited, check with your SGI sales or service representative for available types supported.
- **External PCIe enabled compute blade.** This PCIe enabled board is used in conjunction with an external PCIe expansion enclosure. A x16 adapter card connects from the blade to the external 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 SGI UV 2000 IRU System Components Example

Optional BaseIO SSDs

The BaseIO blade can be configured with one or two internal 1.8-inch solid state drives (SSDs). The SSDs can be configured as JBOD or RAID1. The RAID1 SSD pair is a software RAID1 and two SSDs must be ordered with the system BaseIO to enable this configuration.




Figure 3-6 BaseIO Riser Enabled Blade Front Panel Example

MIC/GPU Enabled Compute Blade

The single-socket MIC/GPU enabled compute blade has one single socket node blade and supports one PCIe accelerator card. The MIC/GPU enabled compute blade has the following features:

- One HARP ASIC based board assembly with twelve NUMALink six (NL6) ports that connect the blade to the backplane and four NL6 ports connecting the blade to external QSFP ports.
- Specialized connectors support the connection to both the bottom compute node and the top MIC or GPU board assembly.
- One Bottom compute node board assembly with a single processor socket also supports eight memory DIMM slots (1600 MT/s memory DIMMs).
- One baseboard management controller (BMC) and one x16 Gen3 PCIe full-height double-wide slot that supports a single MIC or GPU accelerator card.
- The accelerator card connects directly to the bottom compute board assembly via a ribbon cable and draws power from the IRU backplane.




Figure 3-7 MIC/GPU Enabled Compute Blade Example Front View

Bay (Unit) Numbering

Bays in the racks are numbered using standard units. A standard unit (SU) or unit (U) is equal to 1.75 inches (4.445 cm). Because IRUs occupy multiple standard units, IRU locations within a rack are identified by the bottom unit (U) in which the IRU resides. For example, in a 42U rack, an IRU positioned in U01 through U10 is identified as U01.

Rack Numbering

Each rack is numbered with a three-digit number sequentially beginning with 001. A rack contains IRU enclosures, optional mass storage enclosures, 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 2000 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) SGI UV 2000 racks in the following sections:

- “Overview” on page 53
- “SGI UV 2000 Series Rack (42U)” on page 54
- “SGI UV 2000 System Rack Technical Specifications” on page 58

Overview

At the time this document was published only the tall (42U) SGI UV 2000 rack (shown in Figure 4-2) was available from the SGI factory for use with the SGI UV 2000 systems. Other racks may be available to house the system IRUs, check with your SGI sales or service representative for information.

SGI UV 2000 Series Rack (42U)

The tall rack (shown in Figure 4-1 on page 55) 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 56).

Note: The front door and rear door locks are keyed differently. The optional water-chilled rear door (see Figure 4-3 on page 57) 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.** Up to 12 outlets are required for a single-rack IRU system as follows:
 - Allow three outlets for the first IRU
 - Two outlets for a maintenance node SMN (server)
 - Two outlets for each storage or PCIe expansion chassis
 - Allow three more outlets for each additional IRU in the system

Note that an eight outlet single-phase PDU may be used for the system management node and other optional equipment.

Each three-phase power distribution unit has 9 outlet connections.




Figure 4-1 SGI UV 2000 Series Rack Example




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




Figure 4-3 Optional Water-Chilled Cooling Units on Rear of SGI 42U Rack

SGI UV 2000 System Rack Technical Specifications

Table 4-1 lists the technical specifications of the SGI UV 2000 series tall rack.

Table 4-1 Tall 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

Optional Octal Router Chassis Information

This chapter describes the optional NUMAlink router technology available in SGI UV 2000 systems consisting of two or more racks. This router technology is available in an enclosure “package” known as the Octal Router Chassis (ORC). This optional ORC chassis can be mounted on the top of the SGI UV 2000 rack. NUMAlink 6 advanced router technology reduces UV 2000 system data transfer latency and increases bisection bandwidth performance. Router option information is covered in the following sections:

- “Overview” on page 59
- “SGI UV 2000 Series NUMAlink Octal Router Chassis” on page 60
- “SGI UV 2000 External NUMAlink System Technical Specifications” on page 62

Overview

At the time this document was published, external NUMAlink router technology was available to support from 2 to 512 SGI UV 2000 racks. Other “internal” NUMAlink router options are also available for high-speed communication between smaller groups of SGI UV 2000 racks. For more information on these topics, contact your SGI sales or service representative.

The standard routers used in the SGI UV 2000 systems are the NL6 router blades located internally to each IRU. Each of these first level routers contain a single 16-port NL6 HARP router ASIC. Twelve ports are used for internal connections (connecting blades together), the remaining four ports are used for external connections. The NUMAlink ORC enclosure is located at the top of each SGI UV 2000 rack equipped with the option. Each top-mounted NUMAlink ORC enclosure contains one to eight 16-port HARP ASIC based router boards. Each of these router boards has a single NL6 HARP router ASIC. This is the same router ASIC that is used in the NL6 router blades installed inside the system IRUs.

Note that the ORC chassis also contains a chassis management controller (CMC) board, two power supplies and its own cooling fans.

SGI UV 2000 Series NUMAlink Octal Router Chassis

The NUMAlink 6 ORC router is a 7U-high fully self contained chassis that holds up to eight 16-port NL6 router blade assemblies. Figure 5-1 shows an example rear view of the ORC with no power or NUMAlink cables connected.

The NUMAlink ORC is composed of the following:

- 7U-high chassis
- 4 or 8 HARP based router blade assemblies
- Cooling-fan assemblies
- Chassis Management Controller (CMC)/power supply assembly (with two power supplies)




Figure 5-1 SGI UV 2000 Optional NUMAlink ORC (Rear View)




Figure 5-2 SGI UV 2000 Optional ORC Chassis Example (Front View)

Note: The NUMAlink unit's CMC is connected to the CMC in each IRU installed in the rack.

SGI UV 2000 External NUMAlink System Technical Specifications

Table 5-1 lists the basic technical specifications of the SGI UV 2000 series external NUMAlink ORC chassis.

Table 5-1 External NUMAlink Technical Specifications

Characteristic	Specification
Height	7U or 12.25 in. (31.1 cm)
Width	13.83 in. (35.13 cm)
Depth	14.66 in (37.24 cm)
Top-mount NUMAlink router weight (approximate)	53 lbs. (24.1 kg) not including attached cables
Power supply	Three 760-Watt hot-plug power supplies
Voltage range Nominal	North America/International 100-240 VAC /230 VAC
Frequency Nominal	North America/International 60 Hz /50 Hz
Tolerance range	47-63 Hz
Phase required	Single-phase
Power cables	6.5 ft. (2 m) pluggable cords

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 63
- “Adding or Replacing PCIe Cards in the Expansion Enclosure” on page 66
- “Removing and Replacing an IRU Enclosure Power Supply” on page 70
- “Replacing a System Fan (Blower)” on page 72
- “Replacing a System Disk” on page 75

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 64
- “Returning the System to Operation” on page 64



Warning: To avoid problems that could void your warranty, your SGI or other approved service provider should perform all the setup, addition, or replacement of parts, cabling, and service of your SGI UV 2000 system, with the exception of the following:

- 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 1U PCIe expansion chassis.
- Using the ESI/ops panel (operating panel) on optional mass storage.
- Removing and replacing IRU power supplies, cooling fans and system disk drives.

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-On and Off From the SGI Management Center Interface” 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-On and Off From the SGI Management Center Interface” 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 7, 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 UV 2000.

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

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

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

- PCI Express uses a point-to-point serial interface vs. a shared parallel bus interface used in older PCI/PCI-X technology
- PCIe hardware connectors are not compatible with PCI/PCI-X, (see Figure 6-1)
- Potential sustained throughput of x16 PCI Express is approximately four times that of the fastest PCI-X throughputs




Figure 6-1 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 UV 2000 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 6-1 shows this concept.

For information on which slots in the PCIe expansion chassis support what lane levels, see Table 6-2 on page 67.

Table 6-1 SGI UV 2000 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 95.

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 6-2 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 6-2.
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 6-2 The PCIe Expansion Enclosure




Figure 6-3 Card Slot Locations

Removing and Replacing an IRU Enclosure Power Supply

To remove and replace power supplies in an SGI UV 2000 IRU, you do not need any tools. Under most circumstances a single power supply in an IRU can be replaced without shutting down the enclosure or the complete system. In the case of a fully configured (loaded) enclosure, this may not be possible.

Caution: The body of the power supply may be hot; allow time for cooling and handle with care.

Use the following steps to replace a power supply in the blade enclosure box:

1. Open the front door of the rack and locate the power supply that needs replacement.
2. Disengage the power-cord retention clip and disconnect the power cord from the power supply that needs replacement.




Figure 6-4 Removing an Enclosure Power Supply

3. Press the retention latch of the power supply toward the power connector to release the supply from the enclosure, see Figure 6-4 on page 70.
4. Using the power supply handle, pull the power supply straight out until it is partly out of the chassis. Use one hand to support the bottom of the supply as you fully extract it from the enclosure.
5. Align the rear of the replacement power supply with the enclosure opening, see Figure 6-5.
6. Slide the power supply into the chassis until the retention latch engages - you should hear an audible click.
7. Reconnect the power cord to the supply and engage the retention clip.

Note: If AC power to the rear fan assembly is disconnected prior to the replacement procedure, all the fans will come on and run at top speed when power is reapplied. The speeds will readjust when normal communication with the IRU's CMC is fully established. See the section "Replacing a System Fan (Blower)" on page 72 for additional information on fan operation.




Figure 6-5 Replacing an IRU Enclosure Power Supply

Replacing a System Fan (Blower)

Chassis cooling for each UV 2000 IRU is provided by nine, heavy duty, counter-rotating fans. Each fan unit is made up of two fans joined back-to-back, which rotate in opposite directions. The counter-rotating action increases airflow and dampens vibration levels.

If one fan fails, the remaining fans will ramp up to full speed and the overheat/fan fail LED on the control panel will illuminate (the system can continue to run with a failed fan).

Note that each power supply in the system is cooled by an individual internal cooling fan.

The IRU enclosure cooling fans are positioned back-to-back with the blades in the UV enclosure. You will need to access the rack from the back to remove and replace a fan. The enclosure's system controller issues a warning message when a fan is not running properly. This means the fan RPM level is not within tolerance. Depending on system configuration, when a cooling fan fails, some or all of the following things happen:

1. The system console will show a warning indicating the rack and enclosure position

```
001c01 L2> Fan (number) warning limit reached @ 0 RPM
```
2. A line will be added to the L1 system controller's log file indicating the fan warning.
3. If optional SGI Embedded Support Partner (ESP) is used, a warning message will be sent to it also.

The chassis management controller (CMC) monitors the temperature within each enclosure. If the temperature increases due to a failed fan, the remaining fans will run at a higher RPM to compensate for the missing fan. The system will continue running until a scheduled maintenance occurs.

The fan numbers for the enclosure (as viewed from the rear) are shown in Figure 6-6 on page 73.

Note that under most circumstances a fan can be replaced while the system is operating. You will not need any tools to complete the replacement procedure.




Figure 6-6 UV 2000 Rear Fan Assembly (Blowers)

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

1. Open the rack's rear door and identify the fan that has failed, see Figure 6-7.
2. Grasp the failed blower assembly handle and pull the fan straight out.




Figure 6-7 Removing a Fan From the Rear Assembly

3. Slide a new blower assembly completely into the open slot until the fan-interconnect engages and the new unit is flush with the rear of the assembly, see Figure 6-8 on page 75.
4. Confirm that the new fan is operational and close the rack's rear door.

Note: If you disconnected the AC power to the rear fan assembly prior to the replacement procedure, all the fans will come on and run at top speed when power is reapplied. The speeds will readjust when normal communication with the blade pair enclosure CMC is fully established.




Figure 6-8 Replacing an Enclosure Fan

Replacing a System Disk

The dual-disk drive blade is used to house two system disk drive bays as shown in Figure 6-9 on page 76. 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 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 6-9 on page 76 for an example.
2. Pull the handle outward until the locking mechanism is fully cleared.
3. Grasp the disk drive by the side and extract it from the disk riser blade.




Figure 6-9 Removing the UV 2000 System Disk Example

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 with the drive assembly oriented as shown in Figure 6-10 on page 77.
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 inward until the handle clicks into place (this completes the insertion of the disk drive).




Figure 6-10 Replacing the UV 2000 System Disk Example

Troubleshooting and Diagnostics

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

- “Troubleshooting Chart” on page 80
- “LED Status Indicators” on page 81
- “Accessing Online Support Information and Services” on page 83

Troubleshooting Chart

Table 7-1 lists recommended actions for problems that can occur. To solve problems that are not listed in this table, use the SGI Customer Portal website or contact your SGI system support representative. For more information about the SGI Customer Portal options, see “Accessing Online Support Information and Services” on page 83. For an international list of SGI support centers, see: <http://www.sgi.com/support/supportcenters.html>

Table 7-1 Troubleshooting Chart

Problem Description	Recommended Action
The system will not power on.	<p>Ensure that the power cords of the IRU are seated properly in the power receptacles.</p> <p>Ensure that the PDU circuit breakers are on and properly connected to the wall source.</p> <p>If the power cord is plugged in and the circuit breaker is on, contact your technical support organization.</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>
No status LEDs are lighted on an individual blade.	<p>Confirm the blade is firmly seated in the IRU enclosure.</p> <p>See also “Compute/Memory Blade LEDs” on page 82.</p>
The system will not boot the operating system.	<p>Contact your SGI support organization:</p> <p>http://www.sgi.com/support/supportcenters.html</p>
The amber (yellow) status LED of an IRU power supply is lit or the LED is not lit at all. See Table 7-2 on page 81.	<p>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 in. If the green LED does not light, contact your support organization.</p>
The PWR LED of a populated PCIe slot is not illuminated.	<p>Reseat the PCI card.</p>
The Fault LED of a populated PCIe slot is illuminated (on).	<p>Reseat the card. If the fault LED remains on, replace the card.</p>
The amber LED of a disk drive is on.	<p>Replace the disk drive.</p>

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 7-2 for a complete list.

Table 7-2 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 seven LED indicators visible behind the perforated sheetmetal of the blade.

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

- System power good green LED
- BMC heartbeat green LED
- Blue unit identifier (UID) LED
- BMC Ethernet 1 green LED
- BMC Ethernet 0 green LED
- Green 3.3V auxiliary power LED
- Green 12V power good 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. Figure 7-1 shows the locations of the blade LEDs.




Figure 7-1 UV Compute Blade Status LED Locations Example

Accessing Online Support Information and Services

Multiple levels of service, support and troubleshooting information are available through:

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

The site provides access to a variety of on-line and in-person support resources as follows:

SGI Customer Portal

Note: The SGI Customer Portal website is a part of the overall SGI support provided and can be directly accessed at: <http://support.sgi.com>

When you log in to the SGI Customer Portal website, you can access current SGI hardware and software customer manuals. See the section “Obtaining SGI Publications” on page xviii of this document for access instructions. Other Customer Portal offerings include:

- Software downloads and patches
- SGI Knowledgebase
- Service call logging and tracking
- Submission of technical questions

Note that the customer document access on the SGI Customer Portal has replaced the SGI TechPubs Library.

Technical Assistance

This website topic area covers how and where to obtain direct technical assistance from SGI or our contracted service partners and includes:

- Customer Support Centers in North America, the Asia-Pacific zone, Near-East and Europe
- Authorized support partners in areas not directly supported by an SGI Customer Support Center

Other Resources

Topics covered under Other Resources include:

- Support Services descriptions
- Product Warranties
- Customer Replaceable Units (CRU)
- Warranty and Support Contract Transfers
- Service Contracts
- Software Keys

SGI Warranty Levels

The complete SGI Electronic Support services are available to customers who have a valid SGI Warranty or support contract. Additional electronic services may become available after publication of this document; contact your SGI sales or support representative for the latest information.

To purchase a support contract that allows you to use the complete SGI Electronic Support services (such as SGI RS), contact your SGI sales representative.

Optional SGI Remote Services (SGI RS)

The optional SGI RS system automatically detects system conditions that indicate potential failure, see the section “Optional SGI Remote Services (SGI RS)” in Chapter 1 for an overview and description.

Technical Specifications and Pinouts

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

- “System-level Specifications” on page 85
- “Physical Specifications” on page 86
- “Environmental Specifications” on page 87
- “Power Specifications” on page 88
- “I/O Port Specifications” on page 89

System-level Specifications

Table A-1 summarizes the SGI UV 2000 system configuration ranges. Note that while each compute/memory blade holds two processor sockets (one per node board); each processor in a socket supports multiple processor “cores”.

Table A-1 SGI UV 2000 System Configuration Ranges

Category	Minimum	Maximum
Processors	Eight processors (2 blades) ^a	256 processors (per SSI)
Individual Rack Units (IRUs)	1 per rack	4 per rack
Blades per IRU	2 per IRU	8 per IRU
Compute/memory blade DIMM capacity	8 DIMMs per blade	16 DIMMs per blade
CMC units	1 per IRU	1 per IRU
Number of baseIO riser enabled blades	One per SSI	One per SSI

a. Dual-node blades support up to four processors per blade.

Physical Specifications

Table A-2 shows the physical specifications of the SGI UV 2000 system.

Table A-2 SGI UV 2000 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: 52 in. (132.1 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
10U-high Individual Rack Unit (IRU) enclosure specifications	Dimensions: 17.5 in high x 19 in (flange width) wide x 27 in deep 44.45 cm high x 48.26 cm wide x 68.58 cm deep

Note: Racks equipped with optional top-mounted NUMAlink (ORC) routers have an additional weight of 53 lbs. (24.1 kg) plus the weight of additional cables.

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 SGI UV 2000 (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:

- “BaseIO VGA Port Information” on page 89
- “Ethernet Port” on page 91
- “Serial Ports” on page 92
- “USB Type A Connector” on page 94

BaseIO VGA Port Information

The 15-pin VGA port on the baseIO blade (see Figure A-1) has the following features:

- Server-class 2D hardware acceleration support with integrated 24-bit RAMDAC
- Display resolution up to 1600 x 1200 @ 60Hz
- Up to 128 Mbytes DDRII memory interface support
- A single USB keyboard/mouse is supported by the baseIO blade connectors

The baseIO VGA interface (see Table A-5 on page 90) can be used for all basic interaction with your SGI UV 2000 system. Note that it does **not** provide a direct interconnect to the system CMC or SMN.




Figure A-1 VGA Port Pinouts

Table A-5 VGA Pin Functions

Pin Number	Function
1	Red
2	Green
3	Blue
4	N/C
5	Ground
6	Ground
7	Ground
8	Ground
9	Ground
10	Ground
11	N/C
12	DDCDAT
13	HSYNC
14	VSYNC
15	DDCCLK

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-2 shows the Ethernet port.




Figure A-2 Ethernet Port

Table A-6 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-6 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-3 shows a serial port.




Figure A-3 Serial Port Connector

Table A-7 shows pinout assignments for the 9-pin male DB-9 connector.

Table A-7 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-4 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-8 lists the pin assignments for the USB type A connector.




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

Table A-8 Pin Assignments for USB Type A Connector

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

Safety Information and Regulatory Specifications

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

- “Safety Information” on page 95
- “Regulatory Specifications” on page 97

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 97
- “CE Notice and Manufacturer’s Declaration of Conformity” on page 97
- “Electromagnetic Emissions” on page 98
- “Shielded Cables” on page 100
- “Electrostatic Discharge” on page 100
- “Laser Compliance Statements” on page 101
- “Lithium Battery Statement” on page 102

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 når denna del är öppnad åå 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.
