

SGI® Virtu VN200 Graphics Node
Hardware User's Guide

007-5483-001

COPYRIGHT

© 2002-2008 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.

TRADEMARKS AND ATTRIBUTIONS

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

ATI is a trademark of Advanced Micro Devices, Inc. InfiniBand is a registered trademark of the InfiniBand Trade Association. Intel, Itanium, and Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries. InfiniBand is a trademark of InfiniBand Trade Association. Internet Explorer is a registered trademark of Microsoft Corporation. Java and Java Virtual Machine are trademarks or registered trademarks of Sun Microsystems, Inc. Linux is a registered trademark of Linus Torvalds, used with permission by SGI. Mellanox is a registered trademark of Mellanox Technologies, Inc. in the United States and/or other countries. NVIDIA, the NVIDIA logo, Quadro, and TwinView are registered trademarks or trademarks of NVIDIA Corporation in the United States and/or other countries. Phoenix and PhoenixBIOS are registered trademarks of Phoenix Technologies Ltd. Red Hat and all Red Hat-based trademarks are trademarks or registered trademarks of Red Hat, Inc. in the United States and other countries. ScaliManage is a trademark of Scali Inc. Supermicro and Supero Doctor III are registered trademarks or trademarks of Super Micro Computer, Inc. SUSE LINUX and the SUSE logo are registered trademarks of Novell, Inc. UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company, Ltd. Windows is a registered trademark of Microsoft Corporation.

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

Record of Revision

Version	Description
001	April, 2008. Original printing.

Contents

About This Guide	
Audience	xiii
Structure of this Guide	xiii
Related Publications	xiv
Conventions	xv
Product Support	xv
Reader Comments	xvi
1 SGI Virtu VN200 Overview	
Product Description	2
Warnings and Cautions	4
Site Planning	5
Selecting a Delivery Route	5
Selecting a Final Location	6
Unpacking the Rack	7
Unpacking the Graphics Nodes	9
Installing the Rack Shelf	10
System Configurations and Cabling.	11
Graphics Node Dimensions	14
Electrical Requirements	15
Power Consumption for the Graphics Node	15
Environmental Requirements	16
Acoustics	16
Electromagnetic Interference	16
Electrostatic Discharge	17

Regulatory Specifications	17
Server Model Number	17
CE Notice and Manufacturer's Declaration of Conformity	18
Electromagnetic Emissions	18
FCC Notice (USA Only)	18
Industry Canada Notice (Canada Only)	19
VCCI Class A Notice (Japan Only)	19
Shielded Cables	19
2 System Features	
Server Racks.	22
Virtu VN200 Graphics Node	23
Front Panel Components and Status LEDs	24
Rear Panel Components	25
Node Board	26
Node Board Features	27
Jumper Quick-reference	30
Memory Overview	31
Memory DIMM Pairs.	33
Memory Configuration Rules	33
Memory Part Numbers	33
AOC-SIMSO (BMC) Card.	35
Power Supplies	37
Cooling Fans	38
System Disk Drives	39
Graphics Cards.	40
Accelerated Driver Installation	41

3	System Procedures	
	Connecting a Digital Monitor	43
	Connecting a Keyboard and Mouse	44
	Configuring Firmware and Drivers	45
	Power On/Off the Cluster	45
	Powering Off Manually	46
	Verifying the Graphics Card Installation for Linux	47
	Starting the X Window System	48
	Commands	48
	Documentation	49
	Start Graphics	49
	Stop Graphics	49
	Virtual Consoles	49
	xorg.conf Configuration File.	50
	Monitor Section	53
	Modelines.	53
	Device Section	54
	TwinView.	54
	TwinView Edge-matched Display	55
	TwinView Edge Blending	56
	Stereoscopic Viewing	56
	Setting up Multiple X Servers on a Single Graphics Card	57
	ServerLayout Section	59
	Check the NVIDIA Driver	60
	Re-installation of the Graphics Drivers.	60
	Upgrading Drivers	61
	Troubleshooting	61
	General.	61
	Graphics	62
	BMC Configuration	62
	Power	63
	Cooling.	63

Hardware Replacement Procedures	64
Before You Begin	64
Tools and Supplies Needed	64
Removing the Front Bezel	64
Removing the Graphics Node from the Rack	65
Removing the Side Cover.	66
Replacing the Power Supplies	67
Replacing a Disk Drive	68
Replacing the Chassis Fans	69
Replacing the Graphics Card.	69
Replacing the AOC-SIMSO (BMC) Daughter Card	71
Removing the PCI Riser Board	72
Removing the Air Baffles.	73
Replacing Node Board Memory DIMMs	74
Replacing Node Board Processors	75
Install a Processor	77
Install a Heatsink	81
Remove the Heatsink	83
Replace the Thermal Interface Material.	83
Removing TIM from a Processor	84
Removing TIM from a Heatsink	87
Replacing the Node Board	90
Replacing the Control Panel	91

Figures

Figure 1-1	SGI Altix 1300 Cluster with Virtu VN200 Graphics Nodes.	1
Figure 1-2	Unpacking the Rack.	8
Figure 1-3	Unpacking the Graphics Nodes	9
Figure 1-4	Rack Shelf	10
Figure 1-5	Altix XE Cluster Configuration	12
Figure 1-6	SGI Altix ICE Cluster Configuration	13
Figure 1-7	Graphics Node Dimensions	14
Figure 2-1	SGI Virtu VN200 Graphics Nodes.	21
Figure 2-2	SGI Virtu VN200 - Front View.	22
Figure 2-3	Virtu VN200 Graphics Node	23
Figure 2-4	Front Panel Components	24
Figure 2-5	Rear Panel Components	25
Figure 2-6	VN200 Node Board	26
Figure 2-7	Node Board Component Locations	26
Figure 2-8	Node Board Block Diagram	29
Figure 2-9	Memory Block Diagram	32
Figure 2-10	Node Board DIMM Layout	32
Figure 2-11	AOC-SIMSO (BMC) Daughter Card	35
Figure 2-12	Redundant Power Supplies	37
Figure 2-13	Cooling Fans.	38
Figure 2-14	System Disk Drives	39
Figure 2-15	Graphics Card	40
Figure 3-1	SGI Altix XE Cluster Configuration	44
Figure 3-2	Xorg.conf File Organization	50
Figure 3-3	TwinView Edge-Matched Display	55
Figure 3-4	TwinView Edge Blending with 10% Overlap	56
Figure 3-5	Front Bezel Removal	64
Figure 3-6	Removing the Graphics Node from a Rack	65
Figure 3-7	Side Cover Removal	66
Figure 3-8	Power Supply Removal	67

Figure 3-9	Disk Drive Removal	68
Figure 3-10	Cooling Fans Removal	69
Figure 3-11	Graphics Card Removal	70
Figure 3-12	BMC Daughter Card Removal	71
Figure 3-13	PCI Riser Board Removal	72
Figure 3-14	Air Baffle Removal	73
Figure 3-15	DIMM Removal	74
Figure 3-16	Node Board Removal	75
Figure 3-17	Install a CPU Heatsink	82
Figure 3-18	Air Baffles	90
Figure 3-19	Control Panel Removal	91

Tables

Table 1-1	SGI Virtu VN200 System Configurations	11
Table 1-2	Graphics Node Power Consumption.	15
Table 2-1	Node Board Jumper Definitions	30
Table 2-2	Node Board Connector Definitions	30
Table 2-3	Memory DIMM Pairs	33
Table 2-5	DIMM Configurations	34
Table 2-4	Memory Sizes and SGI Part Numbers	34
Table 2-6	Power Supply Specifications	37
Table 2-7	Graphics Card Specifications	40

About This Guide

This guide provides information about how to install, configure, administrate, and repair an SGI Virtu VN200.

Note: Refer to “Related Publications” on page xiv, for a listing of relevant SGI Technical Publications that can provide more detailed information about SGI cluster head nodes, compute nodes, system software, and rack hardware.

Audience

This guide is intended for system administrators. Use this guide to:

- Unpack and install the hardware.
- Gain a basic understanding of the SGI Virtu VN200 hardware and software components.

Structure of this Guide

This guide contains the following chapters:

- Overview — Introduces the SGI Virtu VN200 product and discusses the system level topics such as site planning, installation, and cabling.
- System Features — Describes the components of the SGI Virtu VN200
- Procedures — Describes the procedures for configuring the system and replacing hardware.
- Refer to *SGI Altix XE320 System User's Guide*, publication number 007-5466-00x, for more detailed information about node board BIOS settings, power-on self test (POST) codes, and the baseboard management controller (BMC) intelligent platform management interface (IPMI) management utility.

Related Publications

Your SGI Virtu VN200 system ships with a variety of hardware and software documents in both hard copy and soft copy formats. Hard copy documents are in the packing box and soft copy documents are located on your system hard disk in `/usr/local/Factory-Install/Docs`. Additional third-party documentation may be shipped on removable media (CD/DVD) included with your shipment. You will need to consult the following documentation while using this guide:

- *SGI Altix XE Cluster Quick Reference Guide*, publication number 007-4979-00x
This document is intended as an overview of some of the common operations that system administrators may have to perform to set-up, boot, re-configure (upgrade) or troubleshoot the SGI Altix XE cluster.
- *SGI Altix XE320 System User's Guide*, publication number 007-5466-00x
This guide provides instructions administrating SGI Altix XE320 compute nodes in the SGI Virtu VN200. The SGI Virtu VN200 uses the same node board as the SGI Altix 320 server, so you can also use this guide for information about BIOS settings, power-on self test (POST) codes, and the baseboard management controller (BMC) intelligent platform management interface (IPMI) management utility.
- *SGI Altix XE250 System User's Guide*, publication number 007-5467-00x
This guide provides instructions administrating SGI Altix XE250 nodes in an SGI cluster.
- *SGI Altix XE System Rack Installation Instructions*, publication number 007-4902-00x
This guide provides instructions unpacking and installing an SGI Altix XE cluster rack.
- *Guide to Administration, Programming Environments, and Tools Available on SGI Altix XE Systems*, publication number 007-4901-00x
This guide is a reference document for people who manage the operation of SGI computer systems running SGI ProPack 5 for Linux operating system on Altix XE systems. It explains how to perform general system configuration and operations and describes programming environments and tools available for SGI Altix XE systems.
- *Scali Manage On SGI Altix ICE System Quick Reference Guide*
- *Scali Manage User's Guide*

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.

Product Support

SGI provides a comprehensive product support and maintenance program for its products. SGI also offers services to implement and integrate Linux applications in your environment.

- Refer to <http://www.sgi.com/support/>
- If you are in North America, contact the Technical Assistance Center at +1 800 800 4SGI or contact your authorized service provider.
- If you are outside North America, contact the SGI subsidiary or authorized distributor in your country.

Reader Comments

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

You can contact SGI in any of the following ways:

- Send e-mail to the following address: techpubs@sgi.com
- Contact your customer service representative and ask that an incident be filed in the SGI incident tracking system.
- Send mail to the following address:
 - SGI
 - Technical Publications
 - 1140 East Arques Avenue
 - Sunnyvale, CA 94085-4602

SGI values your comments and will respond to them promptly.

SGI Virtu VN200 Overview

This chapter provides an introduction to the SGI Virtu VN200 graphics node (refer to [Figure 1-1](#)) and introduces information necessary to understand the system.




Figure 1-1 SGI Altix 1300 Cluster with Virtu VN200 Graphics Nodes

Product Description

This document is intended as an overview of some of the common operations that system administrators may have to perform to set-up, boot, re-configure (upgrade) or troubleshoot an SGI Altix XE cluster with SGI Virtu VN200 graphics nodes.

The initial release of the SGI Virtu VN200 system (refer to [Figure 1-1](#)) is one or more SGI Altix XE320 servers (compute nodes), one or more SGI Altix XE250 servers (head nodes), and one more SGI Virtu VN200 graphics nodes (visualization/compute nodes). The nodes are networked together using InfiniBand or Gigabit Ethernet switches and can run parallel programs using a message passing tool like the Message Passing Interface (MPI).

An SGI Virtu VN200 system is a distributed memory server as opposed to a shared memory system like that used in the SGI Altix 450 or SGI Altix 4700 high-performance compute servers. Instead of passing pointers into a shared virtual address space, parallel processes in an application pass messages, with each process requiring its own dedicated processor and address space.

Just like a multi-processor shared memory system, a cluster can be shared among multiple applications. For instance, one application may run on 16 processors in the cluster while another application runs on a different set of 8 processors. Very large clusters may run dozens of separate, independent applications at the same time.

Typically, each process of an MPI job runs exclusively on a processor. Multiple processes can share a single processor, through standard Linux context switching, but this can have a significant effect on application performance. A parallel program can only finish when all of its sub-processes have finished. If one process is delayed because it is sharing a processor and memory with another application, then the entire parallel program is delayed. This gets slightly more complicated when systems have multiple processors (and/or multiple cores) that share memory, but the basic rule is that a process is run on a dedicated processor core.

The primary hardware components in an SGI Virtu VN200 system are:

- Head node(s) (SGI Altix XE250 servers)
- Compute nodes (SGI Altix XE320 servers)
- Graphics nodes (SGI Virtu VN200 graphics nodes)
- Network interconnect components (Gigabit Ethernet switches, InfiniBand switches, PCI cards, and cables)
- System console, monitor, keyboard, and mouse

The head node is connected to the interconnect network and also to the public network, typically via the local area network (LAN). The head node is the point of submittal for all MPI application jobs for the cluster. An MPI job is started from the head node and the sub-processes are distributed to the cluster compute nodes from the head node. The main process on the head node waits for the sub-processes to finish. For large clusters or clusters that run many MPI jobs, multiple head nodes may be used to distribute the load.

The compute nodes are identical computing systems that run the primary processes of MPI applications. These compute nodes are connected to each other through the interconnect network.

A graphics node is similar to a compute node in that it contains processors and memory, but it has an additional high-performance 3D graphics card installed.

The network interconnect components are typically Gigabit Ethernet or InfiniBand. The MPI messages are passed across this network between the processes. This compute node network does not connect directly to the public network because mixing external and internal cluster network traffic could impact application performance. Visualization nodes may be connected to the public network to act as a login or application gateway for remote visualization.

Note: Refer to [“Related Publications” on page xiv](#), for a listing of relevant SGI Technical Publications that can provide more detailed information about SGI cluster head nodes, compute nodes, and system rack.

Warnings and Cautions

Read the following safety and warnings and cautions when operating your SGI Virtu VN200 cluster.



Warning: System power on/off: The power button DOES NOT turn off the system AC power. To remove power from system, you must unplug the AC power cord from the wall outlet. Make sure the AC power cord is unplugged before you open the chassis, add, or remove any components. Failure to do so may result in serious injury or death.



Caution: Electrostatic discharge (ESD) and ESD protection: ESD can damage disk drives, boards, and other parts. We recommend that you perform all replacement procedures on an ESD-safe worksurface. If one is not available, provide some ESD protection by wearing an antistatic wrist strap attached to the chassis ground or any unpainted metal surface on your server when handling parts.



Caution: Always handle boards carefully. They can be extremely sensitive to ESD. Hold boards only by their edges. After removing a board from its protective wrapper or from the server, place the board component side up on a grounded, static free surface. Use a conductive foam pad if available but not the board wrapper. Do not slide board over any surface.



Caution: Do not place the node board on a conductive surface. A conductive surface can damage the BIOS battery and prevent the system from booting. Always place the node board on an ESD-safe worksurface.

Site Planning

Ensure that all site requirements are met before you install and boot your system. If you have questions about the site requirements, contact a site planning representative by e-mail (site@sgi.com).

Selecting a Delivery Route

To ensure that the system can be delivered to the planned location, answer the following questions before you plan a delivery route for the new system:

- Will the shipping crate fit through doorways and hallways and on elevators?

In addition to measuring the width of the hallways along the planned delivery route, measure corners where the system might get stuck, the width and height of doorways and elevators, and other areas that may cause problems.

- **SGI recommends that you use a pallet jack to position the rack as close as possible to its final destination.** If the shipping crate cannot be transported to the final destination, can you unpack the system somewhere else?

Often it is possible to unpack the system in a hallway or on a loading dock, and then roll the system to its final destination.

- Is the floor strong enough to support the weight of the system?
- Is the elevator capable of lifting the system?

If the intended delivery route includes an elevator, check its weight capacity and size against the system specifications. The use of freight blankets can reduce damage to the elevator or the system.

- Are there any steep angles, bumps, changes in level, or thick carpeting along the delivery route?
- Did you ensure that the leveling pads are fully retracted?



Caution: The rack systems have screw-in leveling pads. If you move the rack with these feet extended, severe damage to the chassis can occur. These feet sometimes unscrew during shipment. Before you unpack or move a system, ensure that the leveling pads are fully retracted.

Selecting a Final Location

Consider the following issues when you select a final location for the system:

- Will the system fit in its intended location?
Carefully calculate the total system dimensions to ensure that it will fit in its intended final location.
- Does the intended system location provide adequate access space for maintenance?
- Even if the system will fit in its intended final location, you must have room to maintain it. Ensure that you have enough room to open the doors, remove boards, and accomplish other routine tasks. As a general rule, SGI systems require 48-in of clearance from the front and rear of the rack. Side access to the rack is not required for service.
- Is the intended location subject to flooding, extremes of humidity or temperature, or any other factor that would make it inappropriate for sensitive electronic equipment?

The air temperature should not be too high and should not fluctuate dramatically, air should circulate freely and be relatively dust-free, and the system should not be exposed to any caustic or corrosive chemicals or vapors.

- Will the system interfere with normal traffic through aisles, hallways, or entrance ways in the intended location?
- Will the intended location enable you to conveniently perform routine operations, such as loading and unloading media, attaching cables, and so on?
- Is the floor of the intended final location strong enough to support the weight of the system and any future expansions? Large systems should be installed in computer rooms with raised floors. Pay particular attention to floor loading and weight distribution in this case. Floor-loading specifications are typically calculated by averaging the total chassis weight that is distributed over the entire footprint of the chassis. Because the chassis sits on four casters or four leveling pads, the load at each of these four points is greater.
- Does the intended location of the system allow for future expansion plans?

Unpacking the Rack

Ensure that the temperature of the rack is acclimated to the installation environment and that the system crate is stable, and in an upright position.

Refer to [Figure 1-2](#) as you follow these steps; the numbered illustrations correspond to the numbered steps. Ensure that the temperature of the rack is acclimated to the installation environment and that the system container is stable, and in an upright position.



Caution: Use eye protection when cutting the plastic bands on the shipping container. Cutting the bands under high tension may cause them to snap back and strike your face. Use caution when cutting the metal bands.




Caution: Use caution cutting the horizontal band holding the wooden ramp to the shipping container, if the ramp is not supported, it may fall on you and cause serious personal injury.

1. Use a sharp cutting tool to remove the band that secures the ramp to the shipping container.
2. Position the ramp as shown in [Figure 1-2](#). Holes in the ramp fit over the pegs in the base.
3. Remove the vertical bands that secure the load to the wooden base and remove top corrugated cap, corrugated sleeves, foam top cushion and poly bag from the rack.
4. Remove the four M12 bolts securing the rack to the wooden base:
 - a. Open the front door and locate the two M12 bolts that secure the rack to the large metal bracket.
 - b. Open the rear door, and disconnect the rear door ground wire.
 - c. Remove the rear door by lifting it off the two hinge pins.
 - d. Use a 3/4-in. socket and ratchet, remove the two M12 bolts securing the back of the rack to the wooden base.
5. Remove four 9/16-in. hex head cap screws that secure the large metal bracket to the front of the wooden base and remove the metal bracket.



Warning: Serious personal injury or death could result if the rack becomes unbalanced or gains too much momentum when rolled. Always get the assistance of other qualified personnel when rolling a rack.

6. Get four qualified persons to carefully roll the rack down the ramp and onto the floor.
7. Install the rear door onto rack by aligning the door on the hinge pins pushing downward.
8. Connect the door ground wire.



Unpacking the Graphics Nodes

The SGI Virtu VN200 graphics nodes are packed as shown in [Figure 1-3](#). Unpack each graphics node and place it on a stable ESD safe worksurface or position each graphics node on the rack shelf as shown in [Figure 1-4](#).




Figure 1-3 Unpacking the Graphics Nodes

Installing the Rack Shelf

Install the rack shelf as shown in [Figure 1-4](#). The graphics nodes require 4U of rack space. Five graphics nodes can be placed on the rack shelf. Position each graphics node on the rack shelf as shown in [Figure 1-4](#).




Figure 1-4 Rack Shelf

System Configurations and Cabling

Table 1-1 lists the system configurations for the SGI Virtu VN200. SGI Altix ICE clusters and SGI Altix IA64 configurations will be supported in a future release of the product.

Table 1-1 SGI Virtu VN200 System Configurations

system Configuration	Description
SGI Altix XE cluster with VN200 graphics nodes	Integrated compute/storage/visualization system; racked, provisioned, and managed like a compute node; SGI manufacturing provides software installation template.
SGI Altix ICE cluster with VN200 graphics nodes ^a	Integrated compute/storage/visualization system; racked, provisioned, and managed like a service node; SGI Manufacturing provides software installation template
SGI Altix 4700 cluster with VN200 graphics nodes ^a	Integrated compute/storage/visualization for hybrid compute environments with IA64 SGI Altix systems; SGI Manufacturing provides software installation template

a. Not supported for the initial product release.

Figure 1-5 shows an SGI Virtu VN200 graphics node configured with an SGI Altix XE cluster.

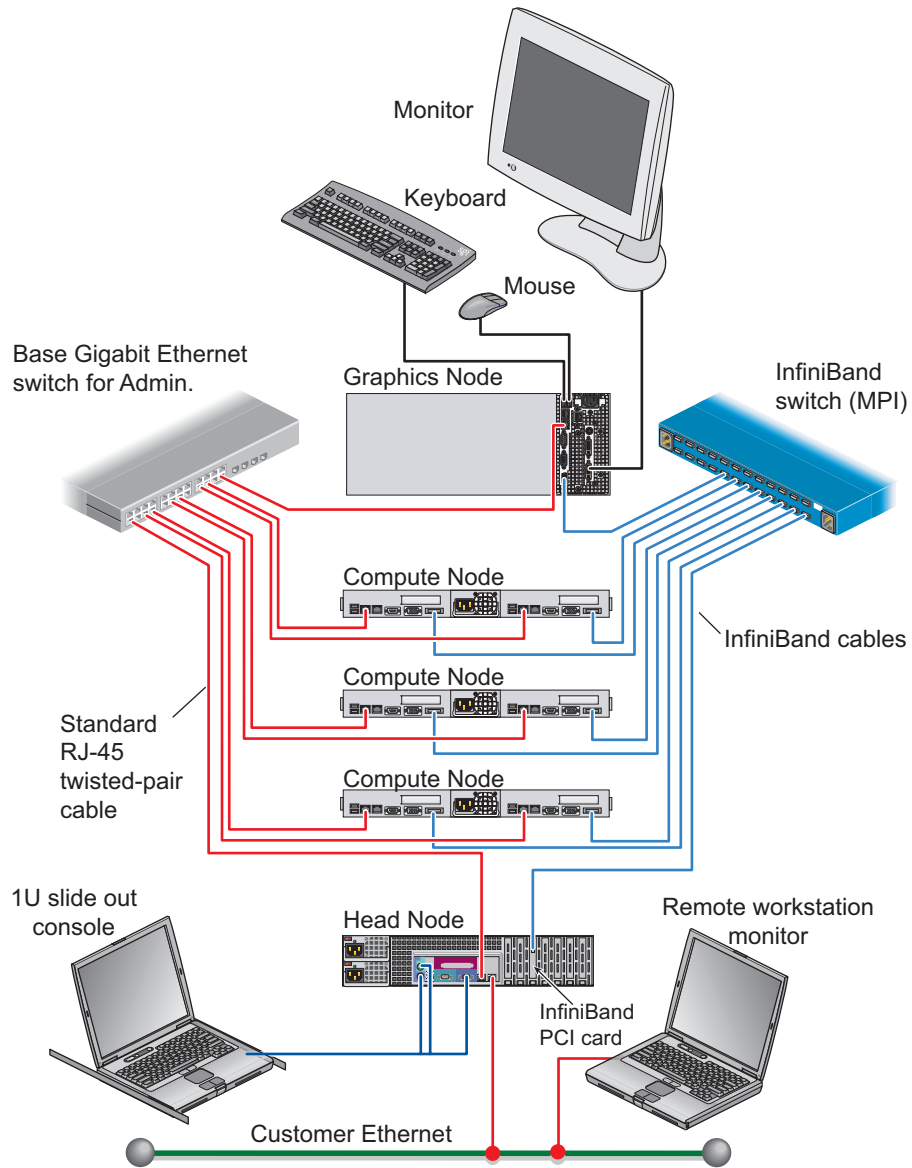


Figure 1-5 Altix XE Cluster Configuration

Figure 1-6 shows an SGI Virtu VN200 graphics node configured with an SGI Altix ICE cluster. This system configuration will be supported in a future product release.

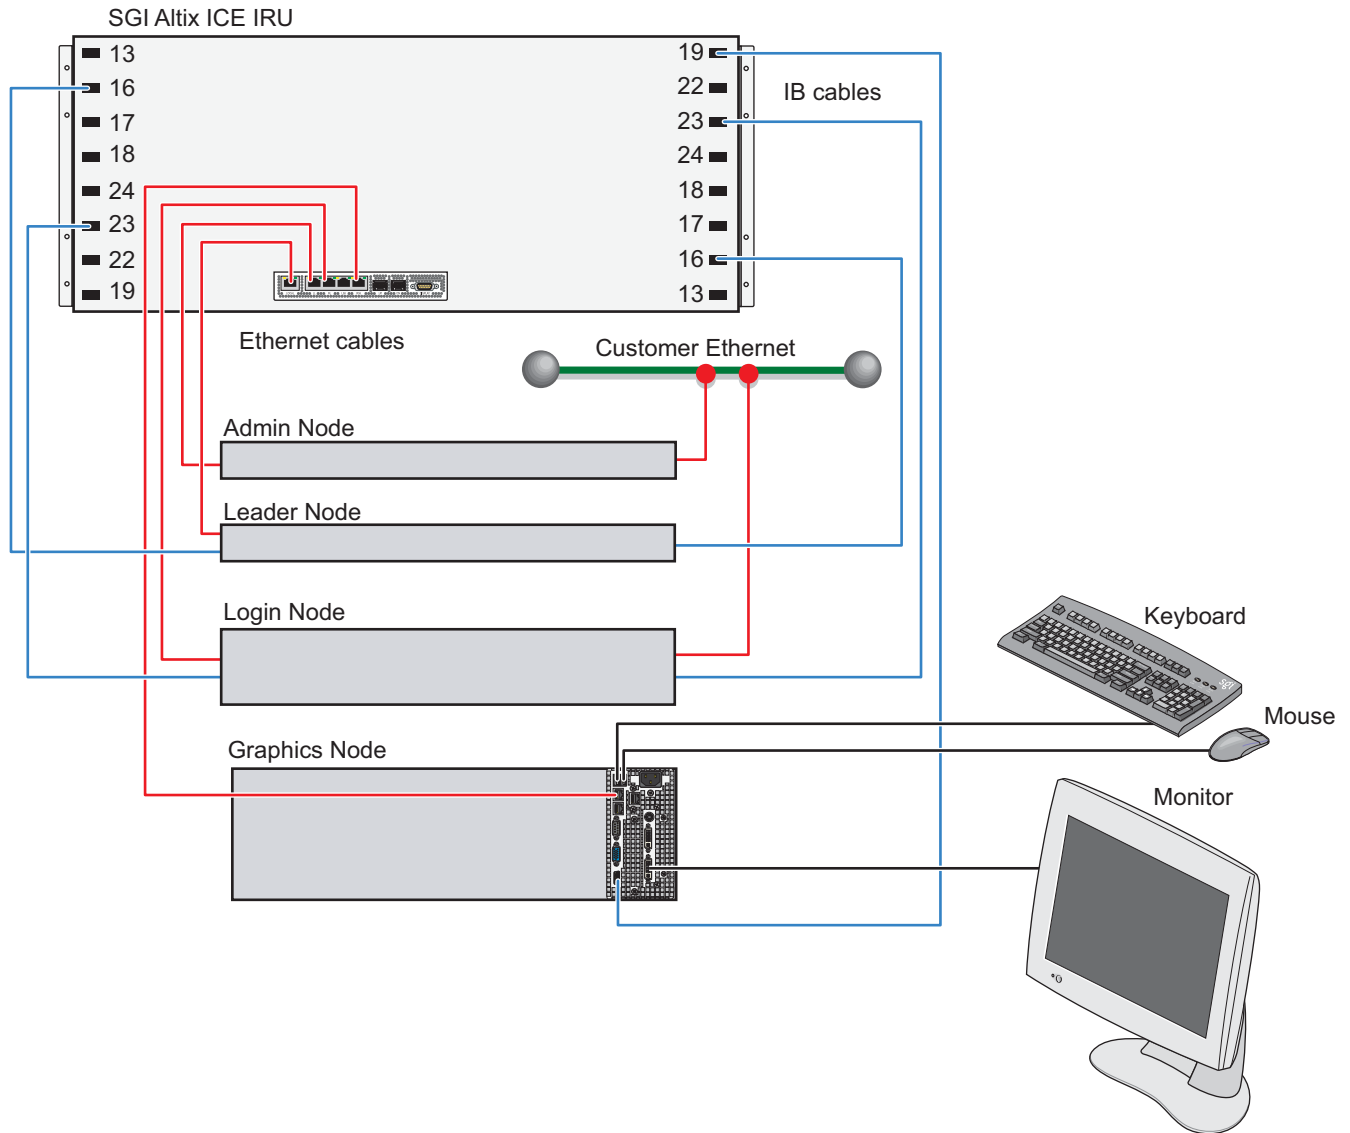


Figure 1-6 SGI Altix ICE Cluster Configuration

Graphics Node Dimensions

The graphics node physical characteristics are shown in [Figure 1-7](#). The graphics node dimensions are as follows:

- Height: 6.72 in. (17.07 cm)
- Width: 3.44 in. (8.74 cm)
- Depth: 32.44 in. (82.39 cm)




Figure 1-7 Graphics Node Dimensions

Electrical Requirements

The SGI Virtu VN200 accepts single-phase input voltage only. It typically connects to a single-phase or 3-phase rack power distribution unit (PDU) that outputs single-phase 220 VAC. Ensure that the required voltage is available and is within a reasonable distance of the intended location. If it is not, the site must be wired for the required voltage.

Power Consumption for the Graphics Node

Note: Refer to “[Related Publications](#)” on page xiv, for a listing of relevant SGI Technical Publications that can provide the electrical requirements for the other node types in the system rack.

The system configuration used to measure the power consumption of the graphics node listed in [Table 1-2](#):

- CPU: 2 Intel 5482 (150W)
- Memory: 16 GB ECC DDR2 800 MHz FBDIMMs using 2 GB DIMMs
- HDD: 1 SATA, 80GB
- OS: Windows XP 64-bit
- Stress software: CPU Burn
- Fans: Two 60 x 38mm 12,000 RPM
- Video: NVIDIA Quadro FX5600

[Table 1-2](#) lists the power consumption for each graphics node.

Table 1-2 Graphics Node Power Consumption

	Power (W)	Current (A)	Voltage (V)	Heat Dissipation (BTU)
Idle	302	2.52	120.5	1,030
Load	468	3.88	120.5	1,597

Environmental Requirements

Electromagnetic interference (EMI), electrostatic discharge (ESD), vibration, and humidity can cause problems for computer systems.

Operating Temperature	10° to 40° C (50° to 104° F)
Non-operating Temperature	-40° to 70° C (-40° to 158° F)
Operating Relative Humidity	8% to 90% (non-condensing)
Non-operating Relative Humidity	5% to 95% (non-condensing)

Acoustics

The acoustic sound level of the graphics node is 62.5 dBA.

Electromagnetic Interference

Electromagnetic interference (EMI) is caused by malfunctioning, incorrectly manufactured, or incorrectly installed devices that radiate electrical signals. Common sources of EMI include electronic, telephone, and communications equipment. EMI transmissions can be conducted or emitted.

Use properly shielded connectors and cables throughout the site.



Caution: Failure to use shielded cables where appropriate may violate FCC regulations and void the manufacturer's warranty.

Electrostatic Discharge

SGI designs and tests its products to ensure that they resist the effects of electrostatic discharge (ESD). However, it is still possible for ESD to cause problems that range from data errors and lockups to permanent component damage. To protect the systems from ESD, follow these precautions:

- Minimize the use of carpeting at computer locations (or consider special static-reducing carpet).
- Ensure that all electronic devices are properly grounded.
- Keep chassis doors and access panels closed while the system is operating.
- Fasten all screws, thumbnail-fasteners, and slide locks securely.
- Use a grounded static wrist strap whenever you work with the chassis or components.
- Use antistatic packing material for storage and transportation.
- Clear the site of all devices that create static electricity or provide possible sources of EMI.

Regulatory Specifications

The SGI Virtu VN200 graphics node conforms to several national and international specifications and European Directives listed on the “Manufacturer’s Declaration of Conformity.” The CE insignia displayed on each device is an indication of conformity to the European requirements.



Caution: Each SGI server 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.

Server Model Number

The CMN (model) number for each server is printed on the system label on the unit.

CE Notice and Manufacturer's Declaration of Conformity

The “CE” symbol indicates compliance of the device to directives of the European Community. A “Declaration of Conformity” in accordance with the standards has been made and is available from SGI upon request.

Electromagnetic Emissions

This section provides the contents of electromagnetic emissions notices for various countries.

FCC Notice (USA Only)

This equipment complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions:

- This device may not cause harmful interference.
- This device must accept any interference received, including interference that may cause undesired operation.

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 technician.



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 Class A 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。

Shielded Cables

The SGI Virtu VN200 product is FCC compliant under test conditions that include the use of shielded cables between the server and its peripherals. Your server and any peripherals that you purchase from SGI have shielded cables. Shielded cables reduce the possibility of interference with radio, television, and other devices. If you use any cables that are not from SGI, ensure that they are shielded. Telephone cables do not require shielding.

Optional monitor cables supplied with your server system use additional filtering molded into the cable jacket to reduce radio frequency interference. Always use the cable that is supplied with your system. If your monitor cable becomes damaged, obtain a replacement cable from SGI.

System Features

An SGI Virtu VN200 (refer to [Figure 2-1](#)) is an integrated compute environment that can scale to 288 (refer to [Figure 2-2](#)).

The standard cooling method for the SGI Virtu VN200 is blower-driven air cooling. Optional supplemental water cooling will be available (in a 42U water-cooled rack) in a future release.




Figure 2-1 SGI Virtu VN200 Graphics Nodes

Server Racks

Each SGI Virtu VN200 rack contains the head nodes, compute nodes, an InfiniBand switch, a Gigabit Ethernet switch, a system console, power distribution units (PDUs) and graphics nodes (refer to [Figure 2-2](#)). The racks are typically air-cooled.




Figure 2-2 SGI Virtu VN200 - Front View

Virtu VN200 Graphics Node

Each Virtu VN200 graphics node (refer to [Figure 2-3](#)) contains the hardware components listed below:

- 1 Supermicro® X7DWT-INF node board that uses the Intel 5400 Seaburg chip set (same node board used in SGI Altix XE320 servers)
- 2 Intel® Xeon® quad-core 5400 sequence processors with 1333- or 1600-MT/s FSB
- 2 to 8 800-MT/s FB-DIMMs
- 2 Gigabit Ethernet ports (GLAN 1 and GLAN2)
- 1 BMC LAN port available from GLAN 1
- 1 4x DDR InfiniBand interface
- 1 PCIe x16 Gen2 graphics board
- 1 AOC-SIMSO IPMI interface board
- 1 or 2 internal 80-, 120-, or 160-GB SATA JBOD disk drives
- Redundant 1000 watt power supplies




Figure 2-3 Virtu VN200 Graphics Node

Front Panel Components and Status LEDs

The front panel of the SGI Virtu VN200 provides status LEDs, system disk access, and power supply access (refer to [Figure 2-4](#)). The front panel status LED definitions are listed below:

- PWR: illuminates white on unit power up
- HDD: illuminates blue on disk drive access
- ALERT: illuminates red on fan or temperature fault
- NIC1 and NIC2: illuminates green on GLAN1 or GLAN2 access
- PWR FAIL: illuminates red on power supply failure

Power Supply Status LED:

- Illuminates yellow when AC input voltage is present and unit is in standby mode
- Illuminates green when power button has been pressed and DC output voltage is present




Figure 2-4 Front Panel Components

Rear Panel Components

Figure 2-5 shows the location of the rear panel components on the SGI Virtu VN200 graphics node. The rear panel component descriptions are listed below:

- Input AC - Connects to system input AC power cord
- USB ports - Universal Serial Bus (USB 2.0) ports
- GLAN1/2 - Gigabit Ethernet ports
- BMC LAN - (not used)
- Serial port - (Disabled)
- VGA port - (Disabled)
- InfiniBand port - 4x DDR InfiniBand interface
- Stereo emitter - Stereo sync connector available on NVIDIA FX3700, 4600 and 5600
- Secondary DVI-I - Secondary display port
- Primary DVI-I - Primary display port




Figure 2-5 Rear Panel Components

Node Board

At the heart of the SGI Virtu VN200 lies a dual-processor socket node board, which is based on the Intel 5400P chipset (refer to [Figure 2-6](#) and [Figure 2-7](#)).




Figure 2-6 VN200 Node Board

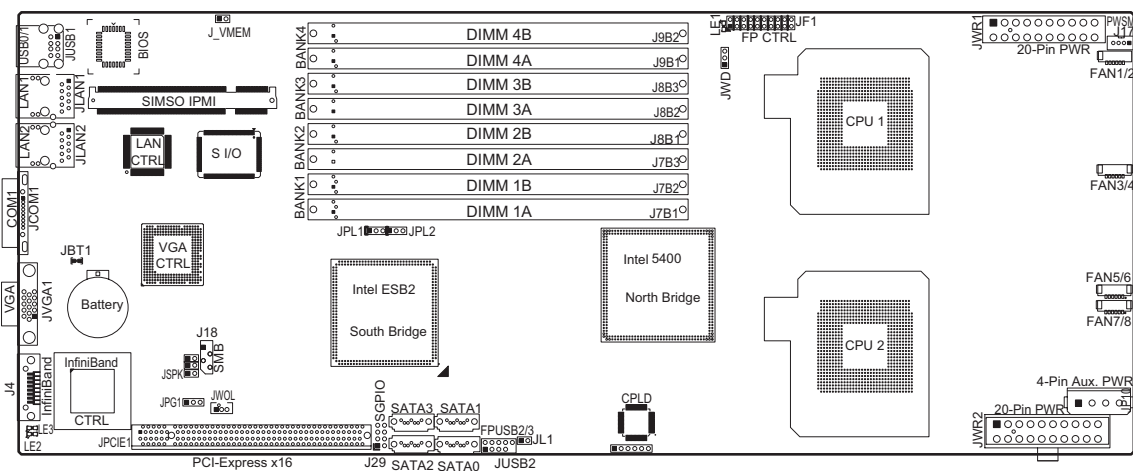


Figure 2-7 Node Board Component Locations

Node Board Features

Refer to the node board block diagram, [Figure 2-8](#).

CPU

- Dual Intel® 64-bit Xeon LGA 771 quad-core/dual-core processors at a front side bus speed of 1.600 GHz or 1.333 GHz

Memory

- Eight 240-pin DIMM sockets with support up to 32 GB DDR2 FBD (fully buffered) ECC 800 MT/s memory

Chipset

- Intel 5400 chipset, including: the 5400 Memory Control Hub (MCH), and the Enterprise South Bridge 2 (ESB2)

Expansion Slots

- One PCI-Express (Gen. 2.0) slot (one PCI-E slot x16: JPCIE1)

BIOS

- 8 MB Phoenix® Flash ROM
- DMI 2.3, PCI 2.2, ACPI 1.0/2.0, Plug and Play (PnP), USB Keyboard Support, SMBIOS 2.3

PC Health Monitoring

- Onboard voltage monitors for CPU cores, chipset voltage (+1.5V), memory voltage, +3.3V, +5V, +12V, -12V, +3.3V Standby, +5V standby, HT and VBAT
- Fan status monitor with firmware control
- CPU/chassis environment monitoring
- Platform Environment Control Interface (PECI) ready
- CPU fan auto-off in sleep mode
- CPU slow-down on temperature overheat
- CPU thermal trip support for processor protection, power LED
- Power-up mode control for recovery from AC power loss
- Auto-switching voltage regulator for CPU cores
- System overheat/Fan Fail LED Indicator and control
- Chassis intrusion detection
- System resource alert via Super Doctor III
- I2C temperature sensing logic
- Thermal Monitor 2 (TM) support

ACPI Features

- Slow blinking LED for suspend state indicator
- Main switch override mechanism
- ACPI Power Management

Onboard I/O

- Four SATA2 system disk ports (only two ports are used)
- One SIMSO IPMI slot
- One Intel Gigabit Ethernet 82563EB controller supports two Gigabit LAN ports
- One COM port
- Up to four Universal Serial Bus 2.0 ports (2 ports, 2 headers, only two are used)
- ATI ES1000 Graphic Controller with 32 MB video memory (not used)
- Super I/O: Winbond W83627HF
- Mellanox MT25204A0-FCC-D InfiniBand controller

Other

- Wake-on-LAN
- Chassis Intrusion
- Console redirection
- Onboard fan speed control by thermal management via BIOS

CD/Diskette Utilities

- BIOS flash upgrade utility and device drivers

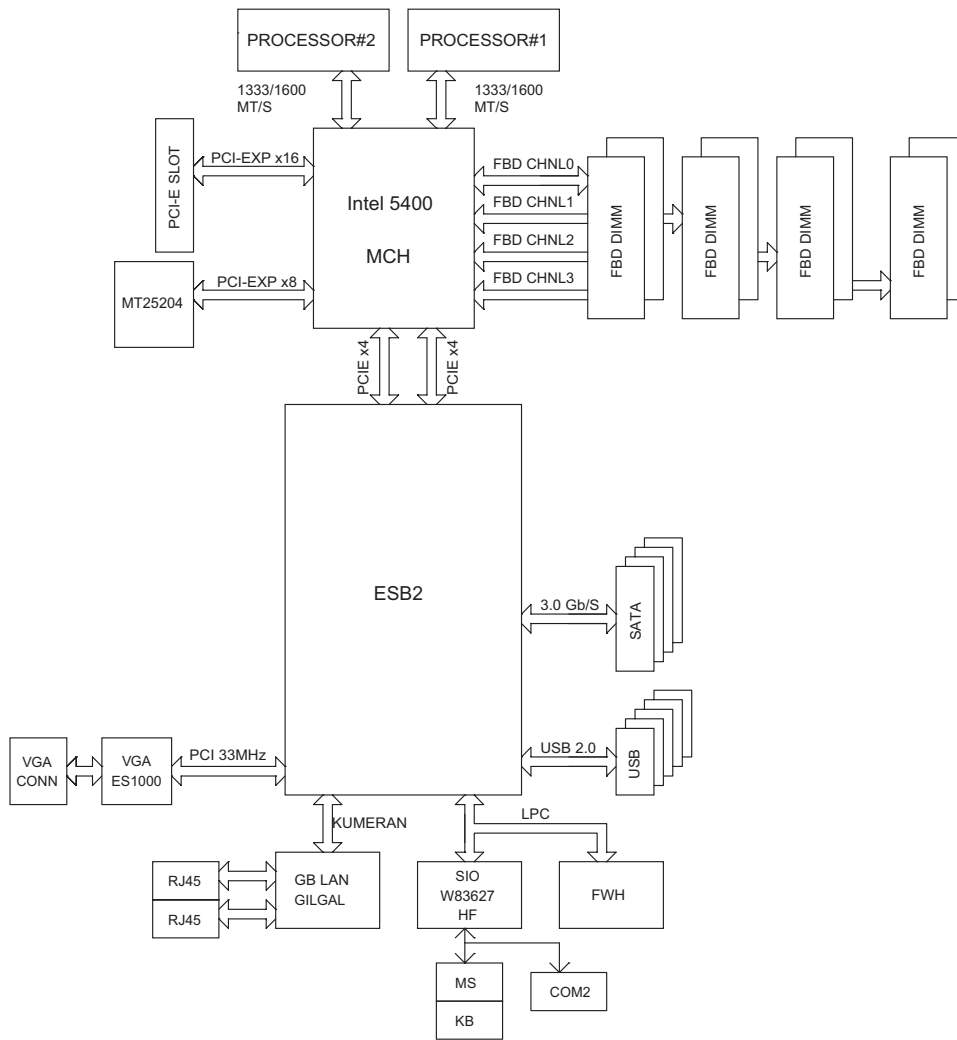


Figure 2-8 Node Board Block Diagram

Jumper Quick-reference

Refer to [Figure 2-7](#) on [page 26](#) when using [Table 2-1](#) and [Table 2-2](#).

Table 2-1 Node Board Jumper Definitions

Jumper	Description	Default Setting
JBT1	CMOS Clear	Open
JI2C1/JI2C2	I2C to PCIe Slot	Open (disabled)
JPG1	VGA Enable Pins 1-2	Pins 1-2 (enabled)
JPL1 GLAN1	Enable Open	Open (Enabled)
JPL2 GLAN2	Enable Open	Open (Enabled)
JWD	Watch Dog	Open (Disabled)

Table 2-2 Node Board Connector Definitions

Connector	Description
FAN 1-8	Fans 1-8 (CPU fans/chassis fans)
J4	InfiniBand connector
J17	Power system management bus (I2C) header
J18	System Management Bus Header
J29	Serial general purpose I/O header
COM1	COM1 (JCOM1) Serial Port
JF1	Front control panel connector
JL1	Chassis intrusion header
LAN1/LAN2	Gigabit Ethernet ports
JP10	4-pin Aux. PWR (for HDDs)
JWR1	Primary 20-Pin PWR connector
JWR2	External 20-Pin PWR connector
JSPK	Speaker header

Table 2-2 Node Board Connector Definitions (**continued**)

Connector	Description
JWOL	Wake-on LAN header
USB	0/1 back panel USB 0/1 for external use (JUSB1)
USB	2/3 front panel USB 2/3
VGA	VGA Connector (JVGA1)
SIMSO	SIMSO IPMI slot
SATA	0-3 Intel SATA 0-3 connectors
LED	Indicator description
LE1	Onboard Power LED Indicator
LE2	InfiniBand Link LED (green)
LE3	InfiniBand Activity LED (yellow)

Memory Overview

Each node board supports two to eight 240-pin DIMMs that can provide up to 32 GB of ECC fully buffered DIMM (FBDIMM) DDR2-800 SDRAM (refer to [Figure 2-9](#) and [Figure 2-10](#)). Memory can be installed in both interleaved (dual-channel) and non-interleaved (single-channel) configurations. All memory modules used to populate the system should be the same size, type, and speed.

The node board memory controller supports memory mirroring to prevent data loss in case of DIMM failure.

Note: The memory mirroring feature is not enabled by default.

The system maintains two copies of all the data in memory; therefore, a mirrored system uses only half of the installed DIMMs. If a DIMM fails, the system recovers the second copy of data from the mirrored DIMM in the opposite channel. If both the primary and the mirrored copy of the data become corrupt, the system memory will fail.

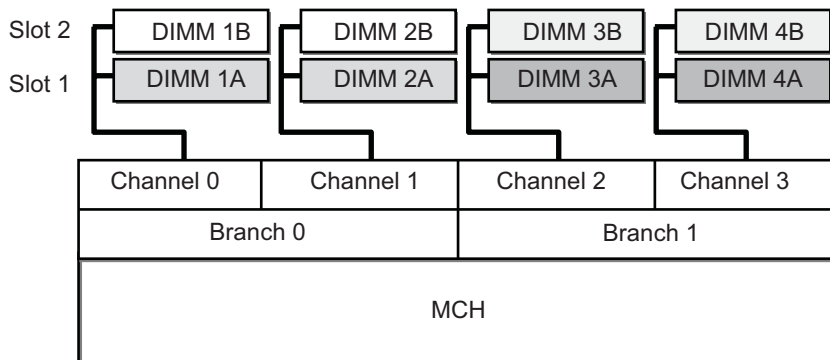


Figure 2-9 Memory Block Diagram




Figure 2-10 Node Board DIMM Layout

Memory DIMM Pairs

[Table 2-3](#) lists the memory DIMM pairs for the SGI Virtu VN200.

Table 2-3 Memory DIMM Pairs

DIMM Pair	DIMM Slot
Pair 1	1A and 2A
Pair 2	3A and 4A
Pair 3	1B and 2B
Pair 4	3B and 4B

Memory Configuration Rules

- You can install a minimum of 2 and a maximum of 8 FBDIMMs on the node board.
- You must install FB-DIMMs in pairs (refer to [Table 2-3](#)).
- Both FB-DIMMs within a pair must have the same capacity and speed.
- 800-MT/s FB-DIMMs require 1600-MT/s processors to run at full speed.
- To configure memory for maximum performance, balance the DIMM capacity across both memory branches; for example, if you install 2GB of memory in Branch 0, you must also configure 2GB of memory in Branch 1 (refer to [Figure 2-9](#)).
- Unbalanced branch configurations can reduce overall memory performance up to 40%.

Memory Part Numbers

[Table 2-4](#) lists the SGI Virtu VN200 memory sizes and their SGI part numbers.

Table 2-4 Memory Sizes and SGI Part Numbers

Memory Size and Type Speed	SGI Part Number
1-GB FB-DIMM 800 MT/s	9010132
2-GB FB-DIMM 800 MT/s	9010133
4-GB FB-DIMM 800 MT/s	9010134

Table 2-5 lists the DIMM configurations for the SGI Virtu VN200.

Table 2-5 DIMM Configurations

DIMM Pair 1 Capacity	DIMM Pair 2 Capacity	DIMM Pair 3 Capacity	DIMM Pair 4 Capacity
2GB ^a			
2GB	2GB		
2GB	2GB	2GB ^a	
2GB	2GB	4GB	
2GB	2GB	2GB	2GB
4GB ^a			
4GB	4GB		
4GB	4GB	4GB ^a	
4GB	4GB	8GB	
4GB	4GB	4GB	4GB
8GB ^a			
8GB	8GB		
8GB	8GB	8GB ^a	
8GB	8GB	8GB	8GB

a. Unbalanced configuration.

AOC-SIMSO (BMC) Card

This section describes the AOC-SIMSO baseboard management controller (BMC). The intelligent platform management interface (IPMI) management utility defines the protocols used for interfacing with the BMC on the node board. Use the GLAN 1 port to connect to the BMC interface (refer to [Figure 2-11](#)).

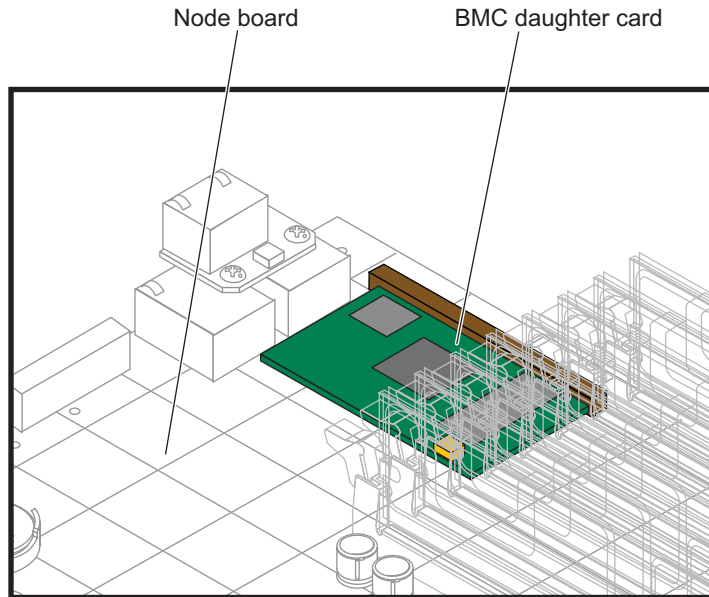


Figure 2-11 AOC-SIMSO (BMC) Daughter Card

Note: Refer to the *SGI Altix XE320 System User's Guide*, publication number 007-5466-00x for more detailed information about the baseboard management controller (BMC) intelligent platform management interface (IPMI) management utility.

Note: The add-on graphics card disables the KVM-over-LAN feature of the BMC.

IPMI controller enables you to communicate with the BMC via a Java-based web interface. The BMC has the following features:

- Installs in the JIPMI slot of the serverboard
- Enables you to:
 - Monitor the Server Event Log (SEL)
 - Monitor system temperature, fan speed, and voltage
 - Power up and power down the system remotely
 - Update BMC firmware images

The BMC monitors onboard instrumentation such as temperature sensors, power status, voltages and fan speed, and provides remote power control capabilities to reboot and/or reset the server. It also includes remote access to the BIOS configuration and operating system console information via serial-over-LAN (SOL). Because the controller is a separate processor, the monitoring and control functions work regardless of CPU operation or system power-on status.

Power Supplies

Each hot-pluggable power supply connects to the power distribution assembly in the center of the chassis. The power supplies are located in the chassis as shown in [Figure 2-12](#). The power supply specifications are listed in [Table 2-6](#).

Table 2-6 Power Supply Specifications

Specification	Description
Power output	1000 Watts
Input power	100-240VAC, 50-60Hz,
Current output	15A
Voltage output	+12 VDC, +5 VDC
Status LED	Illuminates yellow when AC input voltage is present and unit is in standby mode. Illuminates green when power button has been pressed and DC output voltage is present.




Figure 2-12 Redundant Power Supplies

Cooling Fans

Each power supply has built in cooling fans that pull cool air in from the front of the chassis. Two chassis cooling fans pull cool air in from the front of the chassis and blow it across the node board heatsinks and memory DIMMs. The chassis cooling fans receive power from the node board assembly (refer to [Figure 2-6](#)).

The graphics card assembly has its own cooling fan and heatsink assembly. The warm air is discharged out of the rear panel of the chassis (refer to [Figure 2-13](#)).




Figure 2-13 Cooling Fans

System Disk Drives

The system disk drive are controlled by the SATA ports on the node board (refer to [Figure 2-7](#) on [page 26](#). The following system disk drive options are supported by the SGI Virtu VN200 (refer to [Figure 2-14](#)):

- 2.5-in, 80 GB, 7200 RPM SATA
- 2.5-in, 120 GB, 7200 RPM SATA
- 2.5-in, 160 GB, 7200 RPM SATA

Note: Shut down the operating system before replacing a disk drive. Disk drives can be replaced while the system is powered on (hot-plugged).

- To remove a disk drive, release the latch on the drive carrier assembly and slide the drive out of the chassis.
- To install a disk drive, slide the drive into the chassis and close the latch.




Figure 2-14 System Disk Drives

Graphics Cards

The following graphics cards are supported on the SGI Virtu VN200. Other cards may be supported in a future release.

- NVIDIA Quadro FX1700
- NVIDIA Quadro FX3700
- NVIDIA Quadro FX4600
- NVIDIA Quadro FX5600




Figure 2-15 Graphics Card

Table 2-7 lists the SGI Virtu VN200 graphics card specifications.

Table 2-7 Graphics Card Specifications

Type	Memory	Memory Bandwidth	Performance	Display Connectors
NVIDIA Quadro FX1700	512 MB DDR2	12.8 GB/s	191 MTriangles/s	2 DVI-I HD-OUT
NVIDIA Quadro FX3700	512MB GDDR3	51.2 GB/s	250 MTriangles/s	2 DVI-I STEREO
NVIDIA Quadro FX4600	768 MB GDDR3	57.6 GB/s	250 MTriangles/s	2 DVI-I STEREO
NVIDIA Quadro FX5600	1.5GB GDDR3	76.8 GB/s	300 MTriangles/s	2 DVI-I STEREO

Accelerated Driver Installation

Accelerated graphics drivers from NVIDIA are preinstalled by SGI on the SGI Virtu VN200 graphics node. These drivers are not part of the standard SLES or RHEL distribution of Linux. A copy of these drivers is available on the SGI Propack CD and/or supportfolio. SGI distributes these drivers as `.rpm` files.

SGI provides the limited support of the factory installed NVIDIA graphics drivers. SGI has qualified these drivers and they are compatible with the hardware and software configuration of the SGI Virtu VN200. Our support does not extend to bug fixes of advanced features provided by the NVIDIA graphics driver. More recent drivers that have not been qualified by SGI may be available from the NVIDIA website. SGI provides no support for these drivers, however, support may be available directly from NVIDIA.

System Procedures

Connecting a Digital Monitor

The NVIDIA Quadro FX card supports two DVI-I outputs. The DVI-I output farthest away from the stereo connector is the primary channel for the system. By default this port outputs an analog signal when the system is not configured. If you attach digital monitor to the system, you can reset the DVI-I port by performing the following steps.

1. Attach the monitor.
2. Log into the node remotely as `root`.
3. Enter the following commands:

```
init 3
rm /etc/X11/xorg.conf
nvidia-xconfig
init 5
```

The `xorg.conf` file is initialized with the Extended Display Identification Data (EDID) setting from the digital monitor.

Note: The graphics node VGA port (refer to [Figure 2-5](#) on [page 25](#)) is disabled when a NVIDIA Quadro FX card is installed.

Connecting a Keyboard and Mouse

A keyboard/mouse is not supplied as standard for the system but may be provided as an option from SGI. Connect a USB keyboard and a USB mouse to your system, as shown in [Figure 3-1](#). If your keyboard/mouse is further than 10 feet (3 meters) away from the USB ports on the system, then optional USB extension cables are required.

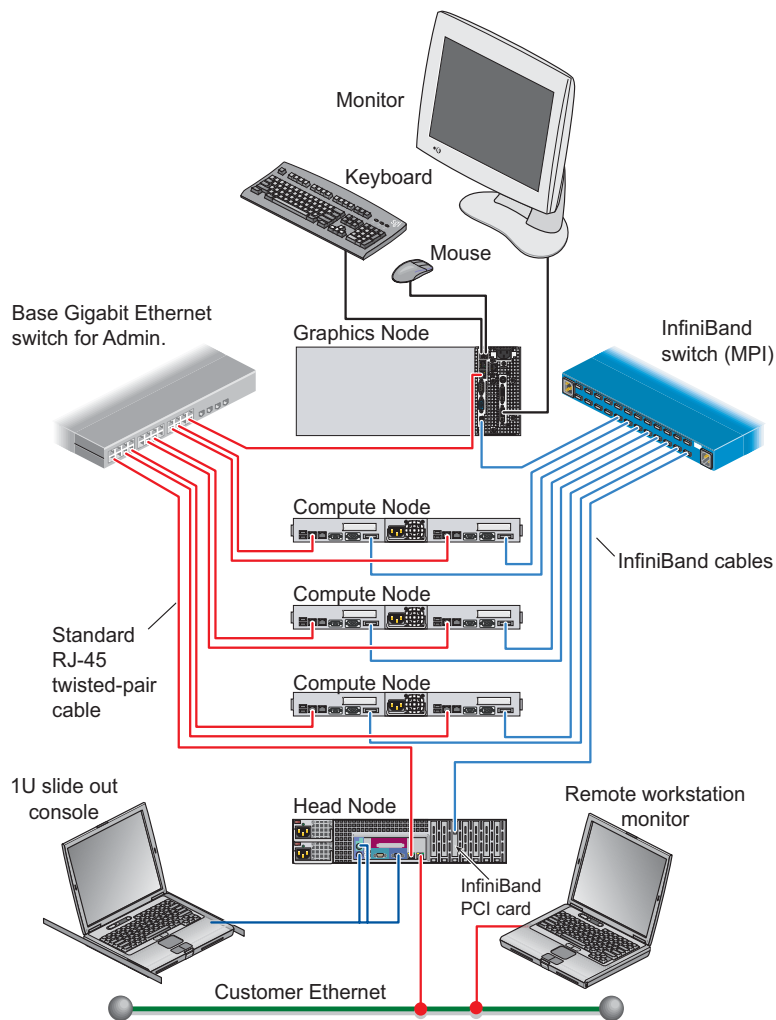


Figure 3-1 SGI Altix XE Cluster Configuration

Configuring Firmware and Drivers

The SGI Virtu VN200 uses the same BIOS images as the SGI Altix XE320 server, with customized settings.

The SGI Virtu VN200 also uses factory-installed, accelerated NVIDIA graphics drivers and X.org windowing system; SGI provides limited support for these SGI qualified NVIDIA graphics drivers. You can obtain the accelerated NVIDIA graphics drivers from the following sources:

- Download the updated firmware images and graphics card drivers from SGI Supportfolio
- Install the drivers from the SGI ProPack CD
- Install the BIOS from a USB flash drive, then install the drivers from the operating system using the appropriate `.rpm` file

Power On/Off the Cluster

You can use the baseboard management controller (BMC) interface to perform power management and other administrative functions. Refer to the *SGI Altix XE320 User's Guide*, publication number 007-5466-00x for information about the BMC.

You can use the Scali parallel shells tools suite to perform remote power management from the head node. Login to the head node as root, and use the `/opt/scali/sbin/power` command to manage the system.

```
# power -H
```

Usage:

```
/opt/scali/sbin/power [option] [nodelist <on|off|cycle|status>]
```

Example:

Use the following command to power cycle cluster nodes 001 through 032:

```
power c11n[001-032] -cycle
```

If your cluster uses the Scali Manage administrative software (release 5.6.x) or later, you can power on/off specific nodes or the entire system using the graphical user interface. Select Management Menu -> Power Mgt ->Power Off. The compute nodes can be halted from the Scali Manage GUI by selecting the nodes and choosing “halt system” and “power down” from the System Management menu. A command line interface is also available to power-on/off or check status.

Note: Refer to the *Scali Manage User’s Guide* for more information. You must have root privileges to perform these types of tasks.

Powering Off Manually

To power off your cluster system manually, follow these steps:



Caution: If you power off the cluster before you halt the operating system, you can lose data.

1. Shut down the operating system by entering the following command:

```
# init 0
```
2. Press the power button on the head node(s) that you want to power off. You may have to hold the button down for up to 5 seconds. You may power off the nodes in any order.
3. To power off the compute nodes, press the power button (for up to 5 seconds) on the front panel of each unit.
4. To power off the graphics nodes, press the power button (for up to 5 seconds) on the front panel of each unit.
5. To power off optional storage units in the cluster rack, press the power button(s) on their rear panel to the OFF (O) position.

Verifying the Graphics Card Installation for Linux

To verify the graphics card is installed and seen correctly by Linux do the following steps.
From a Linux root prompt enter:

```
# /sbin/lspci | grep VGA
```

Check the output for information similar to the following:

```
01:00.0 VGA compatible controller: nVidia Corporation G80 [Quadro FX xxx] (rev xx)
08:01.0 VGA compatible controller: ATI Technologies Inc ES1000 (rev 02)
```

01:00.0 is the primary card for the system.

08:01.0 is the onboard graphics board (disabled when the primary card is connected)

Note: If the specific card for 01:00.0 is not recognized the `pcids` on the operating system may have to be updated. This occurs if the OS is reinstalled or a newer card from NVIDIA is installed in the system.

Starting the X Window System

The SGI Virtu VN200 requires a number of system-specific X Windows System configuration settings for various configurations. The following section provides an overview on how to customize the `xorg.conf` file.

The X Window System is the only window system provided for the SGI Virtu VN200 running Linux. The SGI Virtu VN200 uses an X server (Xorg) from the opensource xorg project (www.x.org).

The GLX extension, which integrates OpenGL and X, is used by X servers that support OpenGL. This extension is installed with the NVIDIA accelerated drivers.

The following graphics commands and configuration information can help you configure the system.

Commands

`nvidia-xconfig`

Generates X configuration files for the NVIDIA driver

`nvidia-settings`

Configures the NVIDIA driver

`nvidia-installer`

Installs, upgrades, or uninstalls the NVIDIA driver

`nvidia-bug-report.sh`

Gathers information about the system and NVIDIA driver

Documentation

NVIDIA driver documentation is located in `/usr/share/doc/NVIDIA_GLX-1.0/README.txt`. Refer to the man pages for the commands listed above and the X.Org website at: <http://www.x.org>.

Start Graphics

The `init 5` command start the system in the graphical mode. When `init 5` is used, the graphics console displays a login screen. Enter the following command to start graphics:

```
# init 5
or
# startx &
```

Stop Graphics

The X Windows System can be stopped using a variety of commands. From the console/keyboard enter the keyboard command:

CTRL-ALT-BACKSPACE

This kills the current X server and either return to login screen or text console. To stop graphics remotely, enter:

```
# init 3
# /etc/init.d/gdm stop
# killall X
```

Virtual Consoles

Use the following keyboard commands to switch consoles:

CTRL + ALT + F1- graphics (X) to virtual console 1
ALT + F2, ALT + F3, ALT + F4 ... - virtual console 2, 3 , 4 ...
ALT + F7 - back to the graphics console (X).

xorg.conf Configuration File

Xorg.conf is the file containing configuration information for the default X server. It is best to think about the organization of the `xorg.conf` file as tree (refer to [Figure 3-2](#)). The file contains information about the following system properties:

- Graphics
- Displays
- Video Formats
- Mouse and keyboard
- Fonts
- Color names
- Extra modules, drivers, and extensions

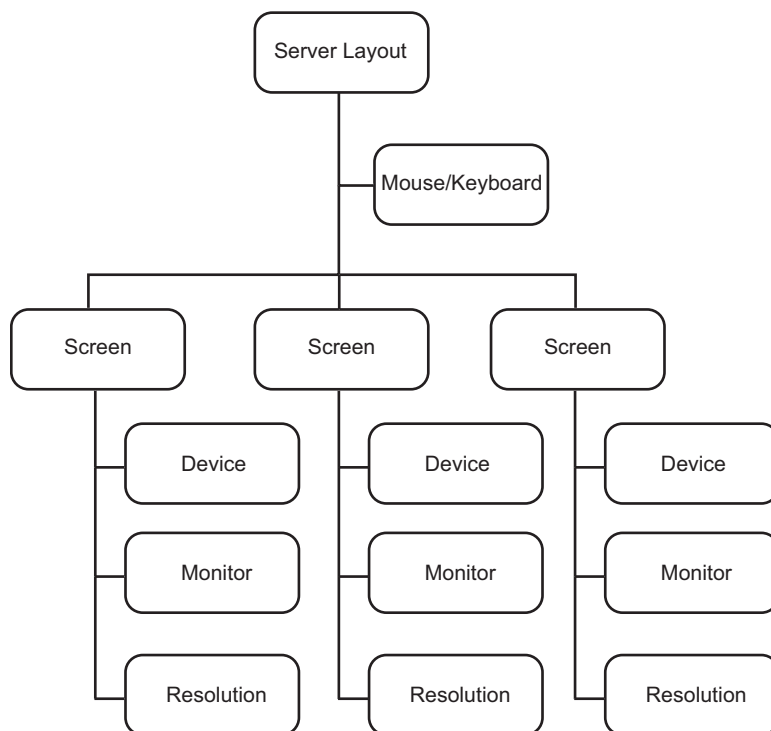


Figure 3-2 Xorg.conf File Organization

The `xorg.conf` file is located in the `/usr/etc/X11` directory. Make a copy of the factory default `xorg.conf` file so that you can make your configuration changes. An example `xorg.conf` file for the SGI Virtu VN200 is listed below for reference.

The X server can be modified by:

- Manually editing `xorg.conf` file
- Using the `/usr/bin/nvidia-xconfig` tool

After editing the `/etc/X11/xorg.conf` file, reboot the system so that the changes can take effect.

Note: Each time X starts a log file is generated in `/var/log/Xorg.0.log`. Common problems that occur after editing the `xorg.conf` file can be solved by examining this file.

Detailed information on editing `xorg.conf` can be obtained by using the `man xorg.conf` command from the Linux prompt. Enter `man nvidia-xconfig` from the linux prompt to learn more about the `xconfig` tool. Also refer to `/usr/share/doc/NVIDIA_GLX-1.0/readme.txt` file.

xorg.conf file for SGI Virtu VN200:

```
# nvidia-xconfig: X configuration file generated by nvidia-xconfig
# nvidia-xconfig: version 1.0 (buildmeister@builder26) Thu Feb 14 18:13:41 PST
# 2008

Section "ServerLayout"
    Identifier      "Layout0"
    Screen          0  "Screen0"
    InputDevice     "Keyboard0" "CoreKeyboard"
    InputDevice     "Mouse0" "CorePointer"
EndSection

Section "Files"
    RgbPath         "/usr/X11R6/lib/X11/rgb"
    FontPath        "/usr/X11R6/lib/X11/fonts/local/"
    FontPath        "/usr/X11R6/lib/X11/fonts/misc/:unscaled"
    FontPath        "/usr/X11R6/lib/X11/fonts/100dpi/:unscaled"
    FontPath        "/usr/X11R6/lib/X11/fonts/75dpi/:unscaled"
    FontPath        "/usr/X11R6/lib/X11/fonts/misc/"
    FontPath        "/usr/X11R6/lib/X11/fonts/Type1/"
    FontPath        "/usr/X11R6/lib/X11/fonts/CID/"
    FontPath        "/usr/X11R6/lib/X11/fonts/100dpi/"
    FontPath        "/usr/X11R6/lib/X11/fonts/75dpi/"
    FontPath        "/usr/X11R6/lib/X11/fonts/truetype/"
EndSection
```

```
Section "Module"
    Load      "dbe"
    Load      "extmod"
    Load      "type1"
    Load      "freetype"
    Load      "glx"
EndSection

Section "InputDevice"
    # generated from default
    Identifier "Mouse0"
    Driver     "mouse"
    Option     "Protocol" "auto"
    Option     "Device"   "/dev/psaux"
    Option     "Emulate3Buttons" "no"
    Option     "ZAxisMapping" "4 5"
EndSection

Section "InputDevice"
    # generated from default
    Identifier "Keyboard0"
    Driver     "kbd"
EndSection

Section "Monitor"
    Identifier "Monitor0"
    VendorName "Unknown"
    ModelName  "Unknown"
    HorizSync  30.0 - 110.0
    VertRefresh 50.0 - 150.0
    Option     "DPMS"
EndSection

Section "Device"
    Identifier "Device0"
    Driver     "nvidia"
    VendorName "NVIDIA Corporation"
EndSection

Section "Screen"
    Identifier "Screen0"
    Device     "Device0"
    Monitor    "Monitor0"
    DefaultDepth 24
    Option     "UseEdid" "True"
    SubSection "Display"
        Depth 24
    EndSubSection
EndSection
```

Monitor Section

The monitor section contains information that defines the capabilities of your display. You can make specific changes to the monitor section of the `xorg.conf` file by referring to the documentation that came with the monitor. An example monitor section is shown below:

```
Section "Monitor"
    Identifier "SDM-P234"
    VendorName "SNY"
    ModelName "3d0"
    HorizSync 28-92
    VertRefresh 57-85

    ModeLine "1920x1200_60n" 154.0 1920 1968 2000 2080 1200 1203 1209 1235
    -hsync -vsync
    ModeLine "1600x1200_60n" 162.0 1600 1664 1856 2160 1200 1201 1204 1250
    -hsync -vsync
EndSection
```

Modelines

Modelines provide video format information to be used with a display. Most displays do not require modelines as the information is returned by the display using Extended Display Identification Data (EDID).

When a display does not provide, or provides incorrect EDID information, a modeline may be required. This is often true for visualization environments with video switching equipment where a EDID information is not returned to the graphics card.

It is possible to create your own modeline using tools such as the one found at:

<http://xtiming.sourceforge.net/cgi-bin/xtiming.pl>

The mode line is contained in the monitor section of the `xorg.conf` file, and is associated with a particular display. An example modeline is shown below:

```
#Modeline description:      pxl  clk  Hres  Hfp  Hsnc  Htot  Vres  Vfp  Vsnc  Vtot
#-----
Modeline "1600x1200_60"    160.875 1600 1704 1808 2160 1200 1201 1205 1242
```

Device Section

The device section defines the configuration parameters for a specific graphics card. There is typically one "Device" section for each graphics card in the system. The device section can make one graphics card with two channels act like two graphics cards using Screen option. A large number of options are available for the setting the device. It is best refer to the NVIDIA README file for specific information about the graphics card configuration parameters. Some examples are given below:

```
Section "Device"
    Identifier    "Device[0]"
    BoardName     "Quadro FX 5600"
    BusID         "1:0:0" # should match up with output of lspci (in decimal)
    Driver        "nvidia" # use "nvidia" (proprietary driver), not "nv" (OSS)!!!
    VendorName    "NVidia"
    Option        "TwinView" "true" # turn on dual channel mode
    Option        "MetaModes" "nvidia-auto-select,nvidia-auto-select"
    Option        "TwinViewOrientation" "LeftOf"
EndSection
```

TwinView

TwinView is an NVIDIA extension with supports multiple displays from a single GPU. It turns on the dual channel mode of the GPU allowing for the a single X server to be configured across the two channels. This has an advantage over other techniques such as Xinerama in that:

- A single X screen is used. The NVIDIA driver conceals all information about multiple display devices from the X server; as far as X is concerned, there is only one screen.
- Both display devices share one frame buffer. Thus, all the functionality present on a single display (e.g. accelerated OpenGL) is available with TwinView.
- No additional overhead is needed to emulate having a single desktop

To enable TwinView, you must specify the following options in the Device section of your `xorg.conf` file:

```
Option "TwinView"
    Option "MetaModes"          "<list of metamodes>"
```

You must also specify either:

```
Option "SecondMonitorHorizSync" "<hsync range(s)>"
Option "SecondMonitorVertRefresh" "<vrefresh range(s)>"
```

or:

```
Option "HorizSync" "<hsync range(s)>"
Option "VertRefresh" "<vrefresh range(s)>"
```

You may also use any of the following options, though they are not required:

```
Option "TwinViewOrientation" "<relationship of head 1 to head 0>"
Option "ConnectedMonitor" "<list of connected display devices>"
```

TwinView Edge-matched Display

The TwinView entries to obtain an edge-matched display (refer to [Figure 3-3](#)) with two channels of 1920x1200 video are:

```
Option "TwinView" "true" # turn on dual channel mode
Option "MetaModes" "1920x1200"
Option "TwinViewOrientation" "RightOf"
```




Figure 3-3 TwinView Edge-Matched Display

TwinView Edge Blending

To configure TwinView displays that require edge-blending (refer to [Figure 3-4](#)) between the two display channels with a 10% overlap, use the following entries:

```
Option      "TwinView" "true" # turn on dual channel mode
Option      "MetaModes" "1920x1200, 1920x1200 +1728+0"
```

The +1728 in the option line above commands the graphics card to start the second channel at 1728 offset from the lower left hand corner of the first channel (the origin).




Figure 3-4 TwinView Edge Blending with 10% Overlap

Stereoscopic Viewing

Configure the following option in the `xorg.conf` file to enable stereoscopic viewing:

```
Option      "stereo"      "integer"
```

For active stereo, you use special shutter glasses. The stereo emitter for shutter glasses connects to the 3-pin DIN connector above the DVI-I outputs on the graphics card (refer to [Figure 2-5](#) on [page 25](#)). Configure the following option in the `xorg.conf` file to configure active stereo for shutter glasses:

```
Option      "Stereo"      "3"
```

Stereo option 3 (active stereo) may be used with TwinView if all modes within each MetaMode have identical timing values.

Use stereo option 4 (passive stereo) to view channel 1 in the left eye and channel two in the right eye.

```
Option      "Stereo"      "4"
```


Passive stereo uses the TwinView clone mode stereo and is not compatible with any other TwinView settings.

Note: Stereo is not compatible with the “Composite extension being set”.

Please refer to the NVIDIA README file for other stereo options and limitations.

Setting up Multiple X Servers on a Single Graphics Card

The NVIDIA graphics card can be configured such that each connected display device functions as a separate X server. To configure multiple X servers:

First, create two separate Device sections, each listing the BusID of the graphics card to be shared and listing the driver as "nvidia", and assign each a separate screen:

```
Section "Device"
    Identifier "nvidia0"
    Driver     "nvidia"
    # Edit the BusID with the location of your graphics card
    BusID     "PCI:1:0:0"
    Screen    0
EndSection

Section "Device"
    Identifier "nvidia1"
    Driver     "nvidia"
    # Edit the BusID with the location of your graphics card
    BusID     "PCI:1:0:0"
    Screen    1
EndSection
```

Then, create two Screen sections, each using one of the Device sections:

```
Section "Screen"
    Identifier "Screen SG-0"
    Device     "nvidia0"
    Monitor    "Monitor0"
    DefaultDepth 24
    Subsection "Display"
        Depth     24
        Modes      "1600x1200" "1024x768" "800x600" "640x480"
    EndSubsection
EndSection
```

```
Section "Screen"
    Identifier "Screen SG-1"
    Device     "nvidial"
    Monitor    "Monitor1"
    DefaultDepth 24
    Subsection "Display"
        Depth     24
        Modes      "1600x1200" "1024x768" "800x600" "640x480"
    EndSubsection
EndSection
```

ServerLayout Section

The `ServerLayout` section defines the input/output keyboard/mouse is associated with a screen and where the multiple screens are located relative to one another. This example shows two screens:

```
Section "ServerLayout"
    Identifier   "Layout [all]"
    InputDevice "Keyboard0" "CoreKeyboard"
    InputDevice "Mouse1" "CorePointer"
    Option      "Xinerama" "off"
    Screen      0 "Screen SG-0"
    Screen      1 "Screen SG-1" RightOf "Screen SG-0"
EndSection
```

This example show a `ServerLayout` section with two separate keyboard/mouse entries:

```
Section "ServerLayout"
    Identifier   "Seat0"
    InputDevice "Keyboard0" "CoreKeyboard"
    InputDevice "Mouse0" "CorePointer"
    Option      "IsolateDevice" "True"
    Screen      0 "Screen SG-0"
EndSection
```

```
Section "ServerLayout"
    Identifier   "Seat1"
    InputDevice "Keyboard1" "CoreKeyboard"
    InputDevice "Mouse1" "CorePointer"
    Option      "IsolateDevice" "True"
    Screen      0 "Screen SG-1"
EndSection
```

Check the NVIDIA Driver

To check that the NVIDIA driver is installed correctly, verify that the `xorg.conf` file contains the following entry under `Devices`:

```
Section "Device"
    Identifier      "Device0"
    Driver          "nvidia"
    VendorName     "NVIDIA Corporation"
EndSection
```

If the driver is not listed as `"nvidia"` then the `xorg.conf` file is not using the NVIDIA accelerated drivers.

Re-installation of the Graphics Drivers

After a clean install of the OS, the `.rpm` files can be re-installed by using one of the following procedures.

- Download the latest 64-bit linux driver from NVIDIA (www.nvidia.com) and follow the instructions provided with the driver.

Note: The NVIDIA installer may require you to install additional packages on the system.

- Install the `.rpm` files supplied by SGI

For SLES 10 with the NVIDIA 169.12 driver, install the following `.rpm` files:

```
rpm -ivh x11-video-nvidiaG01-169.12-3.rpm
nvidia-gfxG01-kmp-smp-169.12_2.6.16.46_0.12-3.x86_64.rpm
```

For Redhat with the NVIDIA 169.12 driver, install the following `.rpm` files:

```
rpm -ivh nvidia-graphics169.12-169.12-96.el5.x86_64.rpm \
nvidia-graphics169.12-devel-169.12-96.el5.x86_64.rpm \
nvidia-graphics169.12-kmdl-2.6.18-53.el5-169.12-96.el5.x86_64.rpm \
nvidia-graphics169.12-libs-169.12-96.el5.x86_64.rpm \
nvidia-graphics-devices-1.0-5.0.el5.noarch.rpm \
nvidia-graphics-helpers-0.0.26-27.el5.x86_64.rpm
```

Upgrading Drivers

SGI provides updated drivers through SupportFolio. If you upgrade drivers directly from the NVIDIA website, remove the `.rpm` files listed on [page 60](#) first to ensure that there are no conflicts when installing the new drivers.

Troubleshooting

General

Use `ipmitool` or the IPMI web interface to check the following resources for signs of a system failure:

- Server Event Log (SEL)
- Console power-on self-test (POST) code messages
- System LEDs

Use `ipmitool` or the IPMI web interface to check the following sensor data for signs of a system failure:

- Fans
- Voltages
- Temperature

If the system hangs during POST remove components in the following order and reset the system to see if the problem goes away:

1. Graphics board
2. Disk drives
3. FB-DIMMs (do not remove pair 1)
4. Processors (do not remove CPU 1)

Graphics

The following list describes basic troubleshooting tips that you can use to isolate graphics problems.

- Each time X starts a log file is generated in `/var/log/Xorg.0.log`. Common problems that occur after editing the `xorg.conf` file can be solved by examining this file.
- Look for visual artifacts on display device:
 - Vertical or horizontal stripes
 - Pixel drop-outs
 - Missing colors
- Run graphics demos to verify performance characteristics of the graphics hardware.
- Run the `glxinfo` command to verify that direct rendering (DRI) is enabled and the OpenGL vendor string is set to NVIDIA Corporation; this confirms the graphics board is using the NVIDIA driver and not the default RHEL/SLES driver
- Run the `nvidia-bug-report.sh` script to verify that the NVIDIA driver installed correctly
- If necessary, reinstall the graphics drivers; after reinstalling, run the `ldconfig` command
- Use the `nvidia-xconfig` command to generate a new `xorg.conf` file
- Swap the graphics card, display device, and display cables to isolate failures

BMC Configuration

The following list describes basic troubleshooting tips that you can use to isolate BMC configuration problems.

- By default, the SGI Virtu VN200 system enables serial-over-LAN (SOL) functionality via serial port B.
- To configure SOL manually, enable serial-console redirection in BIOS, then configure the BMC LAN to assign a static IP address (XE) or configure the BMC as a DHCP client (ICE) via one of the following methods:

- Use the Supermicro `xgicfg` command from a DOS-bootable thumb-drive image
- Use in-band `ipmitool` from the operating-system prompt
- Use the `ipmitool` command or the AOC-SIMSO web interface to communicate with the BMC

Caution: The DEFAULTP.EXE command resets all BMC network settings to their default setting.

- If the BMC does not respond, you can reset it to default values with the SuperMicro DEFAULTP.EXE command from a DOS-bootable thumbdrive image

Power

Use the following tools and techniques to troubleshoot the power components on the SGI Virtu VN200:

- Check the power supply status LEDs. Each power supply has a green LED that illuminates when DC power is present. When a supply fails, the PWR FAIL LED on the control panel illuminates and the power-supply status LED does not illuminate.
- Use the IPMI web interface, the `ipmitool` command, or the power-supply LEDs to determine which supply has failed
- Hot-swap the failing supply without powering the node down (if the node has a redundant power supply) to perform fault isolation

Cooling

If a fan or temperature fault occurs, the red ALERT LED on the control panel illuminates.

- Use `ipmitool` or the IPMI web interface to determine which fan has failed
- Remove the graphics node from the rackmount shelf and open the side access panel to replace the failed fan

Hardware Replacement Procedures

Before You Begin

Before working with your server product, pay close attention to “[Warnings and Cautions](#)” on page 4.

Tools and Supplies Needed

- Phillips (cross head) screwdrivers (#1 bit and #2 bit)
- Needle-nosed pliers
- Antistatic wrist strap and conductive foam pad or ESD-safe worksurface

Removing the Front Bezel

Remove two screws that secure the front bezel to the rack ([Figure 3-5](#)) and remove the bezel.




Figure 3-5 Front Bezel Removal

Removing the Graphics Node from the Rack

1. Power off the system using ipmitool or BMC web interface.
2. Disconnect the AC input power cord from the rear panel of system.
3. Disconnect the video and data cables from the rear panel of the system.
4. Remove the front bezel from the system in accordance with the instructions on [page 64](#).
5. Slide the graphics node out of the rack and place it on an ESD-safe worksurface.




Figure 3-6 Removing the Graphics Node from a Rack

Removing the Side Cover

1. Remove the graphics node from the rack in accordance with the instructions on [page 65](#).
2. Remove the two screws ([Figure 3-7](#)) that secure the side cover to the chassis and remove the side cover.





Figure 3-7 Side Cover Removal

Replacing the Power Supplies

Note: Power supplies can be hot-swapped while the system is in operation in a rack. Have the replacement power supply ready, so that it can be installed quickly without disrupting the airflow through the chassis.

1. If necessary, remove the front bezel from the system in accordance with the instructions on [page 64](#).
2. Move the power supply latch to release the power supply (refer to [Figure 3-8](#)).
3. Grasp the handle and pull the supply out of the chassis.
4. Slide the replacement supply in the chassis, and press it in until the latch can lock it into place.
5. Verify the status LED on the replacement power supply is green.




Figure 3-8 Power Supply Removal

Replacing a Disk Drive

Note: Halt the operating system before replacing a disk drive. Disk drives can be hot-plugged while the system is in operation in a rack. Have the replacement disk drive ready, so that it can be installed quickly without disrupting the airflow through the chassis.

1. If necessary, remove the front bezel from the system in accordance with the instructions on [page 64](#).
2. Release the latch on the front of the disk drive carrier (refer to [Figure 3-9](#)).
3. Pull on the black lever to unseat the drive and slide it out of the chassis.




Figure 3-9 Disk Drive Removal

Replacing the Chassis Fans

1. Remove the side cover from the chassis in accordance with the instructions on [page 66](#).
2. Remove two screws (refer to [Figure 3-10](#)) that secure the fan assembly to the chassis.
3. Disconnect the fan power cables from the node board.
4. Remove the fan assembly.




Figure 3-10 Cooling Fans Removal

Replacing the Graphics Card

Note: The power cable may be plugged into a different location than the power cable shown in [Figure 3-11](#).

1. Remove the side cover from the chassis in accordance with the instructions on [page 66](#).
2. Disconnect the graphics card power cable (refer to [Figure 3-11, A](#)).

Note: Step 3 is not required for the NVIDIA Quadro FX 1700 graphics cards.

3. If necessary, remove the three screws that secure the graphics card to the chassis air baffle (refer to [Figure 3-11, B](#)).
4. Remove two screws (refer to [Figure 3-11, C](#)) that secure the graphics card to the rear panel.
5. Unseat the graphics card from the PCI riser board connector and lift the card out of the chassis (refer to [Figure 3-11, D](#)).
6. Place the graphics card on an ESD-safe worksurface.




Figure 3-11 Graphics Card Removal

Replacing the AOC-SIMSO (BMC) Daughter Card

The AOC-SIMSO daughter card (BMC) resides in a slot attached to the node board. The rear panel BMC LAN port is connected to the BMC using a short USB cable.

1. Remove the graphics card in accordance with the instructions on [page 69](#).
2. Press the retaining bracket arms apart and gently lift the end of the daughter card that is not in the bracket (refer to [Figure 3-12, A](#)).
3. Disconnect the daughter card from the node board connector and remove the daughter card (refer to [Figure 3-12, B](#)).




Figure 3-12 BMC Daughter Card Removal

Removing the PCI Riser Board

The PCI riser board connects between the graphics card and the node board.

1. Remove the graphics card in accordance with the instructions on [page 69](#).
2. Remove the screws that secure the riser board to the chassis (refer to [Figure 3-13](#)).

Two options are available:

- a. Remove the two screws that attach the riser board inside the chassis.
 - b. Remove the two screws that attach to the outside of the chassis.
3. Gently disconnect the PCI riser board from the PCIe x16 slot on the node board and lift the riser board out of the chassis.




Figure 3-13 PCI Riser Board Removal

Removing the Air Baffles

The air baffles deflect the airflow from the chassis fans across the node board and the graphics card.

1. Remove the graphics card in accordance with the instructions on [page 69](#).
2. Remove two screws on each side of the chassis (refer to [Figure 3-14](#)).
3. Lift the plate out of the chassis.




Figure 3-14 Air Baffle Removal

Replacing Node Board Memory DIMMs

1. Remove the graphics card in accordance with the instructions on [page 70](#).
2. Remove the air baffle in accordance with the instructions on [page 73](#), (refer to [Figure 3-15, A and B](#)).
3. Press down on the latches to release the DIMM (refer to [Figure 3-15, C](#)).
4. Carefully grasp the DIMM by its edge and pull it up and out of the guide rails.




Figure 3-15 DIMM Removal

Replacing Node Board Processors



Caution: When handling the processor, avoid placing direct pressure on the label area of the fan.



Caution: Do not place the node board on a conductive surface. A conductive surface can damage the BIOS battery and prevent the system from booting. Always place the node board on an ESD-safe worksurface.

1. Remove the graphics card in accordance with the instructions on [page 69](#).
2. Remove the air baffles from the chassis (refer to [Figure 3-16](#)).






Figure 3-16 Node Board Removal

3. Remove the heatsink assemblies
4. Press the load lever down and away from the retention clasp to release the load plate from its locked position.



5. Gently lift the load lever to open the load plate.



6. Use your thumb and your index finger to hold the CPU at opposite sides.
7. Carefully lift the CPU straight out of the socket.




Note: If you are removing the processor so that you can return the node board to SGI, install a PnP cap in the socket.


Install a Processor

Note: A black pick-and-place (PnP) cap is attached to the load plate to protect the CPU socket.


1. Press the load lever down and away from the retention clasp to release the load plate from its locked position.



2. Gently lift the load lever to open the load plate.



Load plate
released



3. Use your thumb and your index finger to hold the CPU at opposite sides.




- Align pin 1 of the CPU (the corner marked with a triangle) with the notched corner of the CPU socket.




Caution: In step 5, do not drop the CPU on the socket. Do not move the CPU horizontally or vertically, and do not rub the CPU against any surface or any of the contacts which may damage the CPU and/or contacts.

- Once aligned, carefully lower the CPU straight down into the socket.
- Find the corner of the CPU that has a semi-circle cutout below a gold dot (CPU key). This corner should be aligned with the cutout on the socket (socket key).



- With the CPU in the socket, inspect the four corners of the CPU to make sure that it is properly installed.


8. Use your thumb to gently push the load lever down until it snaps into the retention clasp.



9. If the CPU is properly installed into the socket, the PnP cap is automatically released from the load plate when the lever locks.




Caution: Retain the plastic PnP cap for shipping. The node board must be shipped with the PnP cap properly installed to protect the CPU socket. Shipment without the PnP cap properly installed may void the warranty.



10. Remove the PnP cap.

PnP cap released
from load plate



11. Repeat these steps to install a second CPU if desired.

Install a Heatsink

1. Ensure that you have replaced the thermal interface material (TIM) between the heatsink and the CPU die.

Refer to section [“Replace the Thermal Interface Material”](#) on page 83.



Caution: Use caution when you unpack the heat sink so you do not damage the TIM.

2. Place the heatsink on top of the CPU so that the four mounting holes are aligned with those on the (preinstalled) heatsink retention mechanism.



Caution: Do not tighten the screws in step 3 or you may damage the CPU.

3. Install two screws (screw 1 and screw 2, shown in [Figure 3-17](#)) and tighten until just snug.



Caution: Do not overtighten the screws in step 4 or you may damage the CPU.

4. Install the two remaining screws (screw 3 and screw 4, shown in [Figure 3-17](#)); then, gently tighten all four screws.




Figure 3-17 Install a CPU Heatsink

Remove the Heatsink

SGI does not recommend that the CPU or the heatsink be removed from the node board. However, if you do need to remove the heatsink, use the instructions below to prevent damage to the CPU or the CPU socket.

1. Unscrew and remove the heatsink screws from the node board in the sequence as shown in [Figure 3-17](#).



Caution: Do not use excessive force when loosening the heatsink. Using excessive force can damage the heatsink, CPU, or node board.

2. Grasp the heatsink and gently loosen it from the CPU.
3. When the heatsink is loose, remove it from the CPU.
4. Clean the surface of the CPU and the heatsink to remove the old thermal grease.
5. Refer to section “[Replace the Thermal Interface Material](#)” on page 83.
6. Apply a small amount of thermal interface material (TIM) on the surface of the CPU before you install the heatsink.

Replace the Thermal Interface Material

The Intel Xeon processors used in your server require a thermal interface material (TIM) between the processors and heatsinks to properly dissipate the heat that is generated by the processors. If you replace a processor and/or heatsink, you will need to remove the existing TIM from the processor and heatsink. Then, you will need to apply new TIM to the processor before you install the heatsink.

The TIM application instructions accompany the new TIM. The following subsections describe how to remove the existing TIM:

- “[Removing TIM from a Processor](#)” on page 84
- “[Removing TIM from a Heatsink](#)” on page 87

Removing TIM from a Processor

You should always remove the TIM from a processor before you remove it from the node board.



Warning: Always wear gloves when performing this procedure to avoid direct skin contact with the TIM, which is a potentially toxic material.

Tools/Parts Required:

- Dry, lint-free cloth
- Lint-free cloth moistened with isopropyl alcohol
- Gloves



Caution: To prevent the processor from moving in its socket, ensure that the processor load lever remains engaged while you perform the following procedure. If the processor moves while you remove the TIM, the pins on the processor might be damaged.

Perform the following steps to remove the TIM from a processor:

1. Remove the heatsink from the socket.



Note: If you plan to re-use the heatsink with a replacement processor, you must clean the TIM off of the heatsink. (Refer to [“Removing TIM from a Heatsink”](#) on page 87.)


2. Gently push any loose TIM pieces to the center of the processor.



3. Use a dry lint-free cloth to lift the loose TIM pieces off of the processor.



4. Use a dry, lint-free cloth to wipe any remaining TIM pieces off of the processor.



5. Use a different lint-free cloth moistened with isopropyl alcohol (IPA) to remove any remaining TIM from the processor.



Note: Be sure to remove TIM from the gaps between the processor and the load plate.

Removing TIM from a Heatsink

If you plan re-use the heatsink with a replacement processor, you need to remove the TIM from the heatsink.

Warning: Always wear gloves when performing this procedure to avoid direct skin contact with the TIM, which is a potentially toxic material.

Tools/Parts Required:

- Dry, lint-free cloth
- Lint-free cloth moistened with isopropyl alcohol
- Gloves

Perform the following steps to remove the TIM from a heatsink:

1. Remove the heatsink from the socket and place the heatsink on a secure surface with the bottom facing upward.
2. Use a dry, lint-free cloth to wipe any loose TIM pieces off of the heatsink.



3. Use a different lint-free cloth moistened with isopropyl alcohol (IPA) to remove any remaining TIM from the heatsink.



Replacing the Node Board

1. Remove the graphics card in accordance with the instructions on [Figure 3-11](#).
2. Remove the air baffles (refer to [Figure 3-18](#)).




Figure 3-18 Air Baffles

3. Remove the BMC daughter card in accordance with the instructions on [page 71](#).
4. Remove the PCI riser board in accordance with the instructions on [page 72](#).
5. Remove the memory DIMMs in accordance with the instructions on [page 74](#).
6. Remove the heatsinks and processors in accordance with the instructions on [page 75](#).
7. Disconnect the SATA and power cables that attach to the node board.
8. Remove the screws that secure the node board to the chassis.

Replacing the Control Panel

The control panel is located at the front of the system (refer to [Figure 3-19](#)).

1. Remove the disk drives or the front top cover to access the cable that connects to the control panel.
2. Disconnect the cable connector that attaches to the control panel cable.
3. Remove the control panel from the chassis.




Figure 3-19 Control Panel Removal

