

SGI® CloudRack™ X2 System User's Guide

Document Number 007-5573-001

COPYRIGHT

© 2010 SGI. All rights reserved; provided portions may be copyright in third parties, as indicated elsewhere herein. No permission is granted to copy, distribute, or create derivative works from the contents of this electronic documentation in any manner, in whole or in part, without the prior written permission of SGI.

LIMITED RIGHTS LEGEND

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

The electronic (software) version of this document was developed at private expense; if acquired under an agreement with the USA government or any contractor thereto, it is acquired as "commercial computer software" subject to the provisions of its applicable license agreement, as specified in (a) 48 CFR 12.212 of the FAR; or, if acquired for Department of Defense units, (b) 48 CFR 227-7202 of the DoD FAR Supplement; or sections succeeding thereto. Contractor/manufacturer is SGI, 46600 Landing Parkway, Fremont, CA 94538.

TRADEMARKS AND ATTRIBUTIONS

SGI, and the SGI logo are registered trademarks and CloudRack is a trademark of Silicon Graphics International Corporation or its subsidiaries in the United States and/or other countries worldwide.

Intel, Atom and Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Athlon, Opteron and Phenom are trademarks or registered trademarks of Advanced Micro Devices Corporation.

UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company, Ltd.

InfiniBand is a trademark of the InfiniBand Trade Association.

Linux is a registered trademark of Linus Torvalds.

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

Record of Revision

Version	Description
-001	March, 2010 First release

Contents

1. Introduction and Overview	1
ESD and Safety Precautions	1
Upgrading or Replacing Components	2
Overview	2
Example Nodeboard Features	3
Processors	3
DIMM Memory	4
System Power Supplies	4
Compute Tray Drive Subsystem	4
Compute Tray Front I/O Panel	6
Cooling System	7
Main Enclosure Cooling	7
Integral Power Supply Cooling	7
Motherboard Example Diagram	8
2. System Operation and Troubleshooting	11
Unpacking the System and Choosing an Operating Location	11
Placing Rackmounted Systems	12
Choosing a Setup Location	12
Rack Precautions	12
Server Precautions	12
Rack Operating Considerations	13
Ambient Operating Temperature	13
Reduced Airflow	13
Mechanical Loading	13
Circuit Overloading	13
Reliable Ground	13
Providing Power	14

Troubleshooting the System	17
Enclosure Power Supply Status LEDs	18
Individual Tray LEDs	19
System Fan Failure	20
3. System Interfaces Overview	21
NICs 1 and 2 (HPC trays only)	21
HDD	21
Power	21
Tray Control Panel Button Examples	21
Overheat/Fan Fail LED on Back of System.	22
4. HPC Server BIOS Information	23
Starting the BIOS Setup Utility	23
How To Change the Configuration Data.	24
Starting the Setup Utility	24
Main Setup Screen	24
Advanced Setup Configurations.	26
BOOT Features	26
Remote Access Configuration	37
Hardware Health Monitor	38
IPMI Configuration.	41
SEL PEF Configuration	44
The DMI Event Log	45
Security Settings	45
Exit Options.	47
BIOS Error Beep Codes	48
BIOS Error Beep Code List	49
A. Technical Specifications	51
Server Specifications and Features	51
Environmental Specifications	52

Introduction and Overview

This chapter provides an overview of your SGI CloudRack X2 workgroup cluster server's main features.

Operating precautions are provided in this chapter, followed by a general overview of the product.

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

ESD and Safety Precautions

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

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



Warning: Before operating or servicing any part of this product, read the safety precautions.



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.

Upgrading or Replacing Components

The *SGI CloudRack X2 Component Replacement Guide* (P/N 007-5648-00x), describes how to install or replace the following components in an SGI CloudRack X2 cluster server:

- Memory DIMMs
- PCIe cards
- Disk drives
- System fans
- Power supplies

Use the procedures to upgrade system components or replace failing components.



Warning: If a lithium battery is installed in your system as a soldered part, only SGI qualified 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.

Overview

The CloudRack X2 cluster server is comprised of a stand-alone or rack-mounted enclosure and holds up to nine (9) compute “trays” that house serverboards, power supplies, hard drives and graphics or I/O options. Stand-alone enclosures are mounted on castors so they can be moved within the server room or compute lab environment. Check with your sales or service representative before loading any operating system on your server not provided by the SGI factory or service organization.

In addition to the compute trays and chassis, various hardware components may be included as part of your CloudRack X2 configuration as listed below:

- SATA Accessories: A minimum of one disk drive (per serverboard) is required for operation. One (1) internal SATA backplane per compute node. One (1) SATA cable set per compute node. SAS hard drive options are also available.
- Two (2) optional PCI Express x16 riser cards (one per compute serverboard).

- Rackmount hardware for mounting the enclosure in a half-height or full-height rack.
- One (1) CD containing drivers and utilities plus optional CDs depending on order configuration.
- One or two Gbit Ethernet or (optional) InfiniBand switches are used per enclosure.

Example Nodeboard Features

At the heart of each CloudRack X2 compute tray lies one or more multi-processor based node boards (serverboards). These serverboards may be based on Intel or AMD chipsets, depending on the configuration ordered.

Processors

The cluster server system can support a minimum of 8 and a maximum of up to 216 processor cores. The exact type of processors provided with your system depends on the specific configuration you ordered. Check with your sales or service representative for information on processor availability, upgrades and compatibility. The following are examples of types available for the CloudRack X2:

- One or two dual-socket Intel® Xeon® 5500 processor series-based serverboards per tray
- Six single-socket Intel® Atom® processor series-based MicroSlice serverboards per tray
- Two dual-socket, Quad- or Six-Core AMD Opteron™ 2300 or 2400 processor series-based serverboards per tray
- Three single-socket AMD Phenom™ X4 Quad-Core processor series-based MicroSlice serverboards per tray
- Six single-socket AMD Athlon™ X2 Dual-Core processor series-based MicroSlice serverboards per tray

Higher-performance compute (HPC) serverboards support two Intel® Xeon quad-core processors (a total of 4 quad-core processors per compute tray).

DIMM Memory

Memory configuration varies depending on the processor board type ordered with your system tray(s). DIMM population requirements vary from one per serverboard up to 12 per serverboard, using DIMM speeds from 667 MHz to 1333 MHz.

The HPC serverboards in each compute tray have up to six 240-pin DIMM sockets that can support up to 48 GB of registered ECC DDR3-1333/1066/800 SDRAM (96 GB total for the compute tray). As noted, this capacity can vary based on which compute trays your SGI CloudRack X2 is using. In all configurations a minimum of 3GB of DIMM per core installed is recommended for optimum performance.

The *SGI CloudRack X2 Component Replacement Guide* provides details on replacing or installing DIMMs in the system.

Server Chassis Features

The following sections provide a general outline of the main features of the CloudRack X2 system chassis (enclosure).

Figure 1-1 shows the front of the SGI CloudRack X2 server.

System Power Supplies

Your CloudRack X2 server uses three high-efficiency power supplies; each supply provides up to 240 Amps of 12-volt power to the system. The power supplies require a 200-240 Volt power cable input with a 47 Hz to 63 Hz operational range.



Caution: The chassis power supply cord is used as the main disconnect device. Ensure that the socket outlet is located or installed near the equipment and is easily accessible.

Compute Tray Drive Subsystem

Each drive tray in the CloudRack X2 chassis was designed to support up to eight SATA or SAS hard disk drives. Specific optional trays may be available with up to four drives configured for each serverboard in the tray. Note that the standard trays provided in the CloudRack X2 only support SATA disk drives. However, by adding an optional SAS HBA certain versions of the

compute trays can support SAS disk drives. The SAS HBA will also support RAID 0 and 1 on both SAS and SATA disk drives. One HBA is added to each node of the compute tray.

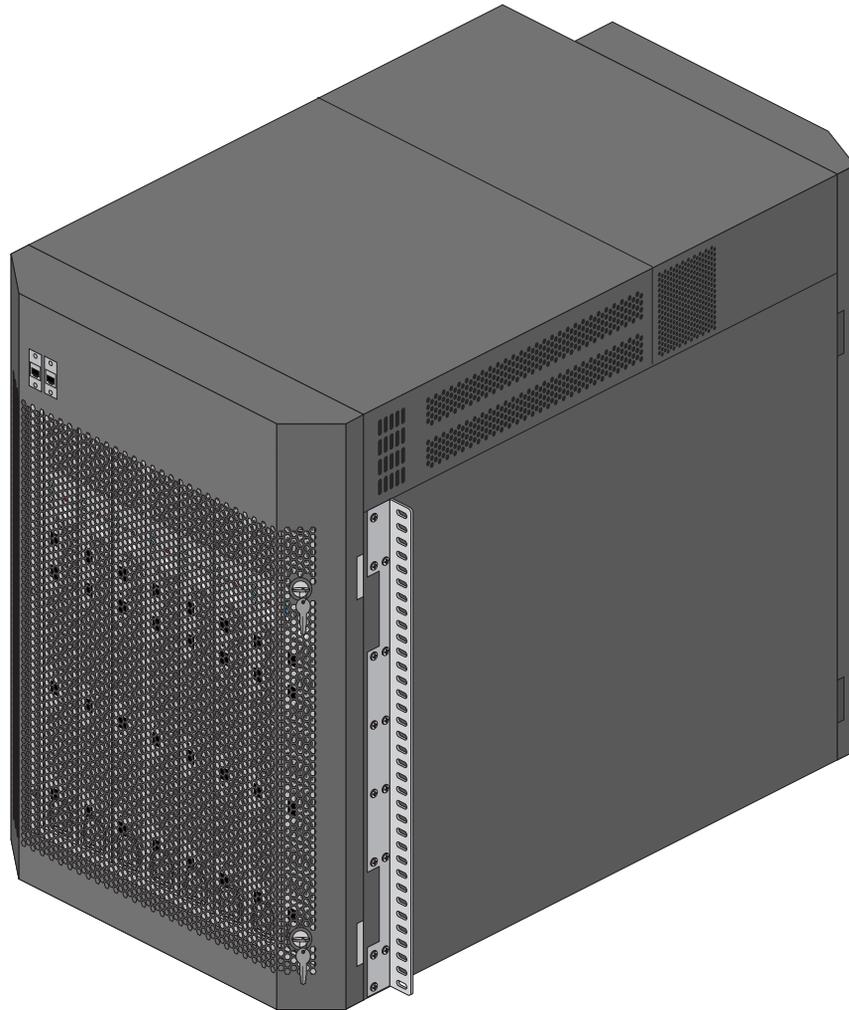


Figure 1-1 CloudRack X2 Enclosure With Front Door Closed

Compute Tray Front I/O Panel

Each system compute tray installs via the front of the CloudRack X2 chassis. Figure 1-2 shows an example compute tray front panel. Its I/O panel typically provides COM ports, status LEDs and Gb Ethernet ports.

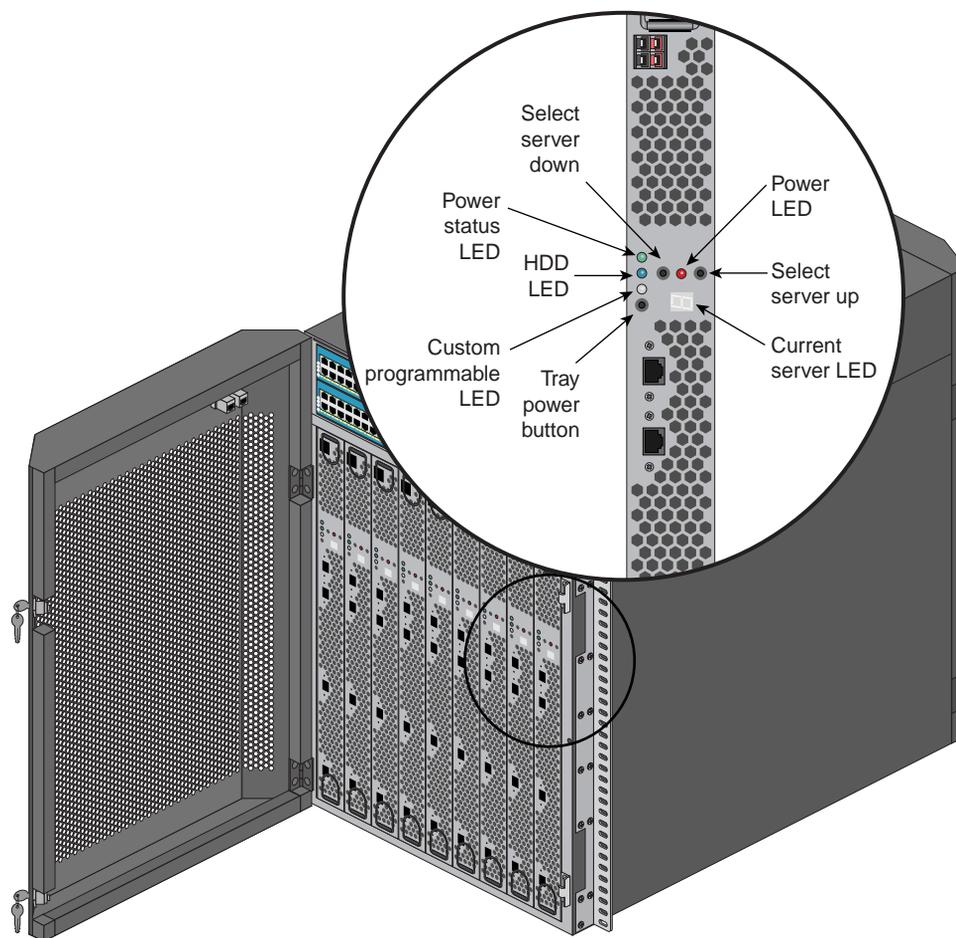


Figure 1-2 Compute Tray Front Panel Example

Cooling System

The server chassis has a staggered cooling design that features an enclosed nine-fan primary air shroud and additional upper cooling array (three power supply fans). A fan speed control setting sensor in the main fan assemblies can increase or decrease fan speed based on ambient temperature.

Main Enclosure Cooling

The primary cooling for the enclosure is an array of three sets of three 5-inch (120-mm) fans. Figure 1-3 on page 8 shows the rear of the system with the door removed. These three fan assemblies (nine fans total) provide the primary cooling for the system compute trays and their supported optional I/O and drive hardware.

Integral Power Supply Cooling

The upper cooling fans (seen in Figure 1-3 on page 8) provide air for the three system power supplies. These 60-mm (2.4-inch) fans are each an integral part of the power supply unit. Each system power supply (located in the upper section of the enclosure) uses a separate, independent cooling fan. See the *SGI CloudRack X2 Component Replacement Guide* for information on replacing a failed power supply.

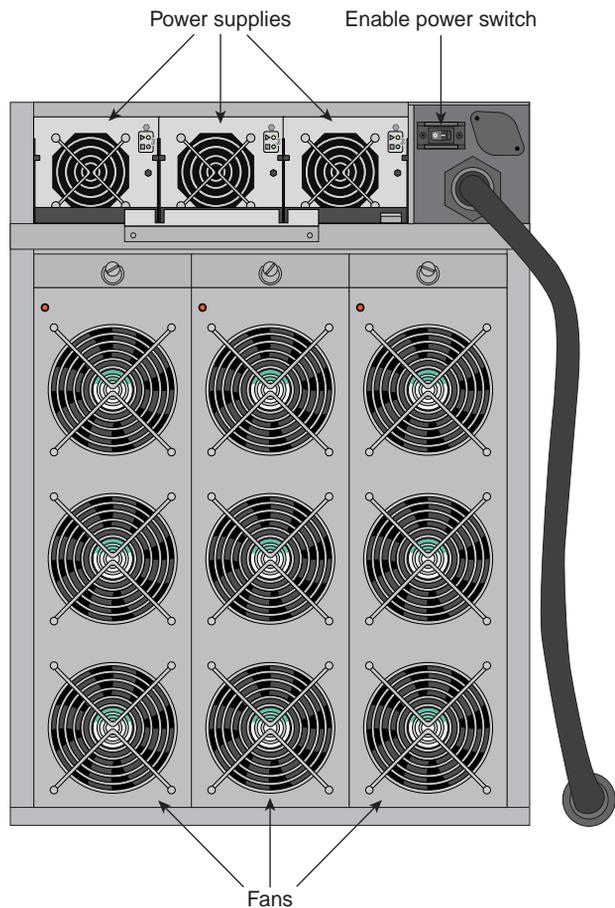


Figure 1-3 CloudRack X2 Cooling Fans on Rear of Enclosure

Motherboard Example Diagram

The SGI CloudRack X2 can be configured with trays that offer higher processor counts or higher compute power. Figure 1-4 on page 9 shows an example diagram of a high-performance computing (HPC) motherboard based on Intel chip sets.

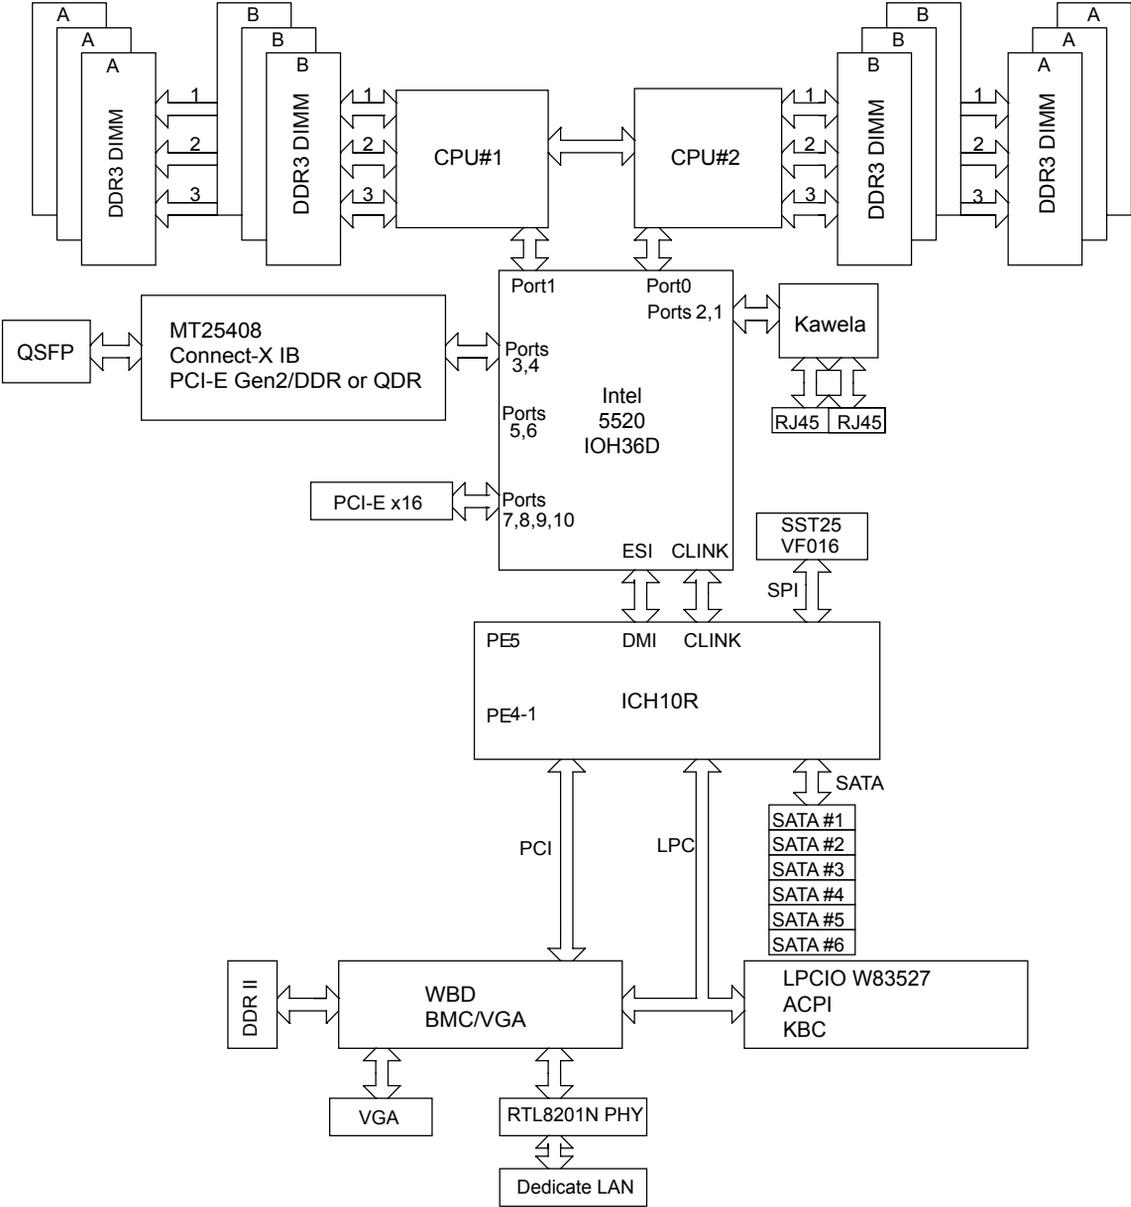


Figure 1-4 HPC Motherboard Block Diagram Example

System Operation and Troubleshooting

The first half of this chapter describes the basic steps needed to get your SGI CloudRack X2 up and running. Following these steps in the order given should enable you to have the system operational within a minimum amount of time. The second half of this chapter provides you with some basic troubleshooting advice. Use these sections to eliminate simple problems or obtain information that may be needed by your service provider.

Unpacking the System and Choosing an Operating Location

You should inspect the box the system was shipped in and note if it was damaged in any way. If the server itself shows damage you should file a damage claim with the carrier who delivered it.

When you decide on a suitable location for the stand-alone or rack unit holding the system, it should be situated in a clean, dust-free area that is well ventilated. Avoid areas where heat, electrical noise and electromagnetic fields are generated. You will also need it placed near a dedicated 200-240 Volt grounded single-phase (L6-30) power outlet or 3-phase power outlet.

The stand-alone version of the CloudRack X2 is designed to fit into a computer lab or server room environment. Take care to maintain the following operating conditions:

- The stand-alone enclosure should have a six-inch (15 cm) minimum top air clearance.
- The stand-alone enclosure should have a minimum of 8 inches (20 cm) of front and rear clearance to allow for proper cooling.

Note: Opening the front or rear cover of the enclosure will require more working space than the minimum cooling space listed above.

- The enclosure should be protected from harsh environments that produce excessive vibration and heat.
- The enclosure should be kept in a clean, dust-free location to reduce maintenance problems.
- Available power must be rated for large computer operation (30 amps at 200-240 Volts).

Placing Rackmounted Systems

Be sure to read the Rack Precautions section if you are having the enclosure installed on site.

Choosing a Setup Location

Leave enough clearance in front of the rack to enable opening the front door completely ~48 inches (1.2 meters).

Leave sufficient clearance in the back of the rack to allow for adequate airflow and ease in servicing.

Rack-mounted systems are generally placed in a Restricted Access Location (dedicated equipment rooms, service closets and the like).

Rack Precautions

Ensure that the leveling jacks on the bottom of the rack are fully extended to the floor with the full weight of the rack resting on them.

The enclosure should be installed in the lowest part of the rack possible.

In a tall rack installation, stabilizers should be attached to the rack if available.

Always make sure the rack is stable before connecting power to the rack and internal inclosure.

Always keep the rack's front door and all panels and components on the servers closed when not servicing to maintain proper cooling.

Server Precautions

Review the electrical and general safety precautions in Chapter 1.

For extra protection, use a regulating uninterruptible power supply (UPS) to protect the cluster server's power supplies from power surges, voltage spikes and to keep your system operating in case of a power failure. This is an optional device not provided by SGI with your system.

Service personnel should always allow the hot plug disk drives and power supply modules to cool before touching them. To maintain proper cooling, always keep the enclosure doors closed when it is not being serviced.

Make sure all power and data cables are properly connected and not blocking the enclosure airflow.

Rack Operating Considerations

Use the guidelines in the following subsections to properly use and maintain a server in a rack.

Ambient Operating Temperature

If installed in a closed or multi-unit rack assembly, the ambient operating temperature of the rack environment may be greater than the ambient temperature of the room. Therefore, consideration should be given to installing the equipment in an environment compatible with the manufacturer's maximum rated ambient temperature.

Reduced Airflow

Equipment should be mounted into a rack so that the amount of airflow required for safe operation is not compromised.

Mechanical Loading

Equipment should be mounted into a rack so that a hazardous condition does not arise due to uneven mechanical loading.

Circuit Overloading

Consideration should be given to the connection of the equipment to the power supply circuitry and the effect that any possible overloading of circuits might have on over-current protection and power supply wiring. Appropriate consideration of equipment nameplate ratings should be used when addressing this concern.

Reliable Ground

A reliable ground for the system must be maintained at all times. To ensure this, the rack itself should be grounded. Particular attention should be given to power supply connections other than the direct connections to the branch circuit (i.e. the use of power strips, etc.). Note that all power and data cables should be routed in such a way that they do not block the airflow generated by the enclosure fans.

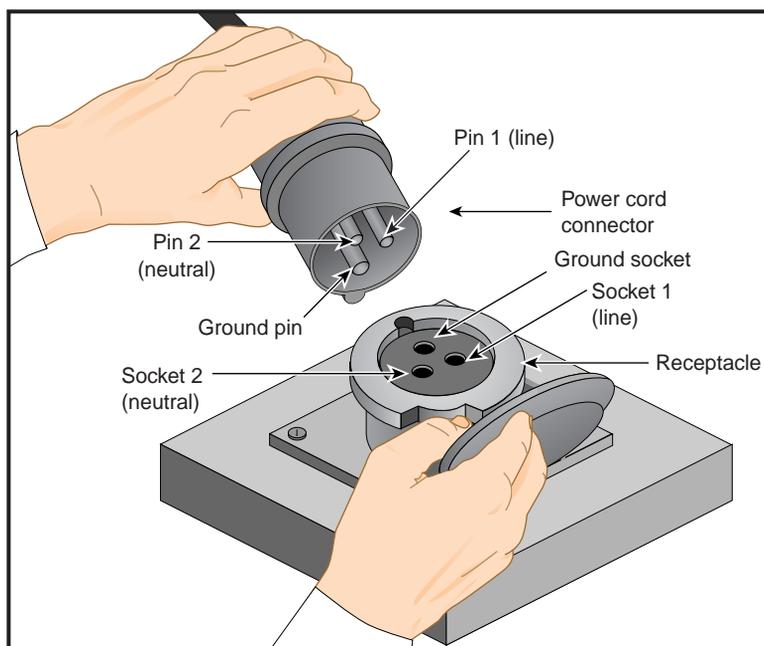


Figure 2-1 Single-Phase Power Plug Example

Providing Power

Plug the power cord from the server power supply array into a rack power distribution unit (PDU) or high-quality power source (see example in Figure 2-1) that offers protection from electrical noise and power surges.

For higher availability it is recommended that you use an optional uninterruptible power supply (UPS) with the cluster server (not provided by SGI).

Finally, press the enable power switch to (1) on the rear of the enclosure, see Figure 2-2 on page 15 for an example.

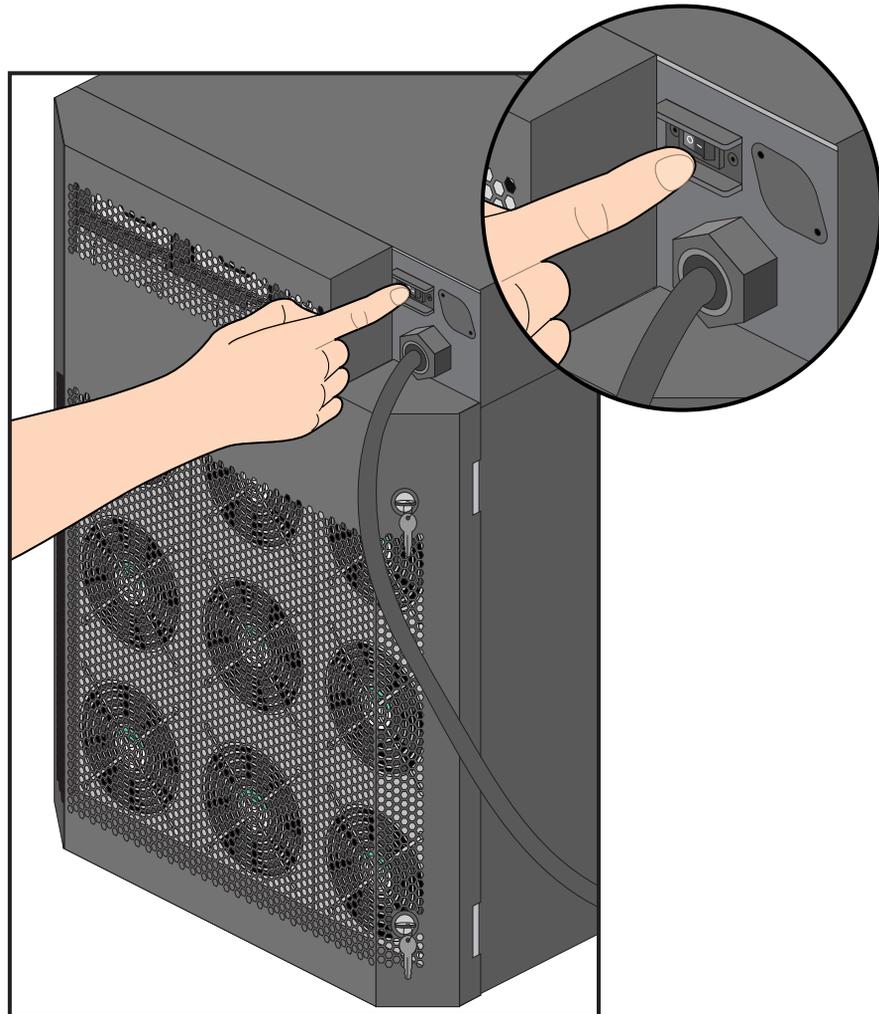


Figure 2-2 Enable Enclosure Power Switch Location

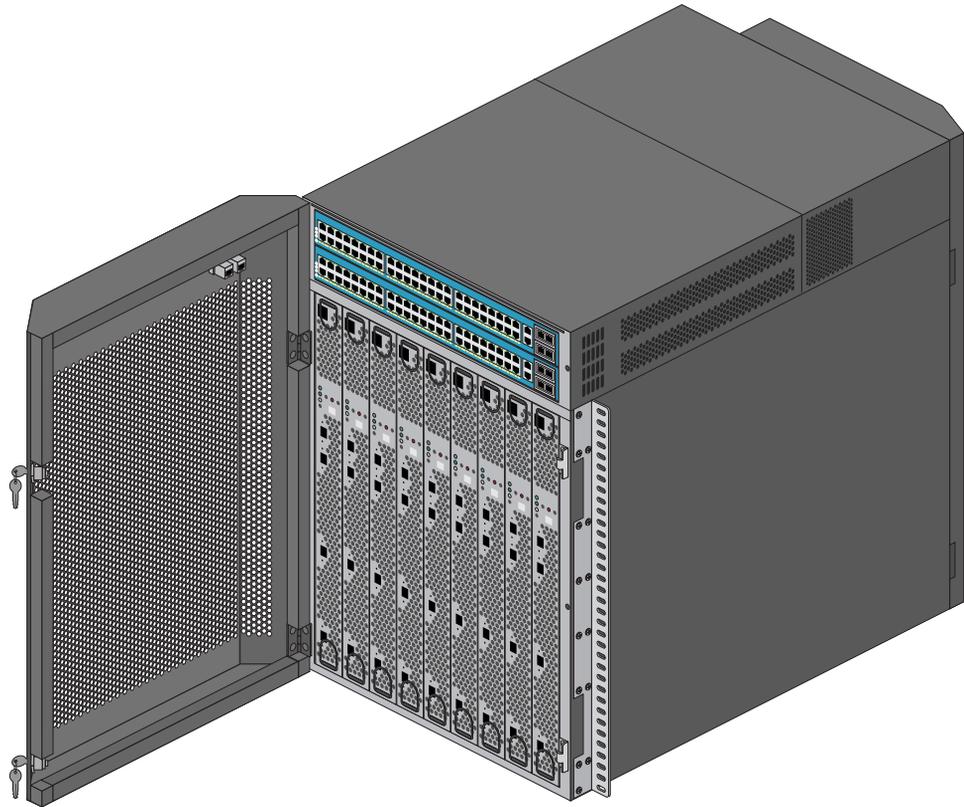


Figure 2-3 Enclosure with Installed Trays Example (Front Panel Open)

Troubleshooting the System

The following table lists recommended actions for problems that can occur. To solve problems that are not listed in this table or in another section of this chapter, contact your SGI system support engineer (SSE) or other approved service provider.

Table 2-1 Troubleshooting Chart

Problem Description	Recommended Action
The system will not power on.	<p>Ensure that the power cord of the enclosure is seated properly in the power receptacle.</p> <p>Ensure the enclosure's power supply switch is set to On (I).</p> <p>Did you push the power button on the "head node" compute tray as well as the "compute node" trays?</p> <p>If the power cord is plugged in and all the power switches are on, contact your support organization or SSE.</p>
An individual compute/memory tray will not power on.	<p>View the LED outputs on the front of the tray, (see also Figure 2-5 on page 19).</p> <p>If the LEDs are not lit, contact your SSE.</p>
The system will not boot the operating system.	Contact your SSE.
The PWR LED of a populated PCI slot in a tray is not illuminated.	Refer to the <i>SGI CloudRack X2 Component Replacement Guide</i> and reseal the PCI card.
The Fault LED of a populated PCI slot is illuminated (on).	Refer to the <i>SGI CloudRack X2 Component Replacement Guide</i> and reseal the PCI card. If the fault LED remains on, replace the PCI card.
The fault LED of a hard disk drive is on.	Refer to the <i>SGI CloudRack X2 Component Replacement Guide</i> and replace the disk drive.

Enclosure Power Supply Status LEDs

Each power supply installed in a CloudRack X2 enclosure has two (green/amber) status LEDs, see Figure 2-4 for an example. These LEDs can be viewed to determine if a problem with the supply exists. If the supply status LED indicates a malfunction, a service technician should replace it as soon as practicable. Service information is available in the *SGI CloudRack X2 Component Replacement Guide*.

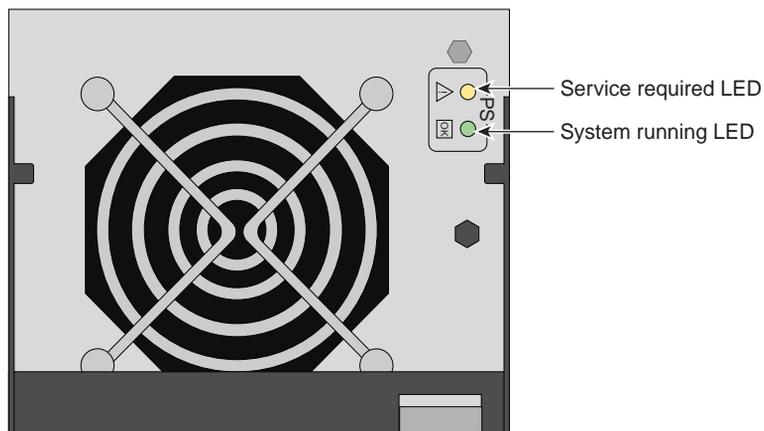


Figure 2-4 System Power Supply Status LED Locations

The LEDs will either light green or amber (yellow), or flash green or yellow to indicate the status of the individual supply. See Table 2-2 for a complete list.

Table 2-2 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 system is off	Blinking	Off
Power supply on (system on)	On	Off

Individual Tray LEDs

Each server tray installed in a CloudRack X2 enclosure has LED indicators to show the operational status of the tray. The LEDs are located on the front section of the tray and are visible when the front cover of the enclosure is open, see the example in Figure 2-5. The functions of the LED status lights on the example tray shown are as follows:

- RED- Power identifier - this red LED shows that power is being supplied to the tray.
- GREEN- Power OK - this green LED lights when the correct power levels are present on the processor(s) and other components used on the tray.
- BLUE- HDD status on the tray. This blue LED lights when functionality is established on the tray's hard disk drive or solid state disk drive.
- WHITE - this custom programmable LED is used for reporting specific tray related status that may differ depending on the build-to-order configuration you purchased.
- NUMERIC - this single-numeral LED shows the server number assigned to the tray.

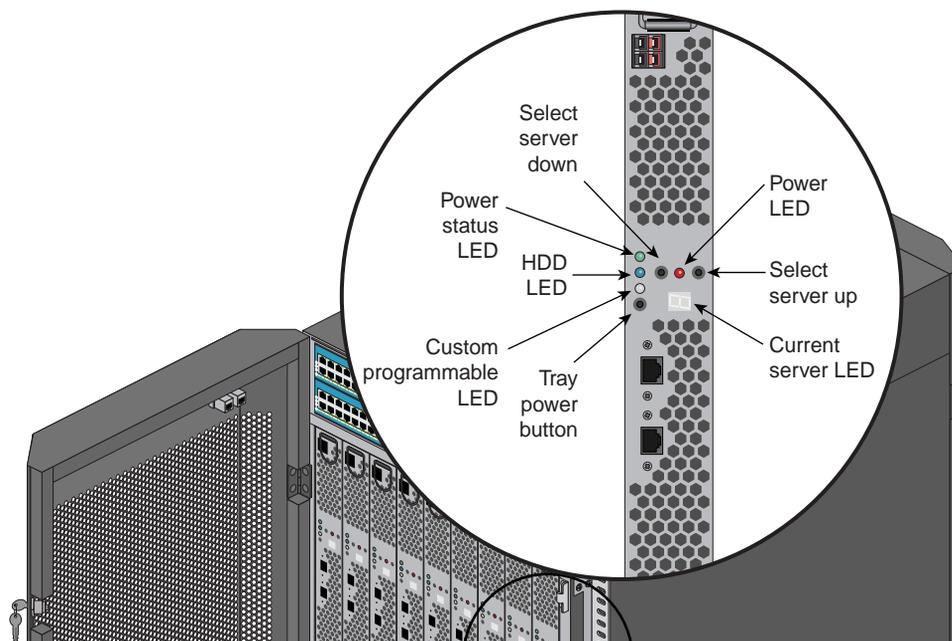


Figure 2-5 Example Compute Tray Status LEDs and Switches

System Fan Failure

The nine fans that cool the main enclosure and compute trays are arranged in three replaceable assemblies on the back of the unit. If a fan in one of the assemblies fails for any reason, a trouble indicator light comes on, see Figure 2-6. Depending on your system's configuration, you may also receive a console warning. Refer to the *SGI CloudRack X2 Component Replacement Guide* for information on replacing a fan assembly. The system can continue to run with a single fan failure, but it should be replaced as soon as possible.

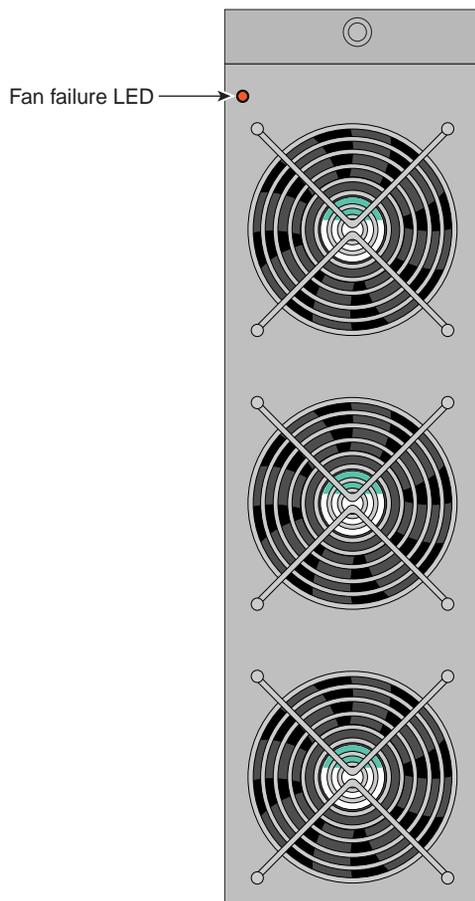


Figure 2-6 Fan Failure LED Location Example

System Interfaces Overview

This chapter provides a brief overview of the standard and optional interfaces available on your SGI CloudRack X2 system. The components of the system are described and illustrated.

NICs 1 and 2 (HPC trays only)

- **NIC1** - Indicates network activity on LAN1 when flashing.
- **NIC2** - Indicates network activity on LAN2 when flashing

HDD

Channel activity for the hard disk drives. This light indicates disk drive activity on the unit when flashing. The drive LED can be seen through the perforated tray front panel in most system configurations.

Power

The LED indicates power is being supplied to the system's power supply unit. This LED should normally be illuminated when the system is operating. Each power supply also has a trouble light indicating a malfunction in either the supply or its cooling fan, see also “Enclosure Power Supply Status LEDs” in Chapter 2.

Tray Control Panel Button Examples

Some compute trays use push-buttons located on the front panel, a power button and server selection buttons are available on specific configurations.

- **Power** - This is the main power button, which is used to apply or turn off the power to the compute tray. Pushing this button removes the main power but keeps standby power supplied to the tray, see Figure 3-1.

- **Select Server** - These buttons select the server number of the individual compute tray.

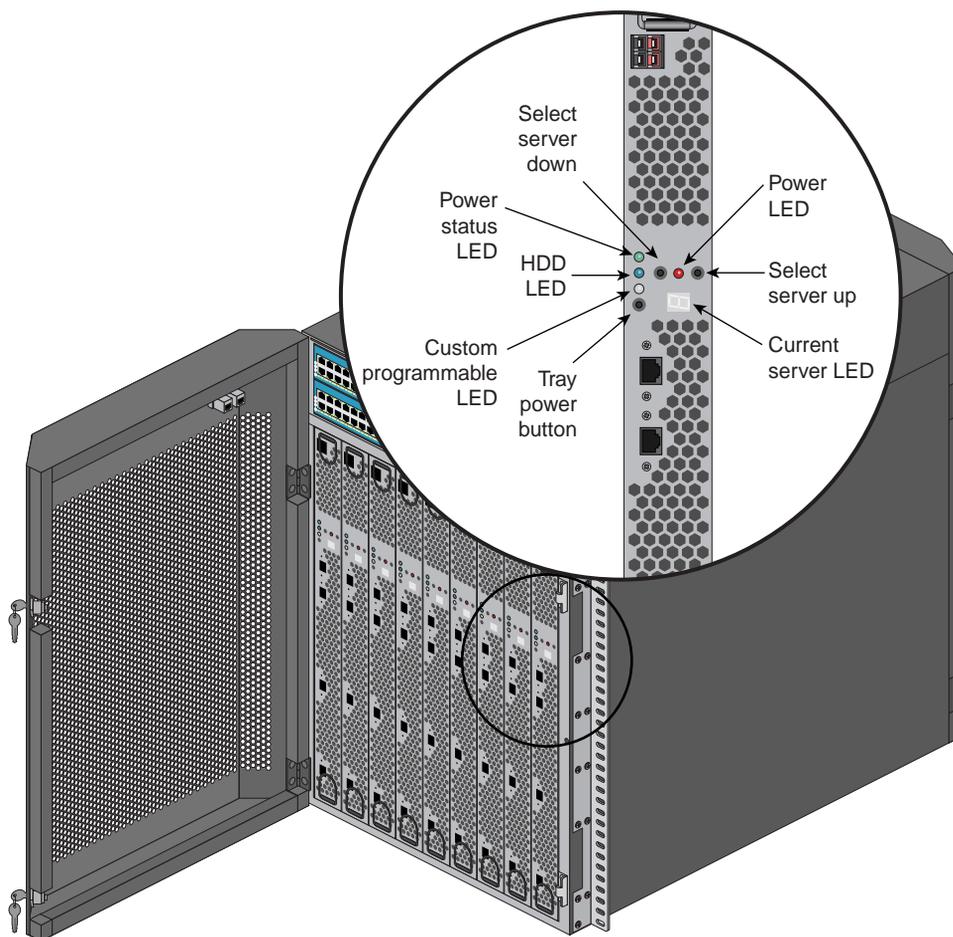


Figure 3-1 Power and Server Select Button Examples on a Compute Tray

Overheat/Fan Fail LED on Back of System

When the red LED on a rear fan assembly flashes, it indicates a fan failure. When on continuously it indicates an error condition, which may be caused by a fan failure, an obstruction of the airflow in the system or the ambient room temperature being too warm. Check the routing of the cables and make sure all fans are present and operating normally.

HPC Server BIOS Information

This chapter describes the functions and features of the AMI BIOS Setup Utility for the SGI HPC version of the CloudRack X2 cluster server. The AMI ROM BIOS is stored in a Flash EEPROM and can be updated as needed; check with your SGI sales or service representative for information on updates. This chapter covers basic navigation of the AMI BIOS Setup Utility screens.

Important: This BIOS information is applicable to Intel Xeon based cluster servers only.

Starting the BIOS Setup Utility

To enter the AMI BIOS Setup Utility screens, press the <Delete> key while the system is booting up.

Note: In most cases, the <Delete> key is used to launch the AMI BIOS setup screen. There are a few cases when other keys are used, such as <F1>, <F2>, etc.

Each main BIOS menu option is described in this manual. The Main BIOS setup menu screen has two main frames. The left frame displays all the options that can be configured. Note that grayed-out options cannot be configured. Options in blue can be configured by the user. The right frame displays the key legend. Above the key legend is an area reserved for a text message. When an option is selected in the left frame, it is highlighted in white. Often a text message will accompany it. Note that the AMI BIOS has default text messages built in. SGI retains the option to include, omit, or change any of these text messages.

The AMI BIOS Setup Utility uses a key-based navigation system called "hot keys". Most of the AMI BIOS setup utility "hot keys" can be used at any time during the setup navigation process. These keys include <F1>, <F10>, <Enter>, <ESC>, arrow keys, etc.

Note: Options printed in Bold are default settings.

How To Change the Configuration Data

The configuration data that determines the system parameters may be changed by entering the AMI BIOS Setup utility. This Setup utility can be accessed by pressing at the appropriate time during system boot.

Starting the Setup Utility

Normally, the only visible Power-On Self-Test (POST) routine is the memory test. As the memory is being tested, press the <Delete> key to enter the main menu of the AMI BIOS Setup Utility. From the main menu, you can access the other setup screens. An AMI BIOS identification string is displayed at the left bottom corner of the screen below the copyright message.

Warning: Do not upgrade the BIOS unless your system has a BIOS-related issue and you have instructions to do the upgrade from your SGI sales or service representative. Flashing the wrong BIOS can cause irreparable damage to the system and may void your warranty. Your warranty may not cover direct, indirect, special, incidental, or consequential damages arising from a BIOS update. If you have to update the BIOS, do not shut down or reset the system while the BIOS is updating. This is to avoid possible boot failure.

Main Setup Screen

When you first enter the AMI BIOS Setup Utility, you will enter the Main setup screen. You can always return to the Main setup screen by selecting the Main tab on the top of the screen. The Main BIOS Setup screen has information similar to that shown below.

System Overview:

The following BIOS information will be displayed:

System Time/System Date:

Use this option to change the system time and date. Highlight System Time or System Date using the arrow keys. Enter new values through the keyboard. Press the <Tab> key or the arrow keys to move between fields. The date must be entered in Day MM/DD/YY format. The time is entered in HH:MM:SS format. (Note that the time is in the 24-hour format. For example, 5:30 P.M. appears as 17:30:00.)

BIOS Build Version:

This item displays the BIOS revision used in your system.

BIOS Build Date:

This item displays the date when this BIOS was completed.

AMI BIOS Core Version:

This item displays the revision number of the AMI BIOS Core upon which your BIOS was built.

Processor:

The AMI BIOS will automatically display the status of the processor used in your system:

CPU Type:

This item displays the type of CPU used in the system motherboard.

Speed:

This item displays the speed of the CPU detected by the BIOS.

Physical Count:

This item displays the number of processors installed in your system as detected by the BIOS.

Logical Count:

This item displays the number of CPU **Cores** installed in your system as detected by the BIOS.

Micro_code Revision:

This item displays the revision number of the BIOS Micro_code used in your system.

System Memory:

This displays the size of memory available in the system:

Size:

This item displays the memory size detected by the BIOS.

Advanced Setup Configurations

Use the arrow keys to select Boot Setup and hit <Enter> to access the submenu items:

BOOT Features

Quick Boot

If enabled, this option will skip certain tests during POST to reduce the time needed for system boot. The options are Enabled and Disabled.

Quiet Boot

This option allows the bootup screen options to be modified between POST messages or the OEM logo. Select **Disabled** to display the POST messages. Select **Enabled** to display the OEM logo instead of the normal POST messages. The options are Enabled and Disabled.

Add-On ROM Display Mode

This sets the display mode for Option ROM. The options are **Force BIOS** and **Keep Current**.

Bootup Num-Lock

This feature selects the Power-on state for Numlock key. The options are Off and On.

Wait For 'F1' If Error

This forces the system to wait until the 'F1' key is pressed if an error occurs. The options are Disabled and Enabled.

Hit 'Del' Message Display

This feature displays "Press DEL to run Setup" during POST. The options are Enabled and Disabled.

Interrupt 19 Capture

Interrupt 19 is the software interrupt that handles the boot disk function. When this item is set to Enabled, the ROM BIOS of the host adaptors will "capture" Interrupt 19 at boot and allow the drives that are attached to these host adaptors to function as bootable disks. If this item is set to Disabled, the ROM BIOS of the host adaptors will not capture Interrupt 19, and the drives attached to these adaptors will not function as bootable devices. The options are Enabled and Disabled.

Power Configuration

Power Button Function

If set to Instant_Off, the system will power off immediately as soon as the user hits the power button. If set to 4_Second_Override, the system will power off when the user presses the power button for 4 seconds or longer. The options are Instant_Off and 4_Second_Override.

Restore on AC Power Loss

Use this feature to set the power state after a power outage. Select Power-Off for the system power to remain off after a power loss. Select Power-On for the system power to be turned on after a power loss. Select Last State to allow the system to resume its last state before a power loss. The options are Power-On, Power-Off and Last State.

Watch Dog Timer

If enabled, the Watch Dog Timer will allow the system to reboot when it is inactive for more than 5 minutes. The options are Enabled and Disabled.

Processor and Clock Options

This submenu allows the user to configure the Processor and Clock settings.

Ratio CMOS Setting

This option allows the user to set the ratio between the CPU Core Clock and the FSB Frequency. (Note: if an invalid ratio is entered, the AMI BIOS will restore the setting to the previous state.) The default setting depends on the type of CPU installed on the motherboard. The default setting

for the CPU installed in your motherboard is [18]. Press "+" or "-" on your keyboard to change this value.

C1E Support

Select Enabled to use the feature of Enhanced Halt State. C1E significantly reduces the CPU's power consumption by reducing the CPU's clock cycle and voltage during a "Halt State." The options are Disabled and Enabled.

Hardware Prefetcher (Available when supported by the CPU)

If set to Enabled, the hardware pre-fetcher will pre-fetch streams of data and instructions from the main memory to the L2 cache in a forward or backward manner to improve CPU performance. The options are Disabled and Enabled.

Adjacent Cache Line Prefetch (Available when supported by the CPU)

The CPU fetches the cache line for 64 bytes if this option is set to Disabled. The CPU fetches both cache lines for 128 bytes as comprised if Enabled.

Intel® Virtualization Technology (Available when supported by the CPU)

Select Enabled to use the feature of Virtualization Technology to allow one platform to run multiple operating systems and applications in independent partitions, creating multiple "virtual" systems in one physical computer. The options are Enabled and Disabled.

Note: Check with your SGI sales or support representative for information before trying to use the Virtualization option. If there is any change to this setting, you will need to power off and restart the system for the change to take effect.

Execute-Disable Bit Capability (Available when supported by the OS and the CPU)

Set to Enabled to enable the Execute Disable Bit which will allow the processor to designate areas in the system memory where an application code can execute and where it cannot, thus preventing a worm or a virus from flooding illegal codes to overwhelm the processor or damage the system during an attack. The default is Enabled. (Check with your SGI sales or service representative for more information before modifying this setting.)

Simultaneous Multi-Threading (Available when supported by the CPU)

Set to Enabled to use the Simultaneous Multi-Threading Technology, which will result in increased CPU performance. The options are Disabled and Enabled.

Active Processor Cores

Set to Enabled to use a processor's Second Core and beyond. (Please refer to Intel's web site for more information.) The options are All, 1 and 2.

Intel® EIST Technology

EIST (Enhanced Intel SpeedStep Technology) allows the system to automatically adjust processor voltage and core frequency in an effort to reduce power consumption and heat dissipation. Check with your SGI sales or service representative for more information on using this option in SGI systems and clusters. The options are Disable (Disable GV3) and Enable (Enable GV3).

Intel® TurboMode Technology

Select Enabled to use the Turbo Mode to boost system performance. The options are Enabled and Disabled.

Intel® C-STATE Tech

If enabled, C-State is set by the system automatically to either C2, C3 or C4 state. The options are Disabled and Enabled.

C-State package limit setting

If set to Auto, the AMI BIOS will automatically set the limit on the C-State package register. The options are Auto, C1, C3, C6 and C7.

C1 Auto Demotion

When enabled, the CPU will conditionally demote C3, C6 or C7 requests to C1 based on un-core auto-demote information. The options are Disabled and Enabled.

C3 Auto Demotion

When enabled, the CPU will conditionally demote C6 or C7 requests to C3 based on un-core auto-demote information. The options are Disabled and Enabled.

Clock Spread Spectrum

Select Enable to use the feature of Clock Spectrum, which will allow the BIOS to monitor and attempt to reduce the level of Electromagnetic Interference caused by the components whenever needed. The options are Disabled and Enabled.

Advanced Chipset Control

The items included in the Advanced Settings submenu are listed below:

CPU Bridge Configuration QPI Links Speed

This feature selects QPI's data transfer speed. The options are Slow-mode, and Full Speed.

QPI Frequency

This selects the desired QPI frequency. The options are Auto, 4.800 GT, 5.866GT, 6.400 GT.

QPI L0s and L1

This enables the QPI power state to low power. L0s and L1 are automatically selected by the motherboard. The options are Disabled and Enabled.

Memory Frequency

This feature forces a DDR3 frequency slower than what the system has detected. The available options are Auto, Force DDR-800, Force DDR-1066, Force DDR-1333.

Memory Mode

The options are Independent, Channel Mirror, Lockstep and Sparing.

- Independent - All DIMMs are available to the operating system.
- Channel Mirror - The motherboard maintains two identical (redundant) copies of all data in memory.
- Lockstep - The motherboard uses two areas of memory to run the same set of operations in parallel.
- Sparing - A preset threshold of correctable errors is used to trigger fail-over. The spare memory is put online and used as active memory in place of the failed memory.

Demand Scrubbing

A memory error-correction scheme where the Processor writes corrected data back into the memory block from where it was read by the Processor. The options are Enabled and Disabled.

Patrol Scrubbing

A memory error-correction scheme that works in the background looking for and correcting resident errors. The options are Enabled and Disabled.

Throttling - Closed Loop/Throttling - Open Loop

Throttling improves reliability and reduces power in the processor by automatic voltage control during processor idle states. Available options are Disabled and Enabled. If Enabled, the following items will appear:

Hysteresis Temperature (For the Closed Loop only)

Temperature Hysteresis is the temperature lag (in degrees Celsius) after the set DIMM temperature threshold is reached before Closed Loop Throttling begins. The options are Disabled, 1.5° C, 3.0° C, and 6.0° C.

Guardband Temperature (For the Closed Loop only)

This is the temperature which applies to the DIMM temperature threshold. Each step is in 0.5°C increment. The default is [006]. Press "+" or "-" on your keyboard to change this value.

Inlet Temperature

This is the temperature detected at the chassis inlet. Each step is in 0.5°C increment. The default is [070]. Press "+" or "-" on your keyboard to change this value.

Temperature Rise

This is the temperature rise to the DIMM thermal zone. Each step is in 0.5°C increment. The default is [020]. Press "+" or "-" on your keyboard to change this value.

Air Flow

This is the airflow speed to the DIMM modules. Each step is one mm/sec. The default is [1500]. Press "+" or "-" on your keyboard to change this value.

Altitude

This feature defines how many meters above or below sea level the system is located. The options are Sea Level or Below, 1~300, 301~600, 601~900, 901~1200, 1201~1500, 1501~1800, 1801~2100, 2101~2400, 2401~2700, 2701~3000.

DIMM Pitch

This is the physical space between each DIMM module. Each step is in 1/1000 of an inch. The default is [400]. Press "+" or "-" on your keyboard to change this value.

North Bridge Configuration

This feature allows the user to configure the settings for the Intel North Bridge chip.

Crystal Beach/DMA

This feature works with the Intel I/O AT (Acceleration Technology) to accelerate the performance of TOE devices. (Note: A TOE device is a specialized, dedicated processor that is installed on an add-on card or a network card to handle some or all packet processing of this add-on card.) When this feature is set to Enabled, it will enhance overall system performance by providing direct memory access for data transferring. The options are Enabled and Disabled. Check with your SGI sales or service representative for information on the availability of this option.

Intel VT-d

Select Enabled to enable Intel's Virtualization Technology support for Direct I/O VT-d by reporting the I/O device assignments to VMM through the DMAR ACPI Tables. This feature offers fully-protected I/O resource-sharing across the Intel platforms, providing the user with greater reliability, security and availability in networking and data-sharing. The settings are Enabled and Disabled.

IOH PCIE Port1 Bifurcation

This feature allows the user to set IOH Bifurcation configuration for the PCI-E Port

The options are X4X4X4X4, X4X4X8, X8X4X4, X8X8.

IOH PCIE Max Payload Size

Some add-on cards perform faster with the coalesce feature, which limits the payload size to 128 MB; while others, with a payload size of 256 MB which inhibits the coalesce feature. Please refer to your add-on card user guide for the desired setting. The options are 256 MB and 128MB.

SouthBridge Configuration

This feature allows the user to configure the settings for the Intel ICH South Bridge chipset.

USB Functions

This feature allows the user to decide the number of on-board USB ports to be enabled. The Options are: Disabled, 2 USB ports, 4 USB ports, 6 USB ports, 8 Ports, 10 Ports and 12 USB ports.

Legacy USB Support

Select Enabled to use Legacy USB devices. If this item is set to Auto, Legacy USB support will be automatically enabled if a legacy USB device is installed on the motherboard, and vice versa. The settings are Disabled, and Enabled.

USB 2.0 Controller

Select Enabled to activate the on-board USB 2.0 controller. The options are Enabled and Disabled.

USB 2.0 Controller Mode

This setting allows you to select the USB 2.0 Controller mode. The options are Hi-Speed (480 Mbps) and Full Speed (12 Mbps).

BIOS EHCI Hand-Off

Select Enabled to enable BIOS Enhanced Host Controller Interface support to provide a workaround solution for an operating system that does not have EHCI Hand-Off support. When enabled, the EHCI Interface will be changed from the BIOS-controlled to the OS-controlled. The options are Disabled and Enabled.

XIDE/SATA Configuration

When this submenu is selected, the AMI BIOS automatically detects the presence of the IDE devices and displays the following items:

- SATA#1 Configuration

If Compatible is selected, it sets SATA#1 to legacy compatibility mode, while selecting Enhanced sets SATA#1 to native SATA mode. The options are Disabled, Compatible and Enhanced.

- Configure SATA#1 as

This feature allows the user to select the drive type for SATA#1. The options are IDE, RAID and AHCI.

- SATA#2 Configuration

Selecting Enhanced will set SATA#2 to native SATA mode. The options are Disabled, and Enhanced.

Primary IDE Master/Slave, Secondary IDE Master/Slave, Third IDE Master, and Fourth IDE Master

These settings allow the user to set the parameters of Primary IDE Master/Slave, Secondary IDE Master/Slave, Third and Fourth IDE Master slots. Hit <Enter> to activate the following submenu screen for detailed options of these items. Set the correct configurations accordingly. The items included in the submenu are:

- Type

Select the type of device connected to the system. The options are Not Installed, Auto, CD/DVD and ARMD.

- LBA/Large Mode

LBA (Logical Block Addressing) is a method of addressing data on a disk drive. In the LBA mode, the maximum drive capacity is 137 GB. For drive capacities over 137 GB, your system must be equipped with a 48-bit LBA mode addressing. If not, contact your manufacturer or install an ATA/133 IDE controller card that supports 48-bit LBA mode. The options are Disabled and Auto.

- Block (Multi-Sector Transfer)

Block Mode boosts the IDE drive performance by increasing the amount of data transferred. Only 512 bytes of data can be transferred per interrupt if Block Mode is not used. Block Mode allows transfers of up to 64 KB per interrupt. Select Disabled to allow data to be transferred from and to the device one sector at a time. Select Auto to allow data transfer from and to the device occur multiple sectors at a time if the device supports it. The options are Auto and Disabled.

- PIO Mode

The IDE PIO (Programmable I/O) Mode programs timing cycles between the IDE drive and the programmable IDE controller. As the PIO mode increases, the cycle time decreases. The options are Auto, 0, 1, 2, 3, and 4.

Select Auto to allow the AMI BIOS to automatically detect the PIO mode. Use this value if the IDE disk drive support cannot be determined. Select 0 to allow the AMI BIOS to use PIO mode 0. It has a data transfer rate of 3.3 MBs. Select 1 to allow the AMI BIOS to use PIO mode 1. It has a data transfer rate of 5.2 MBs.

Select 2 to allow the AMI BIOS to use PIO mode 2. It has a data transfer rate of 8.3 MBs. Select 3 to allow the AMI BIOS to use PIO mode 3. It has a data transfer rate of 11.1 MBs. Select 4 to allow the AMI BIOS to use PIO mode 4. It has a data transfer bandwidth of 32-Bits. Select Enabled to enable 32-Bit data transfer.

- DMA Mode

Select Auto to allow the BIOS to automatically detect IDE DMA mode when the IDE disk drive support cannot be determined.

Select SWDMA0 to allow the BIOS to use Single Word DMA mode 0. It has a data transfer rate of 2.1 MBs.

Select SWDMA1 to allow the BIOS to use Single Word DMA mode 1. It has a data transfer rate of 4.2 MBs.

Select SWDMA2 to allow the BIOS to use Single Word DMA mode 2. It has a data transfer rate of 8.3 MBs.

Select MWDMA0 to allow the BIOS to use Multi Word DMA mode 0. It has a data transfer rate of 4.2 MBs.

Select MWDMA1 to allow the BIOS to use Multi Word DMA mode 1. It has a data transfer rate of 13.3 MBs.

Select MWDMA2 to allow the BIOS to use Multi-Word DMA mode 2. It has a data transfer rate of 16.6 MBs.

Select UDMA0 to allow the BIOS to use Ultra DMA mode 0. It has a data transfer rate of 16.6 MBs. It has the same transfer rate as PIO mode 4 and Multi Word DMA mode 2.

Select UDMA1 to allow the BIOS to use Ultra DMA mode 1. It has a data transfer rate of 25 MBs.

Select UDMA2 to allow the BIOS to use Ultra DMA mode 2. It has a data transfer rate of 33.3 MBs.

Select UDMA3 to allow the BIOS to use Ultra DMA mode 3. It has a data transfer rate of 66.6 MBs.

Select UDMA4 to allow the BIOS to use Ultra DMA mode 4. It has a data transfer rate of 100 MBs. The options are Auto, SWDMAn, MWDMAAn, and UDMAAn.

S.M.A.R.T. For Hard disk drives

Self-Monitoring Analysis and Reporting Technology (SMART) can help predict impending drive failures. Select Auto to allow the AMI BIOS to automatically detect hard disk drive support. Select Disabled to prevent the AMI BIOS from using the S.M.A.R.T. Select Enabled to allow the AMI BIOS to use the S.M.A.R.T. to support hard drive disk. The options are Disabled, Enabled, and Auto.

32Bit Data Transfer

Select Enable to enable the function of 32-bit IDE data transfer. The options are Enabled and Disabled.

IDE Detect Timeout (sec)

Use this feature to set the time-out value for the BIOS to detect the ATA, ATAPI devices installed in the system. The options are 0 (sec), 5, 10, 15, 20, 25, 30, and 35.

Clear NVRAM

This feature clears the NVRAM during system boot. The options are No and Yes.

Plug & Play OS

Selecting Yes allows the OS to configure Plug & Play devices. (This is not required for system boot if your system has an OS that supports Plug & Play.) Select No to allow the AMI BIOS to configure all devices in the system.

PCI Latency Timer

This feature sets the latency Timer of each PCI device installed on a PCI bus. Select 64 to set the PCI latency to 64 PCI clock cycles. The options are 32, 64, 96, 128, 160, 192, 224 and 248.

PCI IDE BusMaster

When enabled, the BIOS uses PCI bus mastering for reading/writing to IDE drives. The options are Disabled and Enabled.

Load Onboard LAN1 Option ROM/Load Onboard LAN2 Option ROM

Select Enabled to enable the onboard LAN1 or LAN2 Option ROM. This is to boot the computer using a network interface. The options are Enabled and Disabled.

Serial Port1 Address/ Serial Port2 Address

This option specifies the base I/O port address and the Interrupt Request address of Serial Port 1 and Serial Port 2. Select Disabled to prevent the serial port from accessing any system resources. When this option is set to Disabled, the serial port physically becomes unavailable. Select 3F8/IRQ4 to allow the serial port to use 3F8 as its I/O port address and IRQ 4 for the interrupt address. The options for Serial Port1 are Disabled, 3F8/IRQ4, 3E8/IRQ4, 2E8/IRQ3. The options for Serial Port2 are Disabled, 2F8/IRQ3, 3E8/IRQ4, and 2E8/IRQ3.

Remote Access Configuration

Remote Access

This allows the user to enable the Remote Access feature. The options are Disabled and Enabled. If Remote Access is set to Enabled, the following items will display:

- Serial Port Number

This feature allows the user decide which serial port to be used for Console Redirection. The options are COM 1 and COM 2.

- Serial Port Mode

This feature allows the user to set the serial port mode for Console Redirection. The options are 115200 8, n 1; 57600 8, n, 1; 38400 8, n, 1; 19200 8, n, 1; and 9600 8, n, 1.

- Flow Control

This feature allows the user to set the flow control for Console Redirection. The options are None, Hardware, and Software.

Redirection After BIOS POST

Select Disabled to turn off Console Redirection after Power-On Self-Test (POST). Select Always to keep Console Redirection active all the time after POST.

Note: This setting may not be supported by some operating systems.

Select Boot Loader to keep Console Redirection active during POST and Boot Loader. The options are Disabled, Boot Loader, and Always.

Terminal Type

This feature allows the user to select the target terminal type for Console Redirection. The options are ANSI, VT100, and VT-UTF8.

VT-UTF8 Combo Key Support

A terminal keyboard definition that provides a way to send commands from a remote console. Available options are Enabled and Disabled.

Sredir Memory Display Delay

This feature defines the length of time in seconds to display memory information. The options are No Delay, Delay 1 Sec, Delay 2 Sec, and Delay 4 Sec.

Hardware Health Monitor

This feature allows the user to monitor system health and review the status of each item as displayed.

CPU Overheat Alarm

This option allows the user to select the CPU Overheat Alarm setting which determines when the CPU OH alarm will be activated to provide warning of possible CPU overheat.

Warning: Any temperature that exceeds the CPU threshold temperature predefined by the CPU manufacturer may result in CPU overheat or system instability. When the CPU temperature reaches this predefined threshold, the CPU and system cooling fans will run at full speed.

The options are:

- The Early Alarm: Select this setting if you want the CPU overheat alarm (including the LED and the buzzer) to be triggered as soon as the CPU temperature reaches the CPU overheat threshold as predefined by the CPU manufacturer.
- The Default Alarm: Select this setting if you want the CPU overheat alarm (including the LED and the buzzer) to be triggered when the CPU temperature reaches about 50C above the threshold temperature as predefined by the CPU manufacturer to give the CPU and system fans additional time needed for CPU and system cooling. In both the alarms above, please take immediate action as shown below. (See the notes on P. 4-18 for more information.)

CPU Temperature/System Temperature

This feature displays current temperature readings for the CPU and the System.

The following items will be displayed for your reference only:

CPU Temperature

The CPU thermal technology that reports absolute temperatures (Celsius/Fahrenheit) has been upgraded to a more advanced feature by Intel in its newer processors. The basic concept is each CPU is embedded by unique temperature information that the motherboard can read. This 'Temperature Threshold' or 'Temperature Tolerance' has been assigned at the factory and is the baseline on which the motherboard takes action during different CPU temperature conditions (i.e., by increasing CPU Fan speed, triggering the Overheat Alarm, etc.) Since CPUs can have different 'Temperature Tolerances', the installed CPU can now send information to the motherboard regarding what its 'Temperature Tolerance' is, and not the other way around. This results in better CPU thermal management.

The manufacturer has leveraged this feature by assigning a temperature status to certain thermal conditions in the processor (Low, Medium and High). This makes it easier for the user to understand the CPU's temperature status, rather than by just simply seeing a temperature reading (i.e., 25°C). The CPU Temperature feature will display the CPU temperature status as detected by the BIOS:

Low – This level is considered as the 'normal' operating state. The CPU temperature is well below the CPU 'Temperature Tolerance'. The motherboard fans and CPU will run normally as configured in the BIOS (Fan Speed Control).

User intervention: No action required.

Medium – The processor is running warmer. This is a ‘precautionary’ level and generally means that there may be factors contributing to this condition, but the CPU is still within its normal operating state and below the CPU ‘Temperature Tolerance’. The motherboard fans and CPU will run normally as configured in the BIOS. The fans may adjust to a faster speed depending on the Fan Speed Control settings.

User intervention: No action is required. However, consider checking the CPU fans and the chassis ventilation for blockage.

High – The processor is running hot. This is a ‘caution’ level since the CPU’s ‘Temperature Tolerance’ has been reached (or has been exceeded) and may activate an overheat alarm.

User intervention: If the system buzzer and Overheat LED has activated, take action immediately by checking the system fans, chassis ventilation and room temperature to correct any problems.

Note: The system may shut down if it continues for a long period to prevent damage to the CPU. The information provided above is for your reference only. For more information on processor thermal management, reference Intel’s Web site at www.Intel.com or contact your support representative.

System Temperature: The system temperature will be displayed (in degrees in Celsius and Fahrenheit) as it is detected by the BIOS.

Fan Speed Control Monitor

This feature allows the user to decide how the system controls the speeds of the on-board fans. The CPU temperature and the fan speed are correlative. When the CPU on-die temperature increases, the fan speed will also increase, and vice versa. Select Workstation if your system is used as a Workstation. Select Server if your system is used as a Server. Select “Disabled, (Full Speed @12V)” to disable the fan speed control function and allow the on-board fans to constantly run at the full speed (12V). The Options are: 1. Disabled (Full Speed), 2. Server Mode, 3. Workstation Mode.

Fan1 ~ Fan 4 Reading

This feature displays the fan speed readings from fan interfaces Fan1 through Fan5.

CPU1 Vcore, CPU2 Vcore, +5Vin, +12Vcc (V), VPI DIMM, VP2 DIMM, 3.3Vcc (V), and Battery Voltage

ACPI Configuration

Use this feature to configure Advanced Configuration and Power Interface (ACPI) power management settings for your system.

ACPI Version Features

The options are ACPI v1.0, ACPI v2.0 and ACPI v3.0. Please refer to ACPI's website for further explanation: <http://www.acpi.info/>.

ACPI APIC Support

Select Enabled to include the ACPI APIC Table Pointer in the RSDT pointer list. The options are Enabled and Disabled.

APIC ACPI SCI IRQ

When this item is set to Enabled, APIC ACPI SCI IRQ is supported by the system. The options are Enabled and Disabled.

USB Device Wakeup from S3/S4

Select to Enabled to allow USB devices to wakeup from S3/S4 state. The options are Enabled and Disabled.

High Performance Event Timer

Select Enabled to activate the High Performance Event Timer (HPET) that produces periodic interrupts at a much higher frequency than a Real-time Clock (RTC) does in synchronizing multimedia streams, providing smooth playback and reducing the dependency on other timestamp calculation devices, such as an x86 RDTSC instruction embedded in the CPU. The High Performance Event Timer is used to replace the 8254 Programmable Interval Timer. The options are Enabled and Disabled.

IPMI Configuration

Intelligent Platform Management Interface (IPMI) is a set of common interfaces that IT administrators can use to monitor system health and to manage the system as a whole. For more information on the IPMI specifications, please visit Intel's website at www.intel.com.

Status of BMC

Baseboard Management Controller (BMC) manages the interface between system management software and platform hardware. This is an informational feature which returns the status code of the BMC micro controller.

View BMC System Event Log

This feature displays the BMC System Event Log (SEL). It shows the total number of entries of BMC System Events. To view an event, select an Entry Number and press <Enter> to display the information as shown in the example below:

- **Total Number of Entries**

SEL Entry Number

SEL Record ID

SEL Record Type

Timestamp, Generator ID

Event Message Format User

Event Sensor Type

Event Sensor Number,

Event Dir Type

Event Data.

Clear BMC System Event Log

This feature is used to clear the BMC System Event Log. Caution: Any cleared information is unrecoverable. Make absolutely sure that you no longer need any data stored in the log before clearing the BMC Event Log.

Set LAN Configuration

Set this feature to configure the IPMI LAN adapter with a network address.

Channel Number - Enter the channel number for the SET LAN Configuration command. This is initially set to [1]. Press "+" or "-" on your keyboard to change the Channel Number.

Channel Number Status -This feature returns the channel status for the Channel Number selected above: "Channel Number is OK" or "Wrong Channel Number".

IP Address Configuration

Enter the IP address for this machine. This should be in decimal and in dotted quad form (i.e., 192.168.10.253). The value of each three-digit number separated by dots should not exceed 255.

Parameter Selector

Use this feature to select the parameter of your IP Address configuration.

IP Address

The BIOS will automatically enter the IP address of this machine; however it may be over-ridden. IP addresses are 6 two-digit hexadecimal numbers (Base 16, 0 ~ 9, A, B, C, D, E, F) separated by dots. (i.e., 00.30.48.D0.D4.60).

Current IP Address in BMC

This item displays the current IP address used for your IPMI connection.

MAC Address Configuration

Enter the Mac address for this machine. This should be in decimal and in dotted quad form (i.e., 192.168.10.253). The value of each three-digit number separated by dots should not exceed 255.

Parameter Selector

Use this feature to select the parameter of your Mac Address configuration.

Mac Address

The BIOS will automatically enter the Mac address of this machine; however it may be over-ridden. Mac addresses are 6 two-digit hexadecimal numbers (Base 16, 0 ~ 9, A, B, C, D, E, F) separated by dots. (i.e., 00.30.48.D0.D4.60).

Current Mac Address in BMC

This item displays the current Mac address used for your IPMI connection.

Subnet Mask Configuration

Subnet masks tell the network which subnet this machine belongs to. The value of each three-digit number separated by dots should not exceed 255.

Parameter Selector

Use this feature to select the parameter of your Subnet Masks configuration.

Subnet Masks

This item displays the current subnet masks setting for your IPMI connection.

SEL PEF Configuration

Set PEF Configuration

Set this feature to configure the Platform Event Filter (PEF). PEF interprets BMC events and performs actions based on pre-determined settings or 'traps' under IPMI 1.5 specifications. For example, powering the system down or sending an alert when a triggering event is detected.

The following will appear if PEF Support is set to Enabled. The default is Disabled.

PEF Action Global Control -These are the different actions based on BMC events. The options are Alert, Power Down, Reset System, Power Cycle, OEM Action, Diagnostic Interface.

Alert Startup Delay - This feature inserts a delay during startup for PEF alerts.

The options are Enabled and Disabled. PEF Alert Startup Delay -This sets the pre-determined time to delay PEF alerts after system power-ups and resets. Refer to Table 24.6 of the IPMI 1.5 Specification for more information at www.intel.com. The options are:

No Delay, 30 sec, 60 sec, 1.5 min, 2.0 min.

Startup Delay - This feature enables or disables startup delay. The options are Enabled and Disabled.

PEF Startup Delay -This sets the pre-determined time to delay PEF after system power-ups and resets. Refer to Table 24.6 of the IPMI 1.5 Specification for more information at www.intel.com. The options are No Delay, 30 sec, 60 sec, 1.5 min, 2.0 min.

Event Message for PEF Action - This enables or disables Event Messages for PEF action. Refer to Table 24.6 of the IPMI 1.5 Specification for more information at www.intel.com. The options are Disabled and Enabled.

BMC Watch Dog Timer Action

Allows the BMC to reset or power down the system if the operating system hangs or crashes. The options are Disabled, Reset System, Power Down, Power Cycle.

BMC Watch Dog TimeOut [Min:Sec]

This option appears if BMC Watch Dog Timer Action (above) is enabled. This is a timed delay in minutes or seconds, before a system power down or reset after an operating system failure is detected. The options are [5 Min], [1 Min], [30 Sec], and [10 Sec].

The DMI Event Log

View Event Log

Use this option to view the System Event Log.

Mark all events as read

This option marks all events as read. The options are OK and Cancel.

Clear event log

This option clears the Event Log memory of all messages. The options are OK and Cancel.

Security Settings

The AMI BIOS provides a Supervisor and a User password. If you use both passwords, the Supervisor password must be set first.

Supervisor Password

This item indicates if a supervisor password has been entered for the system. Clear means such a password has not been used and Set means a supervisor password has been entered for the system.

User Password:

This item indicates if a user password has been entered for the system. Clear means such a password has not been used and Set means a user password has been entered for the system.

Change Supervisor Password

Select this feature and press <Enter> to access the submenu, and then type in a new Supervisor Password.

User Access Level (Available when Supervisor Password is set as above)

Available options are Full Access: grants full User read and write access to the Setup Utility, View Only: allows access to the Setup Utility but the fields cannot be changed, Limited: allows only limited fields to be changed such as Date and Time, No Access: prevents User access to the Setup Utility.

Change User Password

Select this feature and press <Enter> to access the submenu, and then type in a new User Password.

Clear User Password (Available only if User Password has been set)

This item allows you to clear a user password after it has been entered.

Password Check

This item allows you to check a password after it has been entered. The options are Setup and Always.

Boot Sector Virus Protection

When Enabled, the AMI BOIS displays a warning when any program (or virus) issues a Disk Format command or attempts to write to the boot sector of the hard disk drive. The options are Enabled and Disabled.

Boot Configuration

Use this feature to configure boot settings.

Boot Device Priority

This feature allows the user to specify the sequence of priority for the Boot Device. The settings are 1st boot device, 2nd boot device, 3rd boot device, 4th boot device, 5th boot device and Disabled.

1st Boot Device - [USB: XXXXXXXXXXXX]

2nd Boot Device - [CD/DVD: XXXXXXXXXXXX]

Hard Disk Drives

This feature allows the user to specify the boot sequence from all available hard disk drives. The settings are Disabled and a list of all hard disk drives that have been detected (i.e., 1st Drive, 2nd Drive, 3rd Drive, etc.)

- 1st Drive - [SATA: XXXXXXXXXXXX]

Removable Drives

This feature allows the user to specify the boot sequence from available Removable Drives. The settings are 1st boot device, 2nd boot device, and Disabled.

1st Drive - [USB: XXXXXXXXXXXX]

2nd Drive

XCD/DVD Drives

This feature allows the user to specify the boot sequence from available CD/DVD Drives (i.e., 1st Drive, 2nd Drive, etc.)

Exit Options

Select the Exit tab from the AMI BIOS Setup Utility screen to enter the Exit BIOS Setup screen.

Save Changes and Exit

When you have completed the system con.guration changes, select this option to leave the BIOS Setup Utility and reboot the computer, so the new system configuration parameters can take effect. Select Save Changes and Exit from the Exit menu and press <Enter>.

Discard Changes and Exit

Select this option to quit the BIOS Setup without making any permanent changes to the system configuration, and reboot the computer. Select Discard Changes and Exit from the Exit menu and press <Enter>.

Discard Changes

Select this option and press <Enter> to discard all the changes and return to the AMI BIOS Utility Program.

Load Optimal Defaults

To set this feature, select Load Optimal Defaults from the Exit menu and press <Enter>. Then, select OK to allow the AMI BIOS to automatically load Optimal Defaults to the BIOS Settings. The Optimal settings are designed for maximum system performance, but may not work best for all computer applications.

Load Fail-Safe Defaults

To set this feature, select Load Fail-Safe Defaults from the Exit menu and press <Enter>. The Fail-Safe settings are designed for maximum system stability, but not for maximum performance.

BIOS Error Beep Codes

During the POST (Power-On Self-Test) routines, which are performed each time the system is powered on, errors may occur.

Non-fatal errors are those which, in most cases, allow the system to continue the boot-up process. The error messages normally appear on the screen.

Fatal errors are those which will not allow the system to continue the boot-up procedure. If a fatal error occurs, you should consult with your system manufacturer for possible repairs.

BIOS Error Beep Code List

The following list of error codes may be helpful in diagnosing certain system problems.

BIOS Error Beep Codes		
Beep Code	Error Message	Description
1 beep	Refresh	Circuits have been reset. (Ready to power up)
5 short beeps + 1 long beep	Memory error	No memory detected in the system
8 beeps 1 continuous beep (with the front panel OH LED on)	Display memory read/write error System Overheat	Video adapter missing or with faulty memory 1 continuous beep with the front panel OH LED on

Technical Specifications

This appendix contains technical specification information about your system.

Server Specifications and Features

Table A-1 shows the physical specifications of the SGI CloudRack X2 server system.

Table A-1 SGI CloudRack X2 Enclosure Physical Specifications

System Features	Specification
Height	24.4 in. (62 cm) [27.2 in. (69 cm) when mounted on casters]
Width	17.6 in. (44.7 cm)
Depth	41 in. (104 cm)
Weight (full) maximum	~518 lbs. (235.5 kg) approximate (shipping weight will be higher)
Voltage range	200-240 VAC (180-264 VAC tolerance range)
Cycles per second	50 or 60 Hz (single-phase AC)
System Cooling	Nine (9) system fans (+3 dedicated power supply fans)
Phase required	Single-phase
Power supply	Three 2900 Watt supplies
Hard drive bays	Two per (HPC) compute tray - up to six with other tray configurations
Power cable	Single pluggable cord
PCIe slots	Optional low-profile (x16) PCI-Express slots on HPC compute trays
HPC cluster system server board dimensions	6.5 in. (165 mm) x 16.4 in (417 mm). Two serverboards per compute tray

Environmental Specifications

Table A-2 lists the environmental specifications of the system.

Table A-2 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)
Cooling requirement	Ambient air or optional water cooling
Maximum altitude	10,000 ft. (3,049 m) operating 40,000 ft. (12,195 m) non-operating
Acoustical noise level	Less than 65 dBa maximum