

**SGI™ SNIA 3000 Series
Internal Technical Configuration Manual**

Change date: May 10, 2001 11:30 am

Document Number 108-0312-001

Preliminary Information

CONTRIBUTORS

Written by <Dick Brownell>

Illustrated by <Dan Young>

Edited by <editor's name>

Production by <production editor's name>

Engineering contributions by <engineers in order of contribution>. This tag (*Names*) wraps to the next line.

© 2000, Silicon Graphics, Inc.— All Rights Reserved. Specifications subject to change without notice. SGI, the SGI logo, and Origin are trademarks of Silicon Graphics, Inc.

The contents of this document may not be copied or duplicated in any form, in whole or in part, without the prior written permission of Silicon Graphics, Inc.

RESTRICTED RIGHTS LEGEND

Use, duplication, or disclosure of the technical data contained in this document by the Government is subject to restrictions as set forth in subdivision (c) (1) (ii) of the Rights in Technical Data and Computer Software clause at DFARS 52.227-7013 and/or in similar or successor clauses in the FAR, or in the DOD or NASA FAR Supplement. Unpublished rights reserved under the Copyright Laws of the United States. Contractor/manufacturer is Silicon Graphics, Inc., 2011 N. Shoreline Blvd., Mountain View, CA 94043-1389.

<Put REGULATORY INFORMATION here. See Silicon Graphics Editorial Style Guide for content and placement instructions.>

Contents

| | | |
|-----------|---|----|
| | List of Figures | 7 |
| | List of Tables | 9 |
| 1. | Introduction | 1 |
| | Scope | 1 |
| | Related Documentation | 1 |
| | Configuration Guide Sections | 2 |
| | Technical Terms | 3 |
| 2. | Introducing the SGI SNIA 3000 Series Systems | 9 |
| | System Scaling | 10 |
| | System Ranges | 11 |
| | SGI SNIA 3000 Series Model Naming Conventions | 12 |
| 3. | SGI SNIA 3000 Series System Overview | 13 |
| | Tall Rack | 16 |
| | Short Rack | 17 |
| | Utility Tray | 18 |
| | Power Bay | 19 |
| | Power Distribution Unit (PDU) | 21 |
| | Power Distribution Strip (PDS) | 23 |
| | Destination Kits (Power input) | 25 |
| | Compute Node (C-brick) | 26 |
| | Processor | 28 |
| | Memory | 29 |
| | Disk (D-brick) | 30 |
| | System Boot (I-brick) | 32 |

| | | |
|--|-----------|------|
| PCI Expansion (P-brick) | | . 35 |
| Router (R-brick). | | . 38 |
| Peripherals | | . 41 |
| PCI Based I/O Channels | | . 41 |
| System Control | | . 42 |
| L1 Controller | | . 42 |
| L2 Controller | | . 43 |
| System Control Network | | . 44 |
| Ethernet Hub | | . 47 |
| ISDN Router | | . 51 |
| System Cooling | | . 51 |
| 4. SGI SNIA 3000 Series Performance and Bandwidth Characteristics | | . 53 |
| C-brick (Compute Node) | | . 53 |
| D-brick (Fibre Channel Disk) | | . 54 |
| I-brick (System Boot with PCI Bus) | | . 54 |
| P-brick (PCI Expansion) | | . 55 |
| Supported PCI Cards | | . 56 |
| Supported Tape Drives | | . 60 |
| System Bisection Bandwidth | | . 61 |
| 5. System Interconnection Fabric Drawings | | . 63 |
| 2-processor through 16-processor Systems | | . 63 |
| 16-processor through 64-processor Systems. | | . 64 |
| 6. System Partitioning | | . 65 |
| Partitioning Overview | | . 65 |
| Definition of Terms. | | . 66 |
| Partitioning Rules | | . 67 |
| Valid Configurations | | . 68 |
| 7. Brick Locations and Xtown2 Cable Configurations | | . 73 |
| Brick Placement in Single-rack Systems | | . 73 |
| SGI SNIA 3200 System Brick Placement Guidelines | | . 73 |
| Single-rack SGI SNIA 3400 Brick Placement Guidelines | | . 74 |

| | | |
|------------|--|------------|
| | Brick Placement in Multiple-rack Systems | 75 |
| | Multiple Rack SGI SNIA 3400 Brick Placement Guidelines | 75 |
| | SGI SNIA 3800 Brick Placement Guidelines | 77 |
| | Xtown2 Cable Configurations | 79 |
| 8. | Configuration Guidelines | 83 |
| | System Configuration Process | 83 |
| | Guidelines for Configuring the C-brick | 85 |
| | Guidelines for Configuring D-bricks | 89 |
| | Guidelines for Determining I/O Brick Locations | 91 |
| | Order for Installing I/O Bricks in a Single-rack SGI SNIA 3400 System. | 92 |
| | Order for Installing I/O Bricks in an I/O Rack. | 92 |
| | Guidelines for Configuring Controller Cards within I/O Bricks. | 93 |
| | General Guidelines | 93 |
| | Configuring the I-brick | 94 |
| | Configuring the P-brick | 95 |
| | Configuring the Power Bay | 96 |
| | Installing Distributed Power Supplies (DPS) | 97 |
| | Guidelines for Configuring the System Control Network | 98 |
| | Base System Components. | 99 |
| | SGI SNIA 3200 System | 99 |
| | SGI SNIA 3400 System | 101 |
| | SGI SNIA 3800 Systems | 104 |
| 9. | Expanding the SGI SNIA 3000 Series System | 107 |
| | Processor Capacity | 107 |
| | Main Memory Capacity | 108 |
| | Peripheral Capacity | 108 |
| 10. | SGI 3000 Series Software Requirements | 109 |
| | Required Software. | 109 |
| | Optional Software | 110 |
| | Selecting Third-party Applications | 110 |

| | | |
|-----------|---|------------|
| A. | SGI SNIA 3200 System Configuration | 111 |
| B. | SGI SNIA 3400 System Configuration | 115 |
| C. | SGI SNIA 3800 System Configuration | 119 |
| D. | Mechanical and Electrical Specifications | 123 |
| E. | System / Channel Bandwidths | 127 |
| F. | List of Acronyms / Abbreviations | 129 |
| G. | Rear View of Bricks | 131 |
| H. | Component Dimensions, Weights, and Power | 135 |
| | Index. | 137 |

List of Figures

| | | |
|--------------------|--|----|
| Figure 3-1 | Eight-processor SGI SNIA 3000 Series Block Diagram | 14 |
| Figure 3-2 | Sixteen-processor SGI SNIA 3000 Series Block Diagram | 15 |
| Figure 3-3 | Tall Rack (Front View) | 16 |
| Figure 3-4 | Short Rack (Front View) | 17 |
| Figure 3-5 | Utility Tray (Rear View) | 18 |
| Figure 3-6 | Power Bay - Front View | 19 |
| Figure 3-7 | Power Bay - Rear View | 20 |
| Figure 3-8 | Three-Phase and Single-Phase PDUs | 22 |
| Figure 3-9 | Power Distribution Strip | 24 |
| Figure 3-10 | C-brick Block Diagram. | 26 |
| Figure 3-11 | C-brick Logic Carrier Assembly - Front View | 27 |
| Figure 3-12 | D (Disk) Brick Assembly | 30 |
| Figure 3-13 | I-brick Block Diagram | 33 |
| Figure 3-14 | I-brick Assembly - Rear View | 34 |
| Figure 3-15 | P-brick Block Diagram | 36 |
| Figure 3-16 | P-brick Assembly - Rear View | 37 |
| Figure 3-17 | R-brick Block Diagram. | 39 |
| Figure 3-18 | R-brick Assembly | 40 |
| Figure 3-19 | Rear View of L2 Controller | 43 |
| Figure 3-20 | SGI SNIA 3200 System Control Network (Typical). | 44 |
| Figure 3-21 | SGI SNIA 3400 System Control Network (Typical). | 45 |
| Figure 3-22 | SGI SNIA 3800 System Control Network (Typical). | 46 |
| Figure 3-23 | Eight-Port Ethernet Hub | 47 |
| Figure 3-24 | Single-hub Network in 128-processor System | 48 |
| Figure 3-25 | Multiple-hub Network in 256-processor System | 49 |

| | | |
|--------------------|---|----|
| Figure 3-26 | Multiple-hub Network in 512-processor System | 50 |
| Figure 3-27 | ISDN Router - Rear View | 51 |
| Figure 6-1 | SGI SNIA 3800 Partitioned Systems | 71 |
| Figure 6-2 | SGI SNIA 3400 Partitioned Systems | 72 |
| Figure 7-1 | NUMAlink Cable Connecting within a Rack and to an Adjacent Rack 79 | |
| Figure 7-2 | NUMAlink Cable Spanning One Full Rack | 80 |
| Figure 7-3 | NUMAlink Cable Spanning Two Full Racks. | 80 |
| Figure 8-1 | Memory DIMM and Processor Heat Sink Assy locations | 87 |
| Figure 8-2 | Disk Drive Installation Order | 90 |
| Figure 8-3 | DPS Power Bay Locations | 97 |

List of Tables

| | | |
|------------------|--|-----|
| Table 2-1 | Linux SSI Scaling | 10 |
| Table 2-2 | System Configuration Ranges. | 11 |
| Table 2-3 | SGI SNIA 3000 Series Systems | 12 |
| Table 3-1 | PDU Specifications | 21 |
| Table 3-2 | PDS AC Power Cord Specifications | 23 |
| Table 3-3 | Destination Kit Applications | 25 |
| Table 3-4 | Main-memory DIMM Sizes | 29 |
| Table 3-5 | Number of Available PCI Slots | 41 |
| Table 4-1 | Bandwidth Characteristics of the C-brick | 53 |
| Table 4-2 | Bandwidth Characteristics of the D-brick | 54 |
| Table 4-3 | Bandwidth Characteristics of the I-brick. | 55 |
| Table 4-4 | Bandwidth Characteristics of the P-brick | 55 |
| Table 4-5 | Supported PCI Cards | 56 |
| Table 4-6 | PCI Card - Brick and System Maximum Quantities | 58 |
| Table 4-7 | Supported Tape Drive Subsystems | 60 |
| Table 4-8 | Bisection Bandwidths | 61 |
| Table 6-1 | Partitioning Support Timeline | 65 |
| Table 6-1 | Valid System Partitioning | 68 |
| Table 8-1 | C-brick Memory Configurations | 88 |
| Table 9-1 | Memory Marketing Codes | 108 |

Preliminary Information

1. Introduction

1.1 Scope

This guide is intended for Silicon Graphics sales and support staff.

This guide is designed as an informational tool to:

- Provide the Silicon Graphics presales staff with timely, accurate information that enables them to respond promptly and knowledgeably to customer inquiries and requests.
- Enable quick, accurate identification of the specific components required to fit a specific customer solution.
- Establish and support SGI's relationship and credibility with the customer (both new and existing customers) in a way that enables us to expedite order processing while improving the quality and accuracy of the shipped system. The goal of the *SGI SNIA 3000 Series Internal Technical Configuration Manual* is to create immediate confidence that can be built into a positive, ongoing customer relationship.

The organization of the *SGI SNIA 3000 Series Internal Technical Configuration Manual* focuses on the configuration process rather than the hardware architecture. Hardware architecture information is included, but it is placed in the broader sales context of identifying and meeting customer needs and expectations. Once you understand the customer's solution requirements, you can use the configuration guide to create a specific SGI SNIA 3000 series solution.

1.2 Related Documentation

For more information, consult the following document:

- Preparing for a SGI SNIA 3000 Series Installation HR-?????-0A

1.3 Configuration Guide Sections

The configuration guide consists of the following sections:

- **Section 1 Introduction**
Purpose and structure of the configuration guide.
- **Section 2 Introducing the SGI SNIA 3000 Server Computer System**
Overview of the SGI SNIA 3000 series.
- **Section 3 System Overview**
Technical overview of the SGI SNIA 3000 series and its capabilities to support various customer workloads.
- **Section 4 Performance and Bandwidth Characteristics**
Performance and bandwidth information for all bricks.
- **Section 5 System Interconnect Fabric Drawings**
System drawings that show C-brick-to-router and router-to-router cabling.
- **Section 6 System Partitioning**
System partitioning and the system partitioning rules.
- **Section 7 I/O Brick and Xtown2 Cable Drawings**
System drawings that show C-brick-to-I/O brick cabling.
- **Section 8 Configuration Guidelines**
Step-by-step guide that defines the proper configuration for a given customer and application.
- **Section 9 Expanding a System**
Information to expand a system or add to the base configuration.
- **Section 10 Software Requirements**
Description of the basic software as well as:
 - optional software and
 - third-party applications
- **Appendixes**
Supplemental information and drawings for system configuration, and mechanical and electrical specifications.

1.4 Technical Terms

| Term | Description |
|--------------------------|--|
| Bedrock | A crossbar ASIC in the C-brick that provides connections among the processors, the memory, the interconnection fabric, and the I/O subsystem. Each C-brick has one Bedrock ASIC. |
| brick | A functional subrack; for example, C-brick, R-brick; a 19-inch rack-mounted enclosure. |
| cabinet | See rack |
| cable assembly | A group of conductors in a sheath with connectors on both ends. |
| cable management bracket | Secures cables to the side of a rack. |
| cable management clip | A device that secures a cable to a bracket. |
| cable management shelf | A horizontal bracket that supports routing of cables. |
| cable dock | A mechanical sleeve that provides a guide, strain relief and locking mechanism for attaching a connector to a bulkhead. |
| card | A printed circuit board assembly (PCBA) that typically has an edge connector and plugs into a subassembly. |
| carrier | <p>A hardware fixture that attaches to one or more printed circuit board assemblies. The purpose of a carrier is to enable easy installation or removal of PCBAs from an enclosure. The following types of carriers exist:</p> <ul style="list-style-type: none"> • logic carrier Attaches to a motherboard and/or a power board. • PCI carrier Attaches to a PCI card to enable insertion of the PCI card into I/O bricks without removing the I/O brick cover; this enables hot-plugging of PCI cards. |
| carrier actuator | A lever that, when moved, inserts or removes a card from a connector. The PCI carrier has a carrier actuator. |
| C-brick | A 19-inch rackmount enclosure that contains the memory and processors; performs the compute function for the SGI SNIA 3000 series. Also referred to as the <i>compute node</i> . |
| channel | A communication path between two devices that uses a specific protocol. Additional devices cannot be added to the channel without reconfiguration of the channel. A channel provides the highest data rates for transfers. Example: the Xtown2 channel connects the C-brick and an I/O brick. |

| Term | Description |
|-------------------------------------|---|
| chassis | <i>See</i> rack |
| compute node | <i>See</i> C-brick |
| CPU | <i>See</i> processor |
| Crosstalk | 1. An I/O channel protocol name. 2. A single-ended I/O channel that uses the Crosstalk channel protocol. Crosstalk transfers data between Xbridge ASICs or between Xbridge and XC chips within a brick. |
| Crosstown | <i>See</i> Xtown |
| D-brick | The D-brick is a purchased disk enclosure that supports twelve disk drives. The D-brick is 4-U high and mounts in a standard 19-inch rack. |
| DIMM (dual inline memory module) | A printed circuit board assembly (PCBA) that contains main and directory memory. Two DIMMs represent a bank of memory. |
| DPS (distributed power supply) | One of six power supply units that are inserted into the power bay. |
| drive dock | The slot, compartment, or receptacle cavity within the chassis that receives the wrapper and provides connectivity for the hard drive. |
| harness assembly | A group of connectors that are interconnected by two or more cables. |
| hot plug | Hot plug requires that the device be deconfigured from the system prior to removing, adding, or replacing the device. After the new device is installed, action must be taken to reconfigure the system before the new device can be used. The system remains powered on and functioning during this operation. Examples of Hot Pluggable devices: PCI cards, disk drives |
| hot swap | Hot swap is the ability to remove, add, or replace a device without informing the system. This action is taken with power on and the system functioning. Example: the cooling fans on a brick and the power supplies in the power bay are components that can be hot swapped. |
| Bedrock | A crossbar ASIC that is located in the C-brick. The Bedrock provides connections among the processors, the memory, the interconnection fabric, and the I/O subsystem. Each C-brick has one Bedrock chip. |
| I-brick | The rackmount enclosure that contains the electronics and hardware necessary to boot a system and supports four additional PCI cards. The I-brick is 4 U high and mounts in a standard 19-inch rack. |

| Term | Description |
|------------------------|---|
| interconnection fabric | The interconnection fabric consists of a set of cables and routers that link together compute nodes (C-bricks). |
| L1 system controller | The brick-level system controller. It is responsible for power control and sequencing, environmental control and monitoring, initiation of reset, and the storage of identification and configuration information for its host brick. Also referred to as the <i>Level 1</i> system controller. |
| L2 system controller | The rack-level system controller. It manages central communications for the rack and controls all the bricks in that rack. When an L2 is configured in a system, all L1 system controllers are connected to the L2 controller. Also referred to as the <i>Level 2</i> system controller. |
| L2 touchscreen display | A 2.5-inch by 4-inch 70-position touchscreen display that is used to access system control information. |
| L3 system controller | The L3 system controller is a standalone workstation (or laptop) that runs the Linux operating system. The L3 system controller provides a central point of control for the entire system. Also referred to as the <i>Level 3</i> system controller. |
| link | A one-to-one connection between two processors or nodes in a multiprocessor computer system. |
| Merced | An Intel code name/trademark for a class of 64-bit processor chips (IA64) designed by Intel corporation. The new name for Merced is Itanium. |
| Metarouter | The Metarouter is identical to the router, extends the interconnect fabric, and is used for the interconnection of routers. |
| midplane board | A PCBA that is mounted vertically in a frame; it has connectors on both sides into which other PCBAs are inserted. A midplane board generally has minimal logic; the main function of a midplane board is to interconnect multiple PCBAs. |
| MIPS | 1. An acronym for million of instructions per second. 2. MIPS Technologies, Inc. is the name of a former SGI subsidiary that develops processor ASICs. These processor ASICs are referred to as MIPS processors. |
| module | An independent assembly of electronic components with some distinct function. |
| motherboard | A printed circuit board assembly (PCBA) on which other boards or cards can be mounted. |
| mounting bracket | A bracket that attaches the brick to the rack. |
| mounting shelf | An L-shaped bracket that is installed in a rack to support installation of a subassembly. |

| Term | Description |
|--|--|
| network | A network is a communication channel that enables many different devices to be connected. A network is slower than a channel; it enables devices to be added or removed without affecting the operation of the network. |
| node | An addressable device that is attached to a network; can have multiple processors. A C-brick is referred to as a node. |
| NUMALink 3 | A communication channel between C-bricks. The bandwidth is 1.6 GB/s one direction (half duplex) and 3.2 GB/s both directions (full duplex). |
| partition | A software-defined group of processors that work collectively. |
| P-brick | The rackmount enclosure that provides a crosstalk to a PCI interface. It supports 12 PCI cards on 6 PCI buses. The enclosure is 4 U high and mounts in a standard 19-inch rack. |
| PCI (peripheral component interconnect) | An industry standard for connecting peripherals to a processor. |
| POD (power-on diagnostic) | Diagnostics that run automatically at power-up of a system. |
| port | A set of input/output registers on a device. The term port refers to one side of a channel. Example: Connect one end of the Xtown2 cable to port 1 of the C-brick. |
| power bay | A standard 3-U-high, 19-inch rackmount enclosure that contains six removable AC-to-DC power supplies. |
| power board | A PCBA that regulates DC power. |
| processor | A single ASIC that contains a control unit, arithmetic and logic unit, and cache. |
| rack | A frame assembly in which 19-inch components are mounted. Also referred to as <i>cabinet</i> , <i>chassis</i> . <ul style="list-style-type: none"> • short rack A 17-U-high rack. Only SGI™ 3200 series systems use the short rack. • tall rack A 39-U- high rack. The SGI Origin 3400 and SGI Origin 3800 series systems use the tall rack. |
| rack front | The side of a rack that a customer views. Note: The door on the front of the rack includes the SGI logo. |
| rackmount | Describes components that are mounted in a rack system. |

| Term | Description |
|-----------------------------------|--|
| rack rear | The side of a rack that is accessed when configuring a system. The door on the rear of the rack is plain. |
| R-brick | A 19-inch rackmount enclosure that functions as either a router or MetaRouter depending on its physical locations and cabling within the system. When an R-brick functions as a router, it routes information between C-bricks. When the R-brick functions as a MetaRouter, it routes information between R-bricks. The R-brick is a 2U enclosure. |
| router | A device that determines the most efficient connection of receive and send ports. |
| service shelf | A portable shelf that is used to insert and remove bricks within a rack. |
| service shelf lift | A device used in conjunction with a service shelf that allows the brick to be raised or lowered to the desired level for easier insertion or removal. |
| SGI SNIA 3200 | An entry-level system that is contained within a single cabinet and has a maximum of 8 processors. |
| SGI SNIA 3400 | A class of computer systems that has a maximum of 1 processor rack and 32 processors. Additional racks with I/O and disks can be configured in an SGI 3400 system. |
| SGI SNIA 3800 | A class of computer systems that scales from 16 processors to 512 processors. Additional I/O and disk racks are configured in an SGI 3800 system |
| sled | <i>See wrapper</i> |
| U | Unit; one U is equivalent to 1.75 inches of configurable vertical space in a rack. |
| USB (Universal Serial Bus) | An external peripheral interface standard that enables communication between a computer and external peripherals via cables using biserial transmission. The USB has a peak bandwidth of 12 Mb/s. |
| VRM (voltage regulator module) | A printed circuit board assembly that provides a regulated DC voltage from a 48-Vdc power source. |
| wrapped drive | The hard drive that is sheathed within and fastened to its wrapper. |
| wrapper | The mechanical device to which a hard drive is fastened. The wrapper slides into a system chassis to facilitate the proper electrical and mechanical connections between the disk drive and the backplane. Also referred to as a <i>sled</i> . |

| Term | Description |
|--------|---|
| XIO | An SGI proprietary I/O channel, rated at 800 MB/s full duplex. |
| Xtown2 | A differential I/O channel that uses the crosstalk channel protocol. Xtown2 is used to transfer data between the C-brick node and an I/O brick. The bandwidth is 1.2 GB/s one direction (half duplex) and 2.4 GB/s both directions (full duplex). Xtown2 is pronounced <i>Crosstown two</i> . |

2. Introducing the SGI SNIA 3000 Series Systems

The SGI SNIA 3000 series of servers comprises a family of multiprocessor distributed shared memory (DSM) computer systems. The SGI SNIA 3000 series uses a global-address-space cache-coherent multiprocessor that currently scales to 16 processors in a cache-coherent domain.

The processor node (C-brick) consists of four processors, each with private secondary cache, connected at a hub via the Synergy ASIC. The Synergy ASIC provides an interface between the hub ASIC and the processors. The hub ASIC acts as a crossbar between the processors, local SDRAM memory, the network interface, and the I/O interface. Four hubs, each supporting four processors, are connected to an 8-ported router that can connect up to 512 processors in a low latency, three level, Fat Tree topology. The modularity of the DSM approach combines the advantages of low entry-level cost with global scalability in processors, memory, and I/O.

The SGI SNIA 3000 series is based upon the Intel Itanium processor. The Itanium processor is a 64-bit processor that will initially be offered at the following clock speeds and cache sizes:

- 733 MHz with a 2 Mbyte cache (Only available in SGI SNIA 3200 series)
- 800 MHz with a 2 Mbyte cache (Available in SGI SNIA 3400 and 3800 series)
- 800 MHz with a 4 Mbyte cache (Available in SGI SNIA 3400 and 3800 series)

The SGI 3000 series uses a PCI-based I/O system as the I/O protocol.

2.1 System Scaling

The SGI SNIA 3000 series servers use the Linux operating system. Linux is a widely available Open source operating system which has evolved from smaller Intel based systems. In comparison to the IRIX operating system, which scales to hundreds of processors, the scaling ability of the Linux operating system is relatively immature. However, work is underway in the Open source community and within SGI to increase the scaling ability of Linux. It will take some time for linux to scale to large numbers of processors.

When determining the number of processors for a single system, there are two types of system scaling to consider; user intensive scaling and kernel intensive scaling. User intensive scaling or sometimes referred to as CPU intensive is based on applications that spend more than 90 percent of their time in the CPU without making system calls. Many scientific applications fall into this category of CPU intensive. Kernel intensive scaling is scaling that is based on the number of processors the operating system can efficiently support when the application spends less than 90 percent of its time in the CPU. In other words, the application spends more time making kernel calls to the point that the application becomes inefficient.

As the Linux operating system matures the number of processors that are efficiently supported by a single operating system will increase. Refer to Table 2-1 for tentative availability dates on larger SSI sizes.

Table 2-1 Linux SSI Scaling

| SSI Size | User Intensive Tentative Availability Dates for Early Adopter | User Intensive Tentative Availability Dates for Production | Kernel Intensive Tentative Availability Dates for Early Adopter | Kernel Intensive Tentative Availability Dates for Production |
|--------------|---|--|---|--|
| 8 Processor | | | 3Q01 | 4Q01 |
| 16 Processor | 2Q01 | 3Q01 | 4Q01 | 3Q02 |
| 32 Processor | 3Q01 | 4Q01 | 2Q03 | 4Q03 |
| 64 Processor | 4Q01 | 2Q02 | N/A | N/A |

2.2 System Ranges

Table 2-2 describes ranges of system capacities and performance.

Table 2-2 System Configuration Ranges

| CATEGORY | MINIMUM | MAXIMUM |
|-------------------------------|---|--|
| Number of Processors | 2 | 512 ^{note 1} |
| Peak Performance | ~5.86 GFLOPS (one 2-processor node brick) | ~409.6 GFLOPS (128 four-processor node bricks) |
| C-brick Memory Capacity | 0.512 Gbytes | 8 Gbytes |
| System Main Memory Capacity | 0.512 Gbytes (one node brick) | 1024 Gbytes (128 node bricks) |
| Number of I/O Channels | 1 | 128 |
| Aggregated peak I/O bandwidth | 0.8 GB/s (peak) (one I-brick) | ~ 300.8 GB/s (peak) (8 I-brick and 120 P-bricks) |

Note: 1. A 512 processor SGI SNIA 3800 system would be partitioned system. One possible configuration would be to partition a 512 processor system as eight 64 processor partitions.

2.3 SGI SNIA 3000 Series Model Naming Conventions

Table 2-3 describes the three classes of the SGI SNIA 5000 series systems.

Table 2-3 SGI SNIA 3000 Series Systems

| Series | Description |
|---------------|--|
| SGI SNIA 3200 | Contains 2 to 8 processors with a maximum of 2 I/O bricks in a short rack. Additional racks can be added for disks. |
| SGI SNIA 3400 | Contains 4 to 32 processors. The standard processor and I/O configuration is contained in a single tall rack. Additional disk racks can be added as required. |
| SGI SNIA 3800 | Contains from 16 processors (in a single tall compute rack) to 512 processors (sixteen tall compute racks) and a maximum of eight I/O racks. Each I/O rack contains a maximum of eight I/O bricks. Additional disk racks can be added as required. |

3. SGI SNIA 3000 Series System Overview

This section provides a general, technical description of the SGI SNIA 3000 series hardware architecture, especially as it relates to configuration issues.

The SGI SNIA 3000 series systems are distributed shared memory (DSM) computer systems that scale from 2 to 64 processors as a single system image (SSI). In a DSM system, each processor contains memory that it shares with the other processors in the system. The modularity of the DSM systems combine the advantages of low entry-level cost with global scalability in processors, memory, and I/O. Multiple partitions are interconnected to form system processor counts up to 512 processors. Table 2-2 lists the system configuration ranges for the SGI SNIA 3000 series system.

Initial SGI SNIA 3000 series systems use the Intel Itanium™ processor that is a 64-bit processor with an 733 MHz or 800-MHz operating clock frequency.

The processors, each with 2 Mbytes or 4 Mbytes of private secondary cache, are connected at a Bedrock via a Synergy ASIC (refer to Figure 3-1). The Synergy ASIC provides an interface between the Itanium™ processors and the Bedrock ASIC. The Bedrock ASIC acts as a crossbar between the processor interface, local memory interface, the network interface, and the I/O interface.

Four Bedrocks, each containing a maximum four processors, connect to an eight-port router that can connect up to 512 processors in a low latency, three level, Fat Tree topology.

Note: All transfer rates in Figure 3-1 and Figure 3-2 are peak rates.
The “\$” in Figure 3-1 and Figure 3-2 means “cache”.

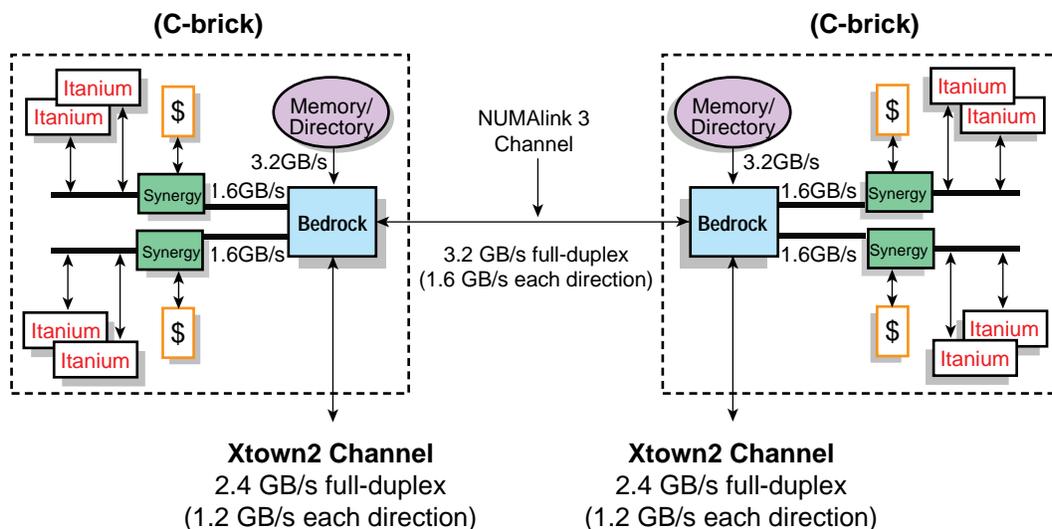


Figure 3-1 Eight-processor SGI SNIA 3000 Series Block Diagram

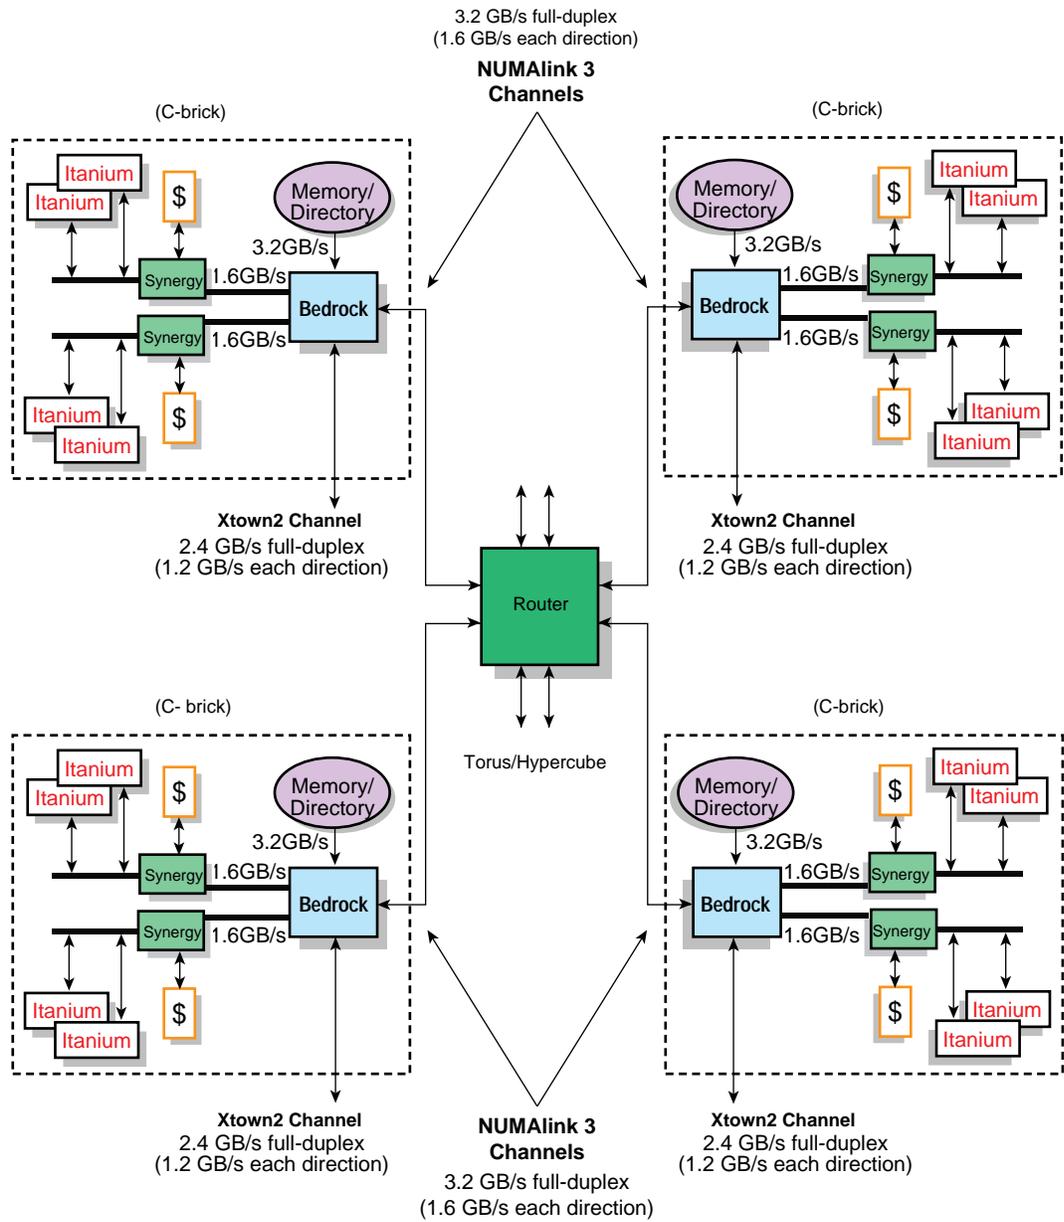
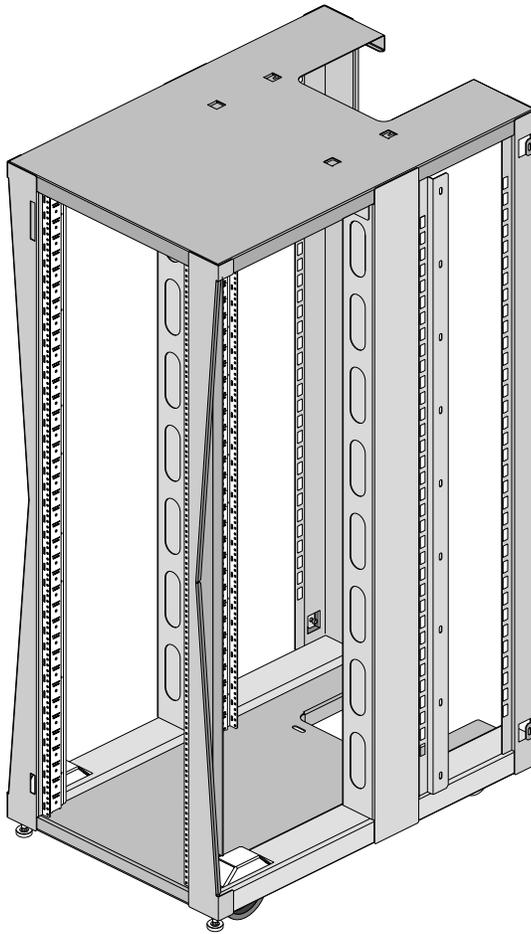


Figure 3-2 Sixteen-processor SGI SNIA 3000 Series Block Diagram

3.1 Tall Rack

The SGI SNIA 3000 series tall rack can house the following standard 19-inch rackmounted subassemblies: C-brick, D-brick, I-brick, P-brick, and R-brick. The tall rack is used by SGI SNIA 3400 and SGI SNIA 3800 systems. The outer dimensions of the tall rack with casters, side panels, and decorative doors does not exceed 74 in. high \times 30 in. wide \times 50 in. deep.



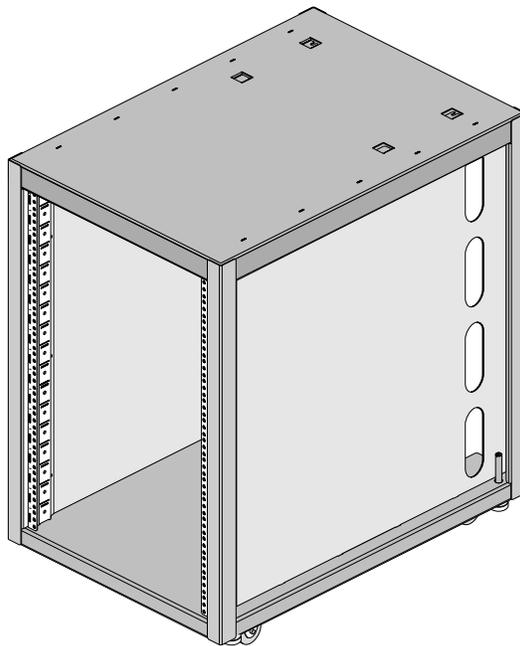
Main features of the tall rack:

- 19-inch EIA standard mounting rails
- 39 U of space (1U = 1.75 inch)
- Mounted on casters
- All subracks must provide self cooling with airflow from front to rear
- Power distribution strip (PDS), 200-250 Vac, 10 A, single-phase power for D-bricks and third-party equipment

Figure 3-3 Tall Rack (Front View)

3.2 Short Rack

The SGI SNIA 3000 series short rack can house the following standard 19-inch rack mounted subassemblies: C-brick, D-brick, I-brick, and P-brick. The short rack is used by SGI SNIA 3200 systems. The outer dimensions of the short rack with casters, side panels, and decorative doors does not exceed 36 in. high \times 26 in. wide \times 41 in. deep.



Main features of the short rack:

- 19-inch EIA standard mounting rails
- 17 U of space (1 U = 1.75 inch)
- Mounted on casters
- All subracks must provide self-cooling with airflow from front to rear
- One power distribution strip (PDS), 200-250 Vac, single-phase, 10 A power for power bay and third-party modules and D-bricks.

Figure 3-4 Short Rack (Front View)

3.3 Utility Tray

The utility tray is a 2 U shelf that mounts in the top 2 U (slot 38 and 39) of the first compute rack (001) of an SGI SNIA 3800 system. The purpose of the utility tray is to hold miscellaneous peripheral devices, such as the Ethernet hub, ISDN router and modem. The standard locations for these peripherals is shown in Figure 3-5. The devices placed on the utility tray will be held in place with velcro straps. A system that is configured with an ISDN router normally will not have a modem. Therefore, if a modem is configured, it will be placed in the same location as the ISDN router. The physical dimensions of the utility tray are 3.25 H × 17.5 W × 27.25 D.

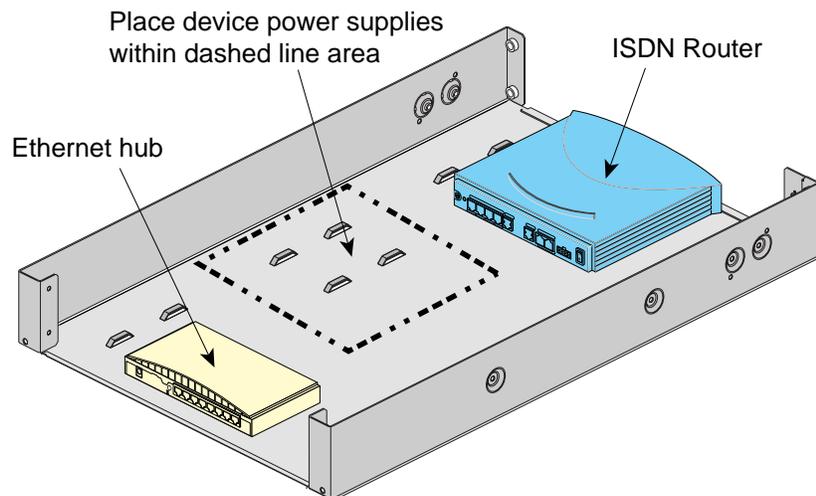


Figure 3-5 Utility Tray (Rear View)

The utility also has a 9-pin Dsub connector in the front of the tray (not shown). A cable is connected from the inside of the tray to this connector and left coiled in the utility tray for future use. This connector and cable is used to connect a PC to the console port of a single C-brick.

3.4 Power Bay

The power bay houses from two to six hot-swap, distributed power supplies (DPSs). It supplies AC power to the DPSs and provides power control and monitoring. A minimum of two DPSs must be present at all times to provide standby 48-Vdc power. The outputs of the DPSs are bused together to provide 4750 watts of available power in an N+1 redundant configuration. DPSs are added when additional subtracks are added to the configuration.

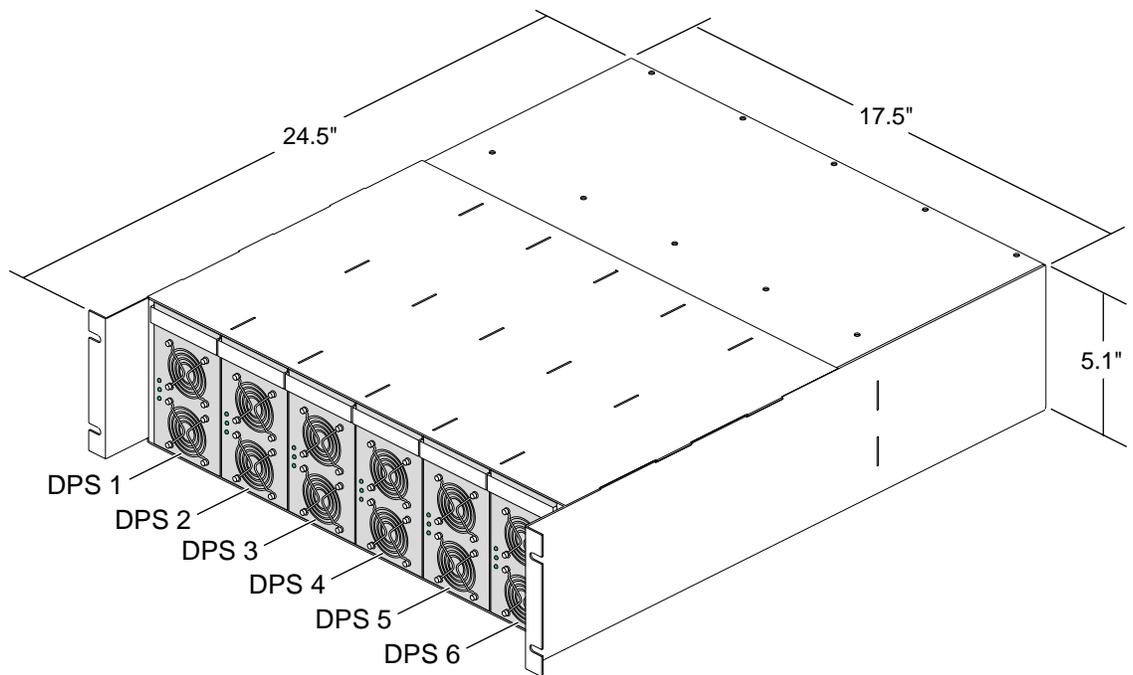


Figure 3-6 Power Bay - Front View

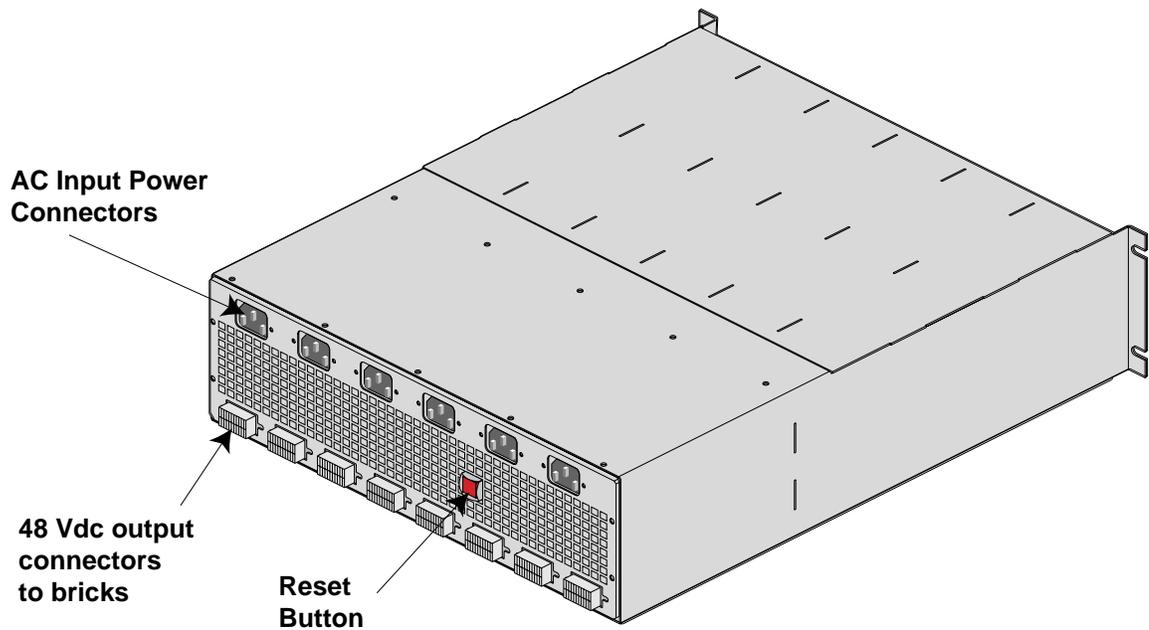


Figure 3-7 Power Bay - Rear View

Main Features of a power bay:

- Height: 3 U
- Dimensions in inches: 5.1 H × 17.5 W × 24.5 D
- Weight: 72 lbs (fully loaded with 6 power supplies)
- Provides 4.75 KW continuous N+1 power
- Provides eight 48-Vdc power output connections that use 21-pin Foxconn connectors

Main Features of a Distributed Power Supply:

- Installs from the front of the rack
- Dimension in inches: 5.0 H × 2.8 W × 13 D
- Approximate weight of 7.5 lbs
- Rated at 950 W maximum output power

3.5 Power Distribution Unit (PDU)

The SGI SNIA 3400 and SGI SNIA 3800 systems support five types of PDUs (refer to Table 3-1). The PDUs protect against over-current conditions and provide an on/off switch to remove power from the rack.

Note: The SGI SNIA 3200 systems do not use PDUs; instead they use a power distribution strip.

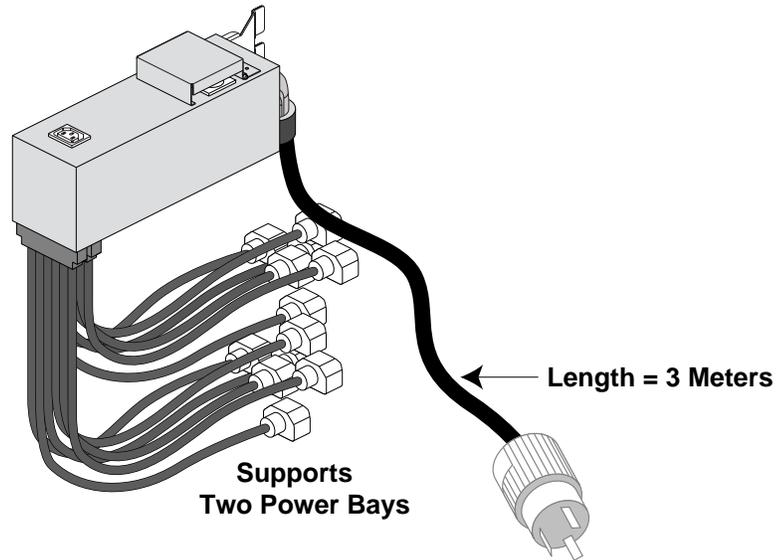
Customer sites that use single-phase power require one single-phase PDU for each power bay. Customer sites that use three-phase power require one three-phase PDU for each rack.

The physical dimensions of the PDU enclosure are $4 \times 5 \times 10$ inches. It mounts in the lower cable management area in the rear of the full-height racks.

Table 3-1PDU Specifications

| Marketing Code | Destination Location | Power Cord | Input Power Connector |
|---|-------------------------------------|---|---|
| DK-N1P-001 (Single-phase NEMA Connector) | Domestic, Mexico, Canada, and Japan | UL Listed Cord 30A 3 Wire | NEMA L6-30 1Ph 30A Plug 208/240 VAC dual input |
| DK-N1P-002 (Single-phase IEC Connector) | Europe and other | 32A Harmonized Cord 3 Wire | IEC 60309 1Ph 32A Plug 208 / 240 VAC dual input |
| DK-N1P-003 (Single-phase IEC Connector) | Domestic, Mexico, Canada, and Japan | UL Listed Cord 30A 3 Wire | IEC 60309 1Ph 30A Plug 208/240 VAC dual input |
| DK-N3P-001 (Three-phase delta) | Domestic, Mexico, Canada, and Japan | UL Listed Cord 60A 4 Wire | IEC 60309 3Ph 60A Plug 200/240 VAC |
| DK-N3P-002 (Three-phase wye) | Europe and other | 5 Wire Harmonized Cord 32A IEC Rated | IEC 60309 32A 3Ph Plug 400 VAC |

Three-Phase PDU



Single-Phase PDU

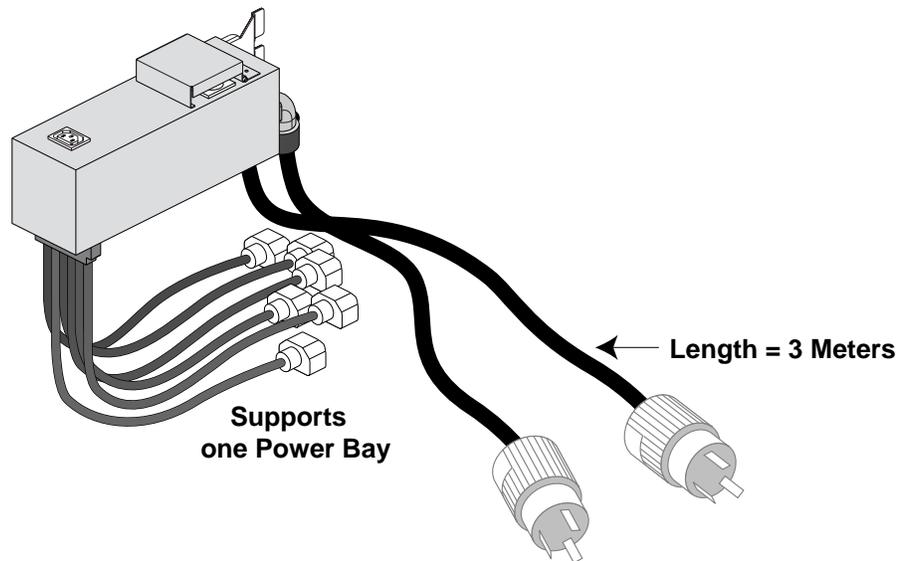


Figure 3-8 Three-Phase and Single-Phase PDUs

3.6 Power Distribution Strip (PDS)

All SGI SNIA 3000 series systems use power distribution strips. In the SGI SNIA 3200 short rack the PDS provides:

- AC power distribution to the power bay
- AC power to a D-brick
- AC power to third party equipment
- Over-current protection for all equipment in the short rack
- An on/off switch to remove power from the short rack

In SGI SNIA 3400 and SGI SNIA 3800 full-height racks the PDS provides:

- AC power to a D-brick
- AC power for third party equipment
- Over-current protection for a D-brick and third party equipment
- An on/off switch to remove power from D-bricks and third party equipment.

In both the short rack and tall rack, the PDS is located on the inside rear wall. Its dimensions are 12 × 2.5 × 3.5 inches. In a short rack, a three-meter power cord connects the PDS to either the AC wall or underfloor outlet. In a tall rack, a 1.5-meter power cord connects the PDS to the PDU.

Refer to Table 3-2 for the types of power cords for a PDS and to Figure 3-9 for a drawing of the PDS.

Table 3-2 PDS AC Power Cord Specifications

| RESOURCE-CODE | Description | Country Required |
|-----------------|----------------------|--------------------|
| DK-S-PWRCBL-001 | 2.5-meter power cord | United States |
| DK-S-PWRCBL-002 | 2.5-meter power cord | United Kingdom |
| DK-S-PWRCBL-003 | 2.5-meter power cord | Italy |
| DK-S-PWRCBL-004 | 2.5-meter power cord | Switzerland |
| DK-S-PWRCBL-005 | 2.5-meter power cord | Continental Europe |
| DK-S-PWRCBL-006 | 2.5-meter power cord | Israel |

Table 3-2 PDS AC Power Cord Specifications

| RESOURCE-CODE | Description | Country Required |
|-----------------|----------------------|------------------|
| DK-S-PWRCBL-007 | 2.5-meter power cord | India |
| DK-S-PWRCBL-008 | 2.5-meter power cord | Australia |

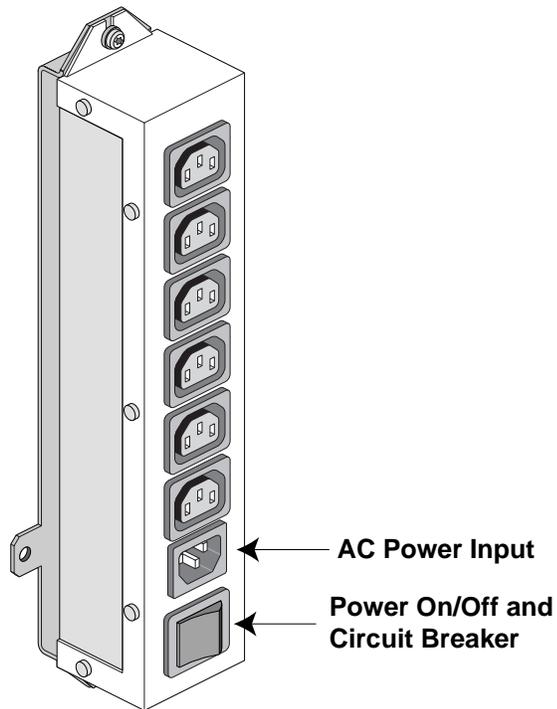


Figure 3-9 Power Distribution Strip

3.7 Destination Kits (Power input)

Destination kits contain the country specific power cords. Table 3-3 defines the type and quantity of destination kits to order based on the country the site is located in and the specific configuration of the system. Refer back to section 3-5 and 3-6 for a detailed description of each destination kit.

Table 3-3 Destination Kit Applications

| Connected to Equipment type | COUNTRY | Destination Kit | Quantity DK Kit Required |
|---------------------------------------|---------------------------------|-----------------|---|
| Power bay in compute rack or I/O rack | U.S.A., Canada Mexico, Japan | DK-N1P-001 | (1) kit per power bay |
| | | DK-N1P-003 | |
| | Europe and other | DK-N3P-001 | (1) kit per rack |
| | | DK-N1P-002 | (1) kit per power bay |
| | | DK-N3P-002 | (1) kit per rack |
| Power distribution strip | U.S.A, Canada Mexico, Japan | DK-S-PWRCBL-001 | (1) kit for each additional PDS added to an I/O rack. (1) kit per 3200 series short rack |
| | United Kingdom | DK-S-PWRCBL-002 | |
| | Italy | DK-S-PWRCBL-003 | |
| | Switzerland | DK-S-PWRCBL-004 | |
| | Continental Europe | DK-S-PWRCBL-005 | |
| | Israel | DK-S-PWRCBL-006 | |
| | India | DK-S-PWRCBL-007 | |
| | Australia | DK-S-PWRCBL-008 | |

3.8 Compute Node (C-brick)

The C-brick is a 3-U-high 19-inch rackmountable enclosure that contains:

- Either two or four 64-bit Intel Itanium™ processors, each with either a 2MB or 4MB L3 cache
- Eight DIMM slots; two DIMM slots comprises one bank of memory
- Node electronics
- One L1 controller

The node electronics, L1 controller, and power regulators are contained on a single printed circuit board (PCB). The processors are housed on separate processor heat sink assemblies. Each processor heat sink assembly contains one processor, L3 cache and heat sink. The L4 caches, Synergy ASICs and Bedrock ASIC are mounted on the IP37 board. Figure 3-10 shows the block diagram of a C-brick.

Note: All transfer rates in Figure 3-10 are peak rates.
The “\$” in Figure 3-10 means “cache.”

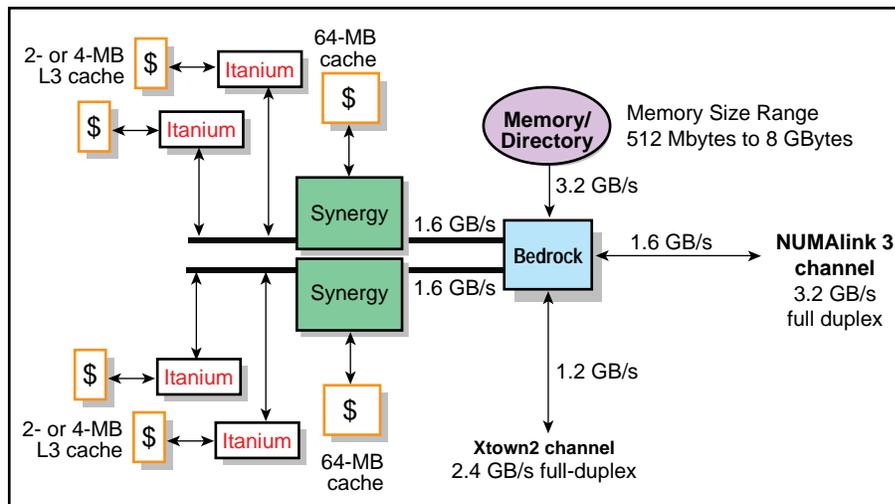


Figure 3-10 C-brick Block Diagram

The C-brick has the following electrical features:

- Configurable as either a 2-processor or 4-processor node
- Configurable from 512 Mbytes to 8 Gbytes of main memory
- Contains one 2MB or 4MB secondary cache per processor
- Contains one 1.6-GB/s (each direction) NUMALink3 channel
- Contains one 1.2-GB/s (each direction) Xtown2 channel
- Contains one USB port that connects to the L2 controller (optional in SGI SNIA 3200 systems)
- Contains one DB9 console port

Figure 3-11 shows a front view of the C-brick logic carrier (the assembly that holds the C-brick components).

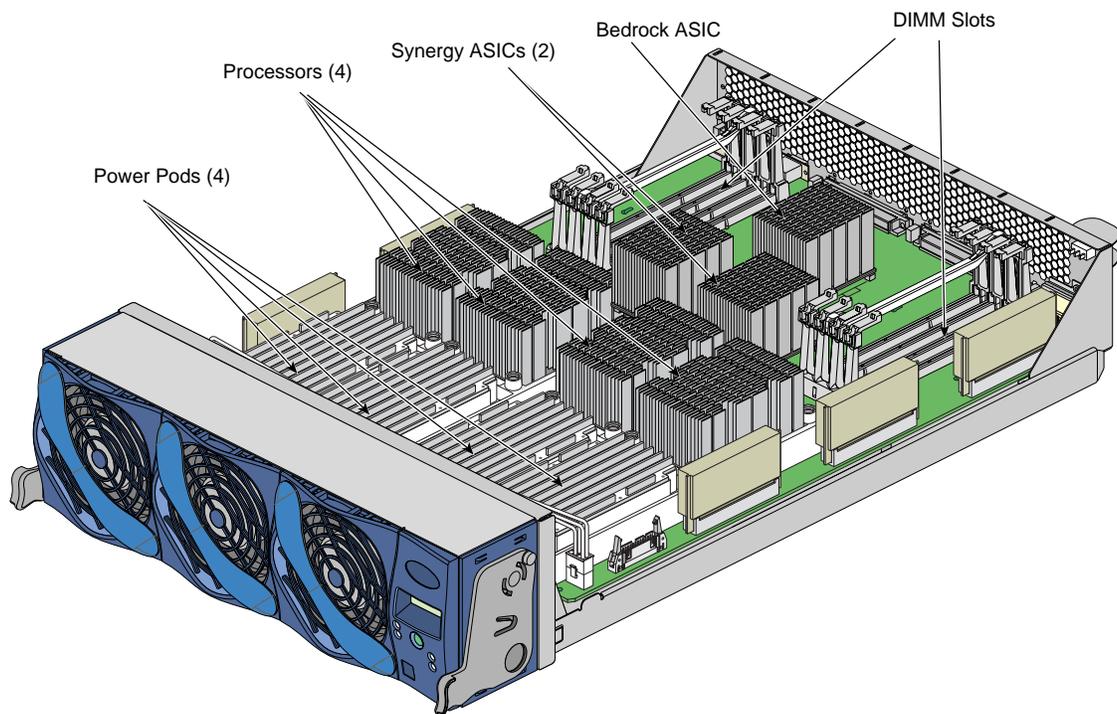


Figure 3-11 C-brick Logic Carrier Assembly - Front View

The C-brick has the following mechanical features:

- Height: 3 U
- Weight: 55 lbs (fully loaded)
- Dimensions in inches: 5.06 H × 17.19 W × 27.80 D
- Installs from the front of the rack
- Cables are located in the rear
- Cooling fans are hot-swappable
- Input power is +48 Vdc (~825 Watts)

3.8.1 Processor

One 733MHz or 800MHz Intel Itanium™ processors, 2MB or 4MB of L3 cache, and heat sink are combined to form an assembly. This assembly is called the processor heat sink Assembly. The C-brick can be configured with two or four processor heat sink assemblies.

The Intel Itanium™ processor has the following features:

- 64-bit design
- 2MB or 4MB secondary cache

3.8.2 Memory

Main memory consists of up to eight memory banks per node; each bank is split between two DIMMs of a DIMM pair with each DIMM pair supporting two banks. Memory must be increased or decreased in two-DIMM increments. The reason for this is that a single bank of memory is contained on two DIMMs and the memory size must be increased or decreased in whole banks. The DIMMs that make up a single bank must be the same memory size; however, each DIMM pair within a brick can be a different memory size. Refer to Table 3-4 for the main memory size matrix. This table does not include systems that have multiple-size memory banks. Refer to Table 8-1, “C-brick Memory Configurations,” on page 88 for memory sizes that use mixed size memory DIMMs.

The clock speed of the memory parts is 100-MHz address and 200-MHz data, which produces a memory bandwidth of 3200 MB/s.

Table 3-4 Main-memory DIMM Sizes

| Marketing Code | DRAM Technology | Single DIMM Size | Minimum Increment (2 Dimms) | 1 Bank Installed | 2 Banks Installed | 3 Banks Installed | 4 Banks Installed |
|----------------|-----------------|------------------|-----------------------------|------------------|-------------------|-------------------|-------------------|
| MEM-512 | 128 Mbits | 256 Mbytes | 512 Mbytes | 512 Mbytes | 1 Gbytes | 1.5 Gbytes | 2 Gbytes |
| MEM-1G | 128 Mbits | 512 Mbytes * | 1 Gbytes | 1 Gbytes | 2 Gbytes | 3 Gbytes | 4 Gbytes |
| MEM-1G-D | 128 Mbits | 512 Mbytes * | 1 Gbytes | 1 Gbytes | 2 Gbytes | 3 Gbytes | 4 Gbytes |
| MEM-2G-D | 256 Mbits | 1 Gbytes * | 2 Gbytes | 2 Gbytes | 4 Gbytes | 6 Gbytes | 8 Gbytes |

* The 512 Mbyte DIMMs are available as standard and premium. The 1 Gbyte DIMMs are only available as premium.

There are two DIMM types used:

- Standard memory DIMM - for systems with a maximum of 128 processors
- Premium memory DIMM - Contains one additional memory chip per DIMM to provide additional directory memory for building configurations larger than 128 processors. Customers who plan to upgrade their systems beyond 128 processors should order the premium memory DIMMs. The cost to remove the existing standard memory and replace it with premium memory is high.

Note: SGI SNIA 3000 series DIMMs are not compatible with the DIMMs used in SGI Origin 200, SGI Origin 2000, or SGI Octane systems. However, they are compatible with the SGI Origin 3000 memory DIMMs.

3.9 Disk (D-brick)

The D-brick is a 4-U high third-party disk enclosure that supports JBOD (just a bunch of disks) and RAID within an SGI SNIA 3000 series rack. The D-bricks configured as RAID will be dropped shipped to the customer from the vendor. The D-bricks configured as JBOD will be factory installed and ship as part of the system.

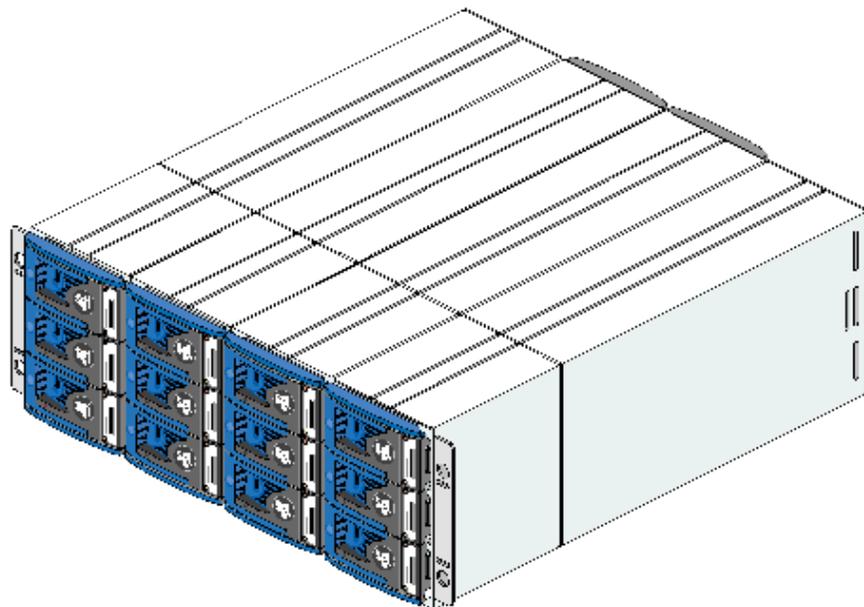


Figure 3-12 D (Disk) Brick Assembly

The SGI TP-9100 storage system is another storage solution for SGI SNIA 3000 series systems. For more information about the SGI TP-9100 storage system, refer to the *SGI Total Performance 9100 Storage System Owner's Guide*, publication number 007-4068-xxx.

The D-brick has the following features:

- Height: 4 U
- Weight: 94 lbs (fully loaded)
- Dimensions in inches: 6.95 H × 17.50 W × 23.00 D
- Maximum number of disk drives: 12
- Requires a minimum of two disk drives
- Disks are hot-pluggable and sled-mounted
- Mounts in a standard 19-inch rack
- Occupies a fixed position; does not slide out
- Supports 2 Fibre Channel loops (disk arrays)
- Input power is 200 to 230 Vac single-phase, 50/60 Hz
- Typical power consumption is 400 VA or less

3.10 System Boot (I-brick)

The I-brick provides the boot I/O functions for all SGI SNIA 3000 series systems. It has two Xtown2 ports that are configurable as either 800 MB/s or 1.2 GB/s each direction; each Xtown2 port can connect to an Xtown2 port on a C-brick. The standard configuration for the Xtown2 ports in the I-brick is 800 MB/s.

The I-brick supports five hot-pluggable PCI cards, two sled-mounted 3.5 inch Fibre Channel disk drives, and a specialized slot for a DVD/CD-ROM. The five hot-pluggable PCI slots support full-length cards with 64-bit data/addressing. Refer to Figure 3-13 for the I-brick block diagram.

The five PCI slots are configured on two buses: bus 1 supports three 33-MHz PCI slots and bus 2 supports two 66-MHz PCI slots. Separate buses enable the I-brick to run 33-MHz and 66-MHz devices in the same brick. Various types of PCI cards can be used in the I-brick, such as SCSI, Fibre Channel, ATM, Gigabit Ethernet, etc. Refer to Table 4-5 on page 56 for a list of supported PCI cards.

The I-brick also provides access to a network via a 10/100BaseT Ethernet port, and access to peripherals via one 1394 channel and two USB channels. Refer to Section 8 for configuration guidelines.

Note: The I-brick used with the SGI SNIA 3000 series systems is identical with the I-brick used with the SGI Origin 3000 series systems.

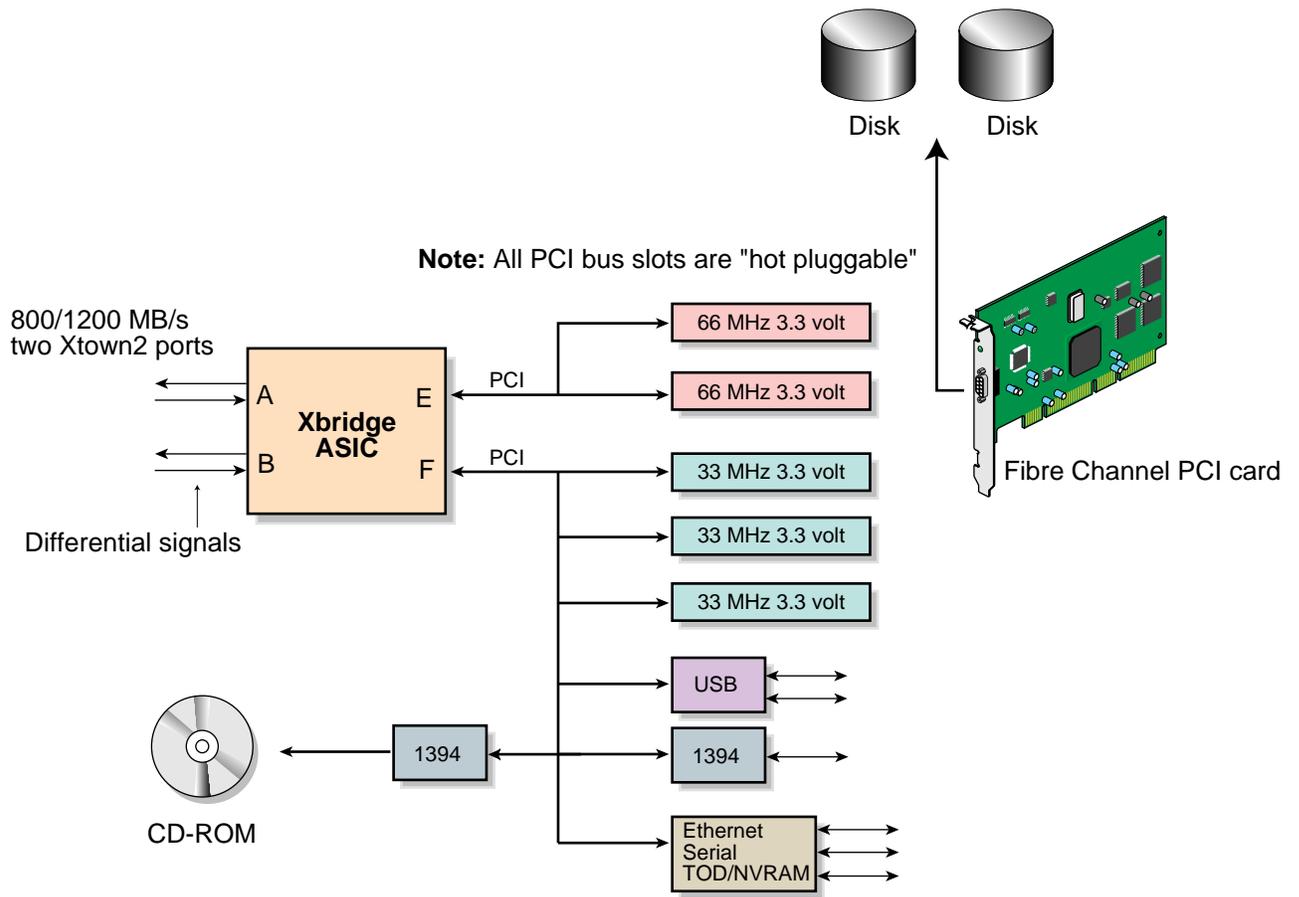


Figure 3-13 I-brick Block Diagram

Figure 3-14 shows a rear view of the I-brick with the cover removed.

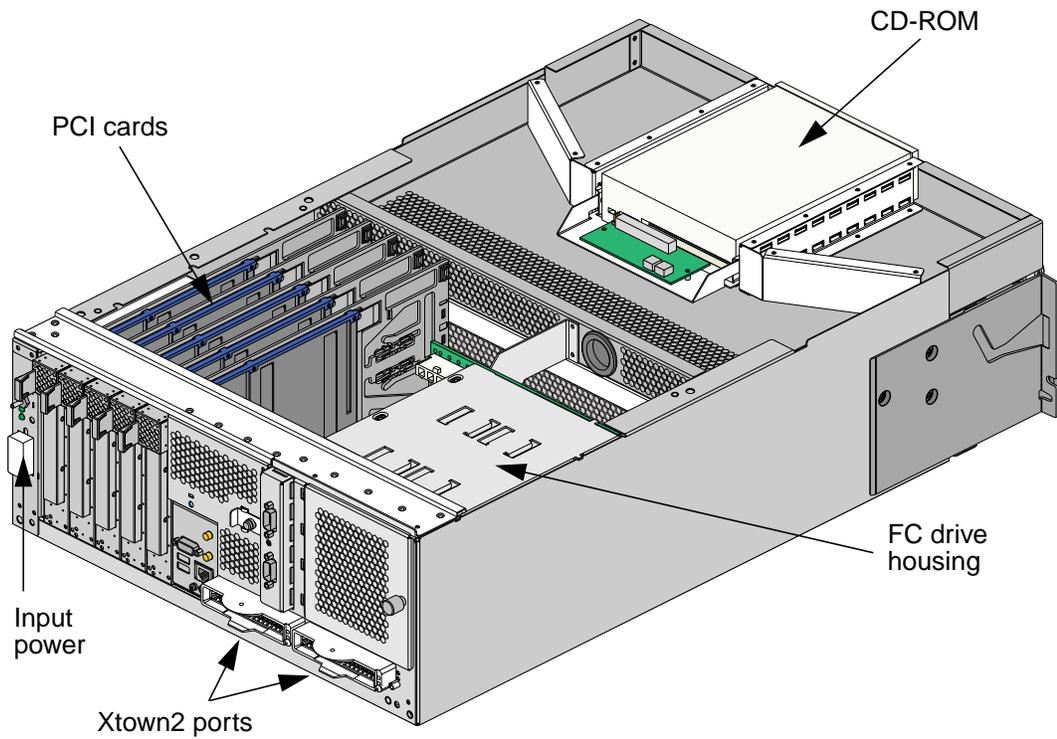


Figure 3-14 I-brick Assembly - Rear View

The I-brick has the following electrical features:

- One serial port (RS-232 or RS-422)
- Two 1394 ports (one internal, one external)
- Two external USB ports
- One 10/100BaseT Ethernet port
- RTI and RTO connections (Real Time sync I/O)
- Two Xtown2 ports - configurable as 800 MB/s or 1.2 GB/s (each direction)
Note: Standard configuration for the Xtown2 ports is 800 MB/s.

The I-brick has the following mechanical features:

- Height: 4 U
- Weight: 69 lbs (fully loaded)
- Dimensions in inches: 6.64 H × 17.50 W × 27.74 D
- Hard mounts in a standard 19-inch rack (does not slide out of the rack)
- Supports a CD-ROM
- Supports five 3.3-Vdc PCI cards (3 PCI cards at 33 MHz and two PCI cards at 66 MHz)
Note: one of the five PCI slots is reserved for a Fibre Channel disk controller.
- Supports two 3.5-inch sled-mounted Fibre Channel disk drives
- Cooling fans are hot-swappable
- Input power is +48 Vdc (~190 watts)

3.11 PCI Expansion (P-brick)

The P-brick is a Crosstalk-to-PCI based I/O expansion subsystem that supports a maximum of 12 hot-pluggable PCI cards. It has two Xtown2 ports that are configurable as either 800 MB/s or 1.2 GB/s each direction; each Xtown2 port can connect to an Xtown2 port on a C brick. The standard configuration for the Xtown2 ports in the P-brick is 1.2 GB/s.

The 12 PCI slots are configured on six buses; each bus supports two 33- or 66-MHz slots.

Refer to Table 4-5 on page 56 for a list of supported PCI cards and to Section 7 for configuration guidelines.

Note: The P-brick used with the SGI SNIA 3000 series systems is identical with the P-brick used with the SGI Origin 3000 series systems.

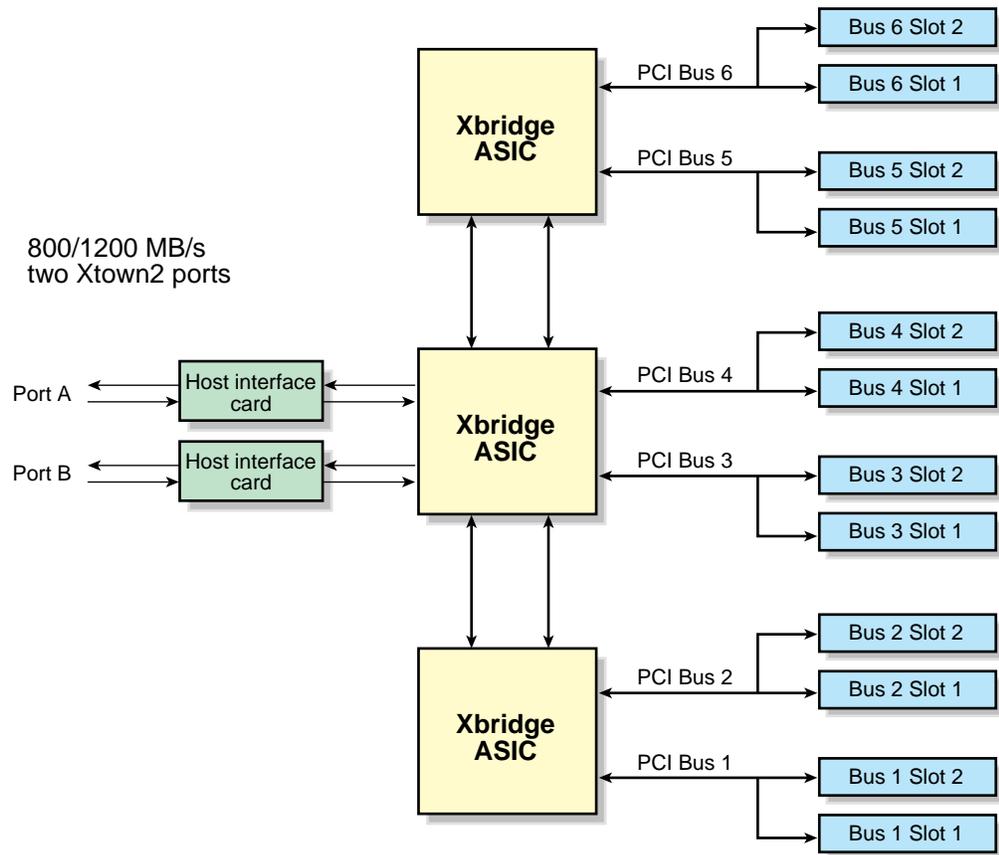


Figure 3-15 P-brick Block Diagram

Figure 3-16 shows a rear view of a P-brick with the cover removed.

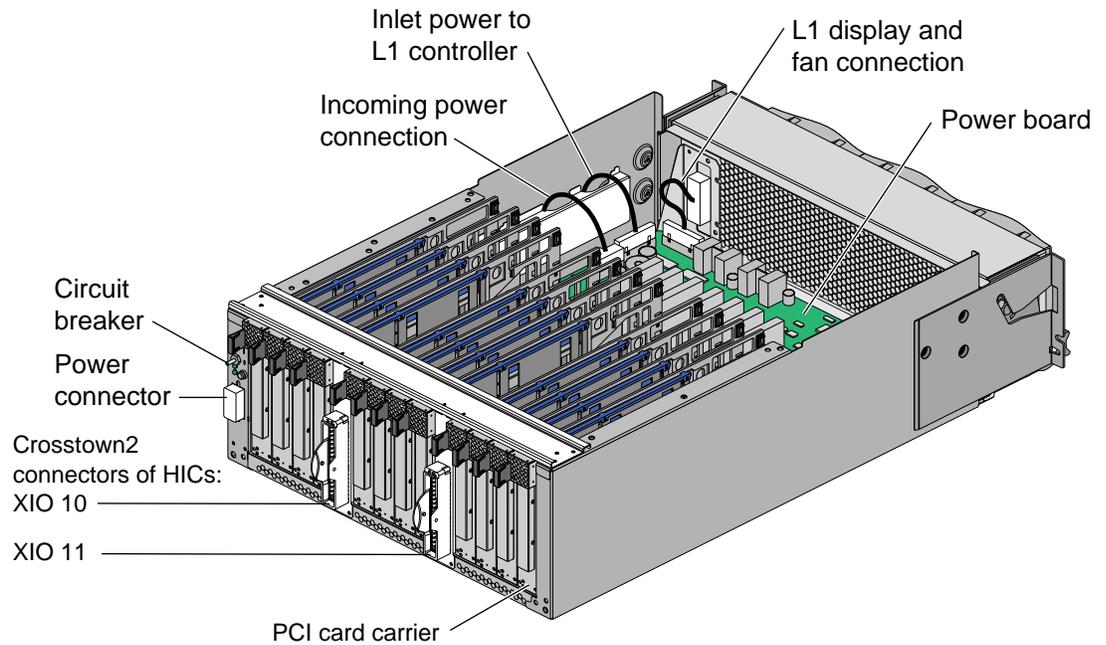


Figure 3-16 P-brick Assembly - Rear View

The P-brick has the following electrical feature:

- Two Xtown2 ports (configurable as 800 MB/s or 1.2 GB/s each direction)

The P-brick has the following mechanical features:

- Height: 4 U
- Weight: 70 lbs (fully loaded)
- Dimensions in inches: 6.64 H × 17.50 W × 27.74 D
- Hard mounts in a standard 19-inch rack (does not slide out of the rack)
- Supports twelve 3.3-volt PCI cards
- Cooling fans are hot-swappable
- Input power is +48 Vdc (~225 watts)

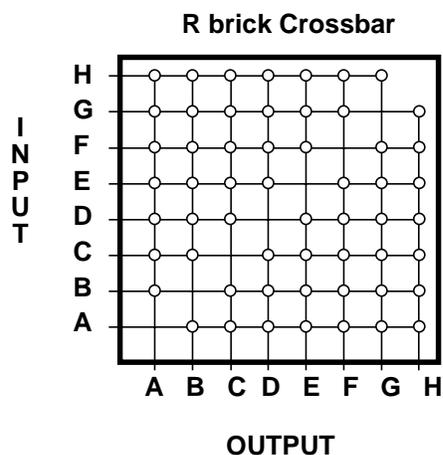
3.12 Router (R-brick)

The R-brick is an eight-port crossbar that connects any input-link channel to any of seven possible output-Link channels. It contains a router ASIC that is mounted on a PCB with its associated power circuitry, L1 controller, and a USB hub. The hub fans out USB signals from the L2 controller to the L1 controller inside the brick and to the four nodes (C-bricks) that may be connected to the router.

The R-brick has a total of eight 100-pin link connectors located on its rear panel. Four of these connect to C-bricks and carry USB signals as well as link signals. The others are only for connection to other routers and do not carry USB signals. Refer to Figure 3-17. Metarouters and repeat routers use all eight ports to connect to other R-bricks.

When an R-brick-to-R-brick connection is made through ports that carry USB signals, the USB signals are ignored. USB signals to the C-bricks are distributed over the network cables. Because an R-brick can have a maximum of four C-bricks attached to it, only four of the R-brick's 100-pin network connectors have USB signals routed to them. Ports 2, 3, 4, and 5 carry USB signals. Therefore, a C-brick must connect to an R-brick via port 2, 3, 4, or 5.

Each R-brick has a dedicated USB connection to the L2 controller through a 4-pin USB connector on its rear panel. Therefore, it is not necessary for an R-brick to distribute USB signals to other R bricks. R-brick-to-R-brick network connections are normally made through the four port connectors that do not carry USB signals; however, they are not restricted to these four ports.



Note: The R-brick used with the SGI SNIA 3000 series systems is identical with the R-brick used with the SGI Origin 3000 series systems.

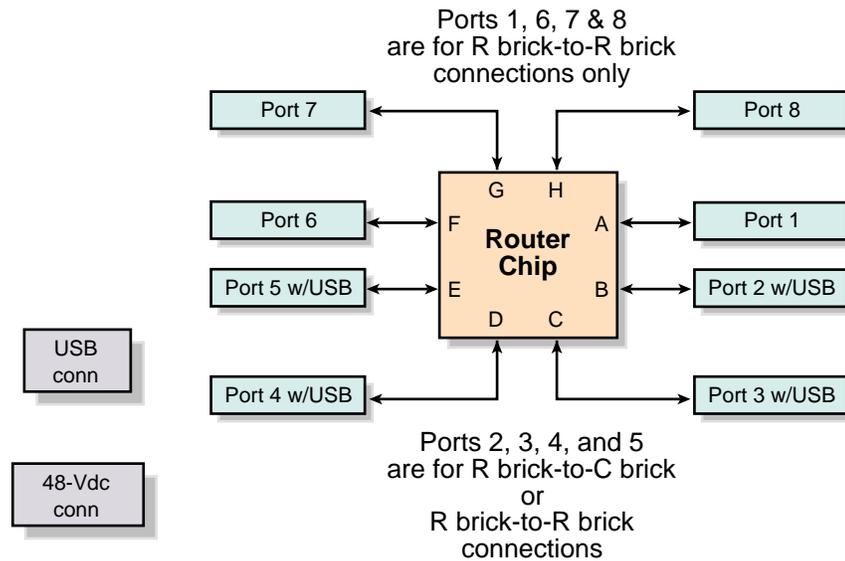


Figure 3-17 R-brick Block Diagram

The R-brick has the following electrical features:

- One USB port (connects to the L2 controller)
- Eight NUMalink I/O connectors (located in the rear)

The R-brick has the following mechanical features:

- Height: 2 U
- Weight: 18 lbs
- Dimensions in inches: 3.35 H × 17.38 W × 27.5 D
- Installs from the front of the rack
- Mounts in a standard 19-inch rack
- Occupies a fixed position, does not slide out
- Cooling fans are hot-swappable
- Input power is +48 Vdc (~60 watts)

Figure 3-18 shows a rear view of the R-brick enclosure with the cover removed.

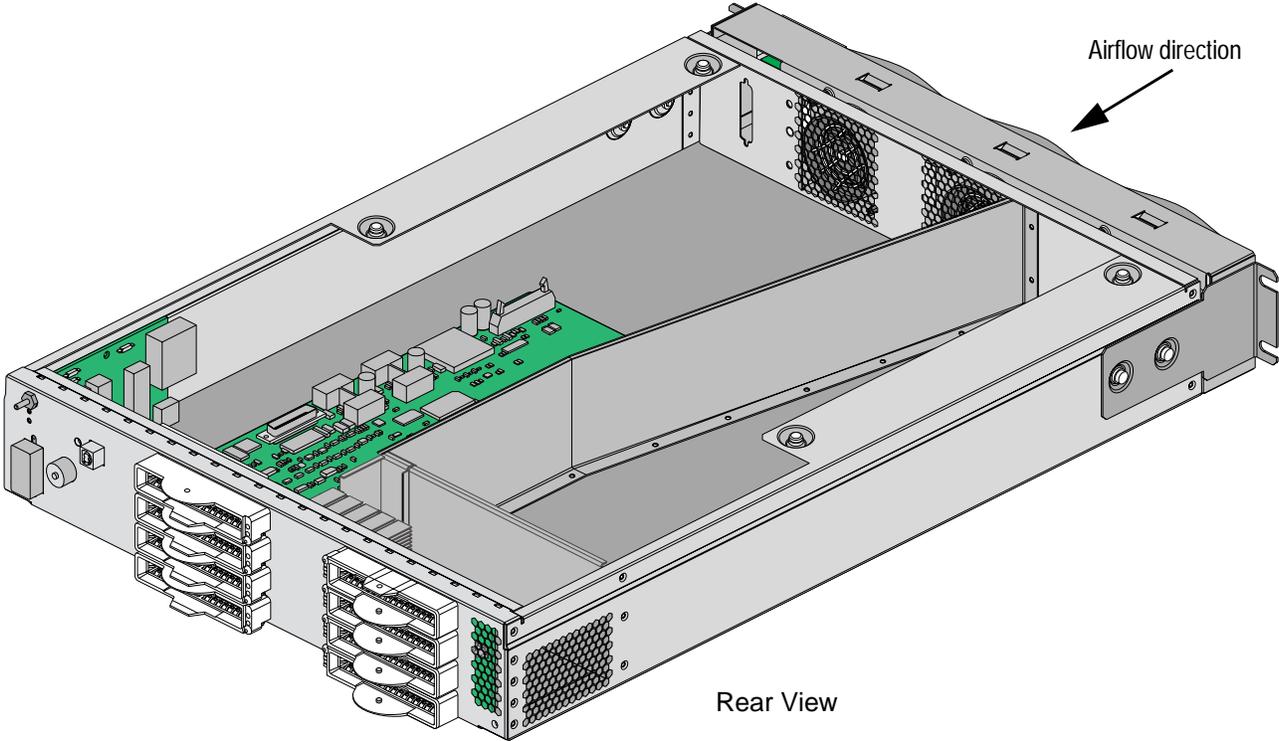


Figure 3-18 R-brick Assembly

3.13 Peripherals

An SGI SNIA 3000 series system supports a wide range of peripherals to deliver unmatched performance for large supercomputing workloads. Refer to Table 4-5 for a list of the supported PCI cards.

3.14 PCI Based I/O Channels

The I-brick and P-brick are the subsystems that provide PCI support for the SGI SNIA 3000 series systems. The I-brick provides five PCI slots: four slots for customer-defined PCI cards and one slot for a Fibre Channel disk controller PCI card. When additional PCI slots are required, a P-brick is configured into the system. The P-brick supports 12 PCI slots. Table 3-5 describes the maximum number of PCI slots for each single SSI that are available in SGI SNIA 3200, SGI SNIA 3400 and SGI SNIA 3800 systems.

Table 3-5 Number of Available PCI Slots

| System | I/O Bricks | Maximum Number of PCI Slots Available |
|---------------|--------------------------------|---------------------------------------|
| SGI SNIA 3200 | one I-brick and one P-brick | 17 |
| SGI SNIA 3400 | one I-brick and seven P-bricks | 89 |
| SGI SNIA 3800 | one I-brick and 15 P-bricks | 185 |

3.15 System Control

The purpose of the SGI SNIA 3000 series control system is to:

- Manage power control and sequencing
- Provide environmental control and monitoring
- Initiate system resets
- Provide storage for identification and configuration information
- Provide a console/diagnostic and scan interface

The SGI SNIA 3000 series control system is composed of three levels:

- L1 controller - brick-level system controller
- L2 controller - rack-level system controller

3.15.1 L1 Controller

The L1 controller is not configurable; it is designed into all bricks except the D-brick.

3.15.2 L2 Controller

The L2 controller is optional in SGI SNIA 3200 systems, however, it is required with all SGI SNIA 3400, and SGI SNIA 3800 systems. The L2 controller is a 5.5 in. × 11.1 in. PCB assembly that is mounted in the top of the rack. The L2 controller does not use configurable rack space. It receives 48-Vdc power (~30 watts) from the power bay.

The L2 provides the following communication ports (refer to Figure 3-19):

- One USB host controller with 4 ports
- One 10/100BaseT auto-negotiating Ethernet port with RJ45 connector
- Two RS-232 ports with a modem control that is capable of 115 Kbaud
- One RS-485 port that is capable of 19.2 Kbaud
- One rack display connector

The L2 controller is required in a rack when:

- The rack contains an R-brick
- Remote maintenance of the system is required (SGI SNIA 3200 systems)
- A rack display is desired

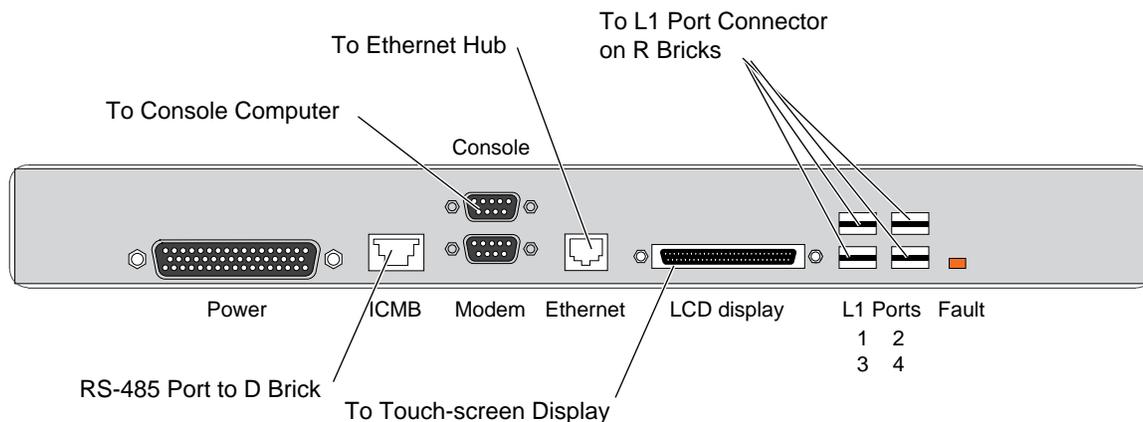


Figure 3-19 Rear View of L2 Controller

3.15.3 System Control Network

The system control network configuration depends on the class of the SGI SNIA 3000 series system. Figure 3-20 through Figure 3-22 show typical SGI SNIA 3200, SGI SNIA 3400, and SGI SNIA 3800 system control configurations.

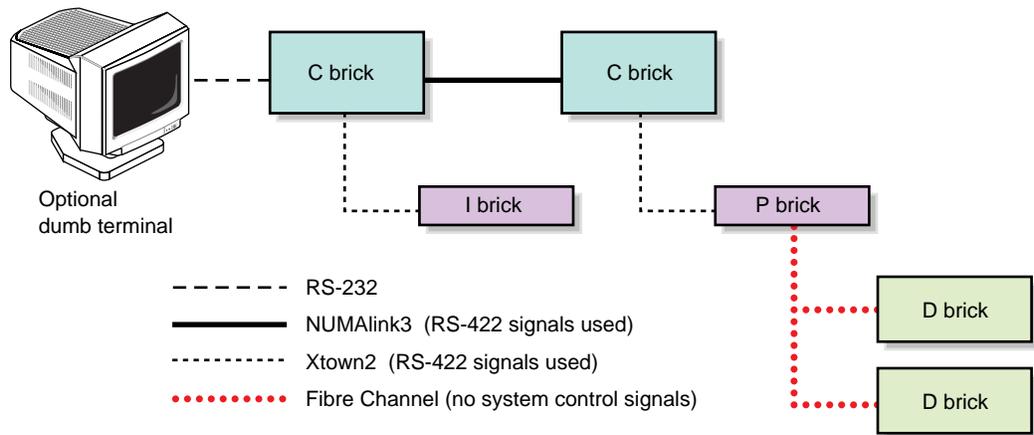
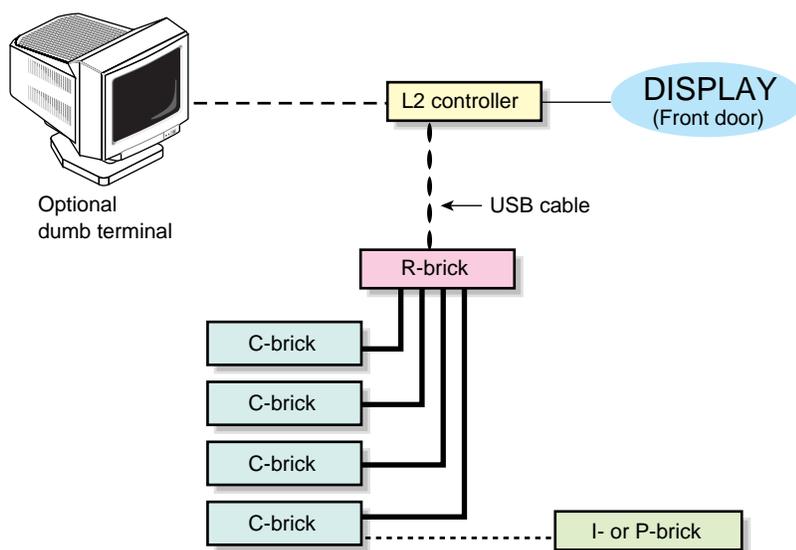


Figure 3-20 SGI SNIA 3200 System Control Network (Typical)



- RS-232
- NUMalink 3 (USB signals used)
- Xtown2 (RS-422 signals used)
- USB

Figure 3-21 SGI SNIA 3400 System Control Network (Typical)

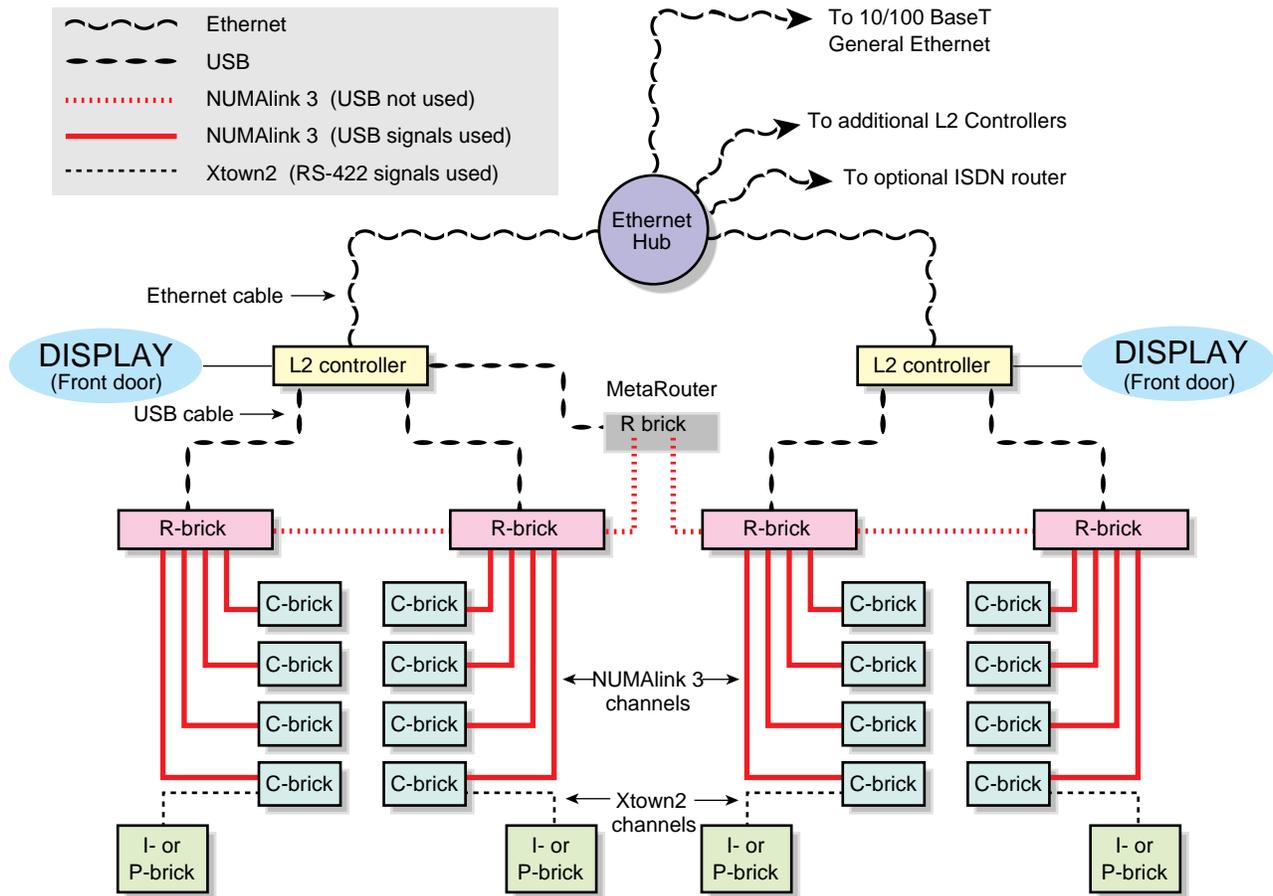


Figure 3-22 SGI SNIA 3800 System Control Network (Typical)

3.15.4 Ethernet Hub

Systems that have multiple compute racks require an Ethernet hub. The Ethernet hub is used to interconnect L2 controllers.

- One Ethernet hub is required for systems up to 128 processors
- Two Ethernet hubs are required for systems greater than 128 and less than or equal to 256 processors
- Three Ethernet hubs are required for systems greater than 256 processors

The Ethernet hub is located on a 2-U utility shelf in the top two locations of a compute rack.

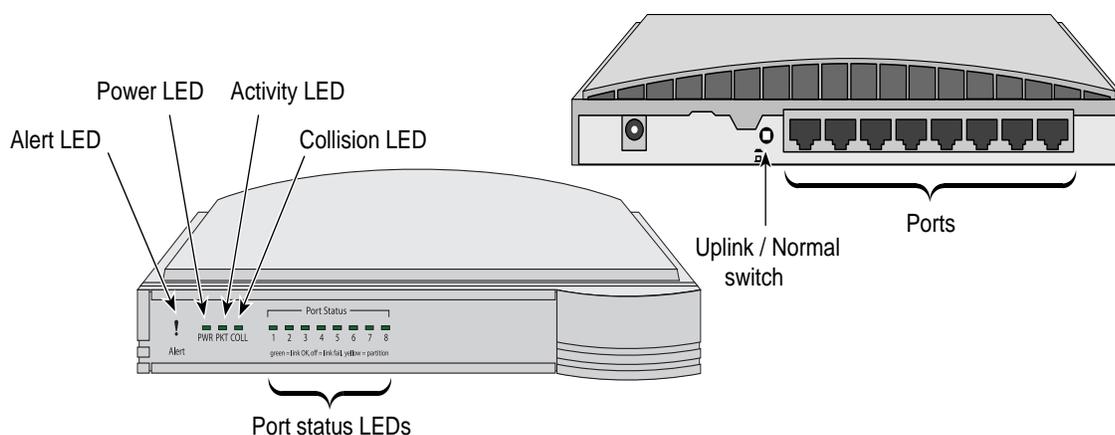


Figure 3-23 Eight-Port Ethernet Hub

The Ethernet hub has the following features:

- Eight RJ-45 ports (port 8 is controlled by an Uplink / Normal switch)
- Maximum 100 meter cable length
- Weight: 1.1 lbs
- Dimensions in inches: 1.4 H × 9.0 W × 5.3 D
- Input power: requires an adapter which converts AC wall power to 8-9 Vac, 50-60 Hz, 1000 mA maximum. The standard power adapter for the Ethernet hub plugs into a power strip located on the inside rear wall of the rack.

Figure 3-24, Figure 3-25, and Figure 3-26 illustrate how Ethernet hubs are cabled to the L2 controllers in various system sizes. These configurations do not show the possibility of L2 controllers in the I/O racks.

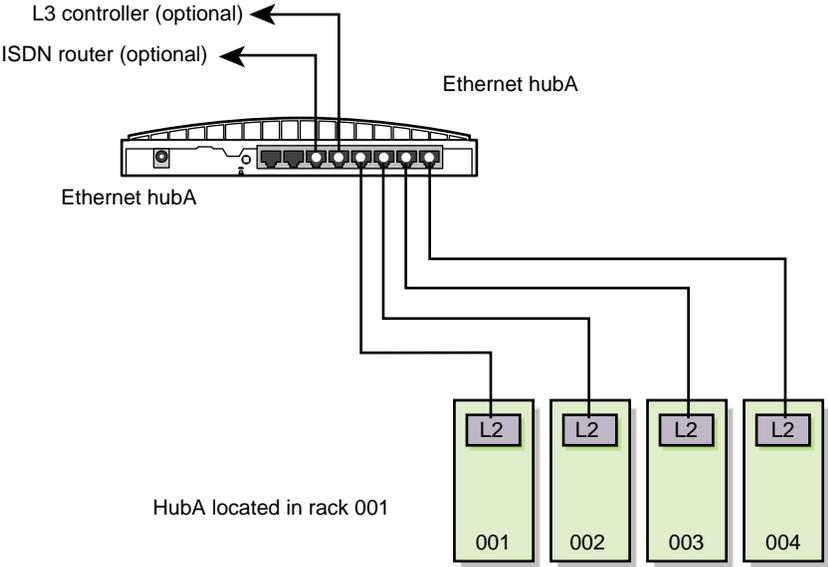


Figure 3-24 Single-hub Network in 128-processor System

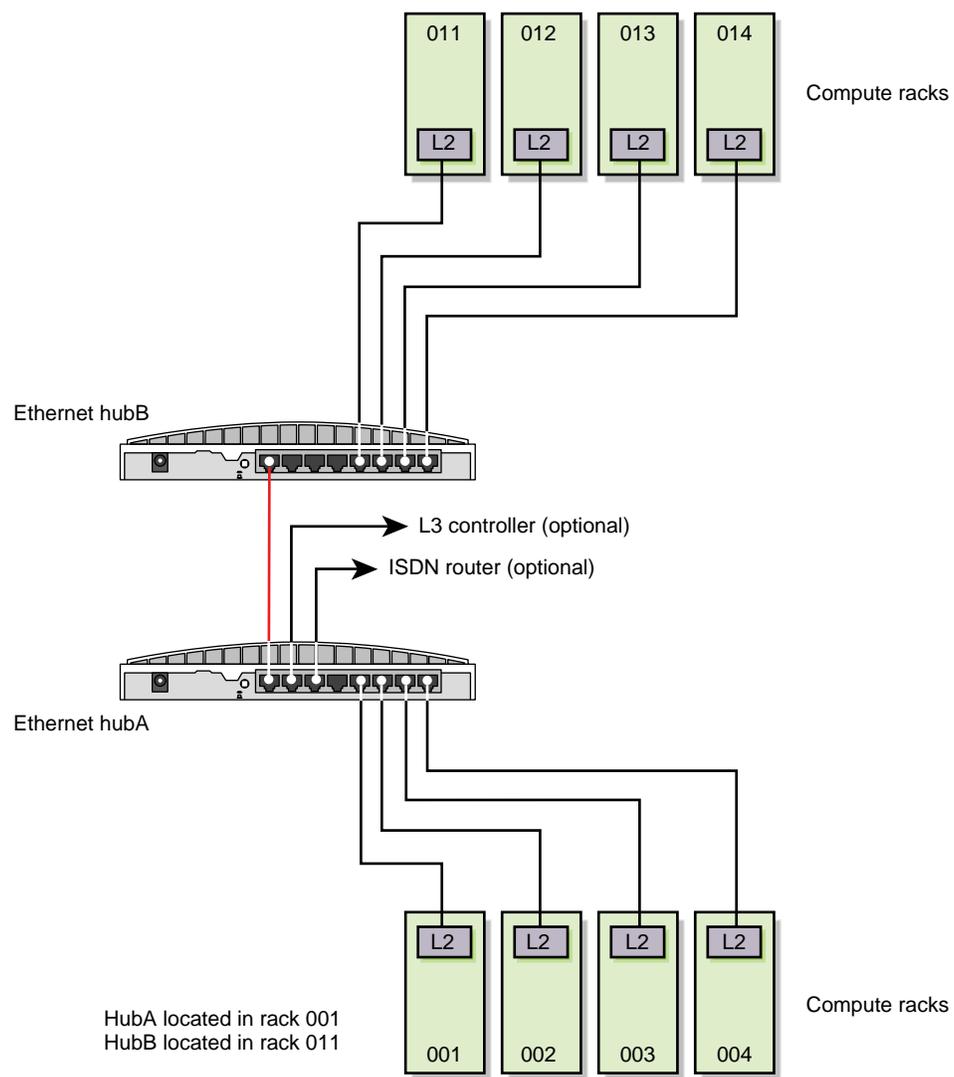


Figure 3-25 Multiple-hub Network in 256-processor System

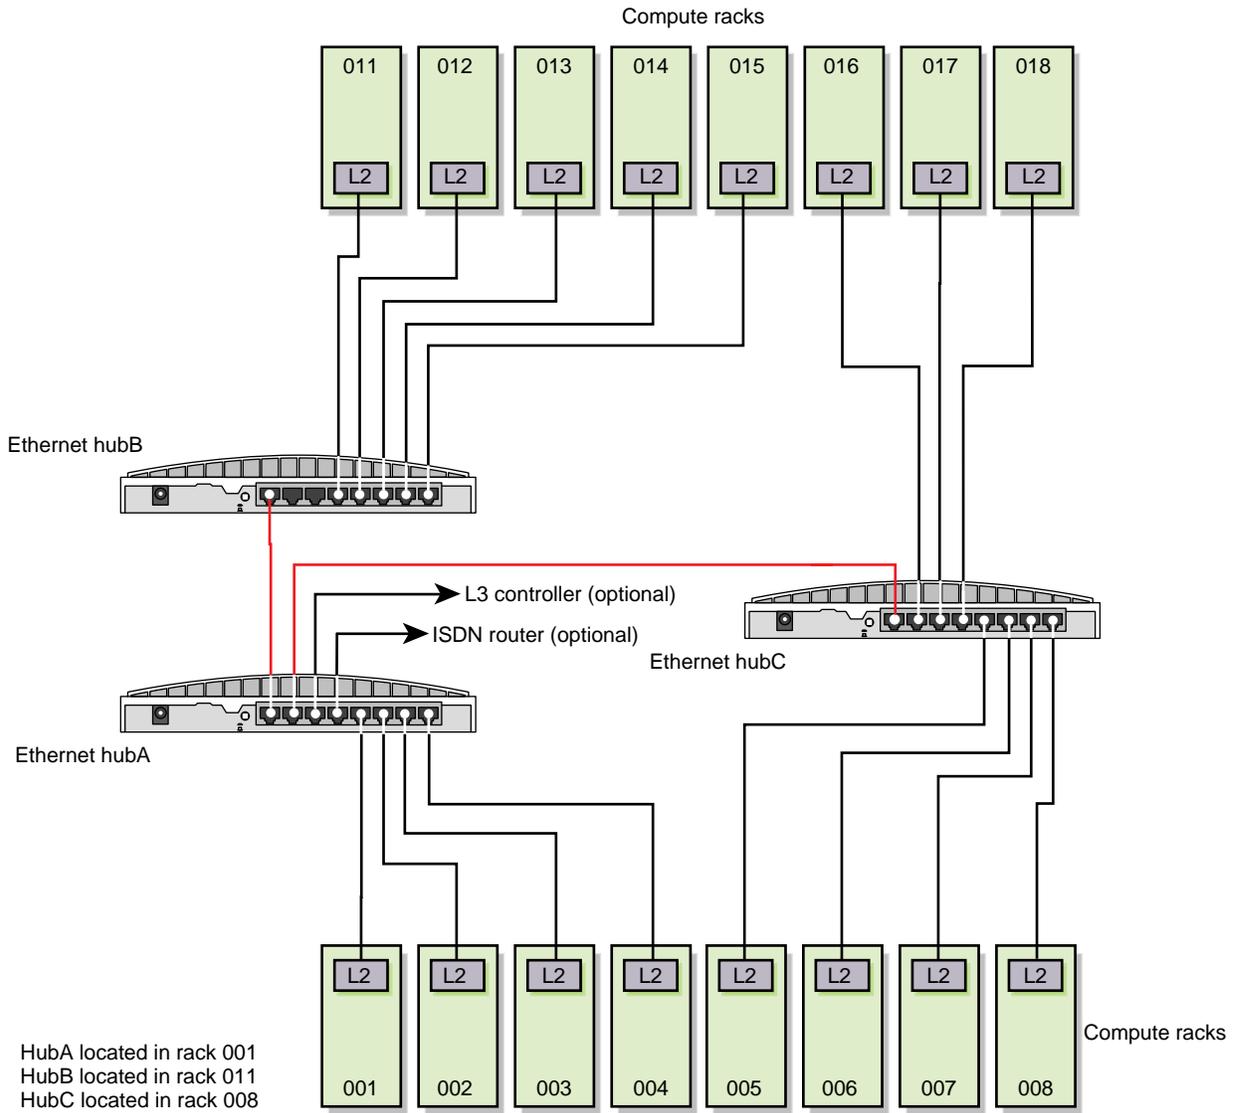


Figure 3-26 Multiple-hub Network in 512-processor System

3.16 ISDN Router

An optional ISDN router can be added to the Ethernet hub network to allow a secure remote connection. The ISDN router is located on the utility shelf in rack 001 and is connected to the Ethernet hub via an Ethernet cable. The physical dimensions of the router are 2.0 in. (5.1cm) H × 9.7 in. (24.6 cm) W × 8.3 in. (21.1 cm) D. The router weighs 1.45 lbs.



Figure 3-27 ISDN Router - Rear View

3.17 System Cooling

All SGI SNIA 3000 series systems are air-cooled. Each brick has fans to provide cooling. Air flows from the front of the brick to the rear.

4. SGI SNIA 3000 Series Performance and Bandwidth Characteristics

This section provides channel bandwidth tables for each of the bricks to help you create a configuration that meets customer requirements for performance, capacity and connectivity. When you configure a system for performance, evaluate the sustained bandwidths of the compute nodes (C-bricks) and the I/O nodes (I-, P- bricks) to determine the sustainable channel bandwidth.

4.1 C-brick (Compute Node)

Each C-brick consists of two or four Intel Itanium™ processors running at 733MHz or 800MHz, each with a 2MB or 4MB secondary cache. Each processor is capable of executing four floating-point instructions per cycle, which supports a peak speed of 3.2 GFLOP/s. The peak and sustained bandwidths for each of the C-brick's channels are listed in Table 4-1. Refer to "Compute Node (C-brick)" on page 26 for detailed information on the C-brick.

Table 4-1 Bandwidth Characteristics of the C-brick

| Description | PEAK Bandwidth | Sustainable Bandwidth |
|--|---|---|
| Link Channel Bandwidth | 3.2 GB/s full-duplex 1.6 GB/s each direction | ~ 1420 MB/s each direction |
| Xtown2 Channel Bandwidth (600 MHz mode) | 2.4 GB/s full-duplex 1.2 GB/s each direction | ~ 1066 MB/s half duplex ~ 1744 MB/s full-duplex ~ 872 MB/s each direction |
| Xtown2 Channel Bandwidth (400 MHz mode) | 1.6 GB/s full-duplex 800 MB/s each direction | ~700 MB/s half duplex ~1120 MB/s/ full-duplex ~560 MB/s each direction |

Table 4-1 Bandwidth Characteristics of the C-brick

| Description | PEAK Bandwidth | Sustainable Bandwidth |
|------------------------------|----------------|-----------------------|
| Main Memory Bandwidth | 3200 MB/s | 3200 MB/s |
| SYSAD Bandwidth (each SYSAD) | 1600 MB/s | ~ 1400 MB/s |

4.2 D-brick (Fibre Channel Disk)

The D-brick supports a maximum of twelve 3.5-inch Fibre Channel (FC) disk drives. The twelve dual-ported disk drives connect to two Fibre Channels. The aggregate channel bandwidth of a disk Fibre Channel depends on the bandwidth capability of the FC controller and the number and type of FC disk drives on the channel. Refer to Table 4-5 for controller bandwidth values and to Table 4-2 for bandwidth values of individual drives.

Table 4-2 Bandwidth Characteristics of the D-brick

| Description | PEAK Transfer Rates |
|-------------------------------|---------------------|
| 18-GB FC Disk Drive | ~ 29 MB/s |
| 36-GB FC Disk Drive (10K RPM) | ~ 30 MB/s |
| 73-GB FC Disk Drive (10K RPM) | ~ 33 MB/s |

4.3 I-brick (System Boot with PCI Bus)

The I-brick is a Crosstalk-to-PCI based I/O subsystem. It has two 800-MB/s Xtown2 ports that connect to C-bricks. There are five PCI slots that are configured on two buses and two drive bays that support Fibre Channel disk drives. For more information about

the I-brick refer to the “System Boot (I-brick)” on page 32. Refer to Table 4-3 for the peak and sustained bandwidth values.

Table 4-3 Bandwidth Characteristics of the I-brick

| Description | Peak Bandwidth |
|---------------------------------------|--|
| Xtown2 Ports A and B | 800 MB/s |
| PCI Bus 1 Slots 1 through 3 at 33 MHz | 128 MB/s in 32-bit mode 256 MB/s in 64-bit mode |
| PCI Bus 2 slots 1 and 2 at 66 MHz | 256 MB/s in 32-bit mode 512 MB/s in 64-bit mode |

4.4 P-brick (PCI Expansion)

The P-brick has two 1200-MB/s Xtown2 ports that connect to C-bricks. There are 12 PCI slots that are configured on six buses. For more information about the P-brick, refer to the “PCI Expansion (P-brick)” on page 35. Refer to Table 4-4 for the peak and sustained bandwidth values.

Table 4-4 Bandwidth Characteristics of the P-brick

| Description | Peak Bandwidth |
|--|--|
| Xtown2 Ports A and B | 2 modes software selectable 800 MB/s or 1200 MB/s |
| PCI Bus 1-6 at 66MHz (two slots per bus) | 256 MB/s in 32-bit mode 512 MB/s in 64-bit mode |

4.5 Supported PCI Cards

The PCI cards listed in Table 4-5 are supported in the I, and P-bricks for First Customer Ship (FCS). Additional PCI cards will be added to this table as they qualify.

Table 4-5 Supported PCI Cards

| Marketing Code | Vendor's P/N | PCI Bus Clock / Data | Description | Bandwidth (Peak) (per port) |
|----------------------|--------------------------------|----------------------|---|------------------------------------|
| SCSI | | | | |
| PCI-SCSI-DF-2P | QLogic QLA1240D | 33-MHz / 64-bit | Two-port ultra SCSI high voltage differential | 40 MB/s (wide) 20 MB/s (narrow) |
| PCI-SCSI-U3-2P | QLogic QLA12160 | 66-MHz / 64-bit | Two-port ultra3 SCSI low voltage differential | 160 MB/s 80 MB/s 40 MB/s |
| Fibre Channel | | | | |
| PCI-FC-1POPT-A | QLogic QLA2200F/66 | 66-MHz / 64-bit | Single-port 1-Gb Fiber Channel with fibre optic cable | 100 MB/s |
| PCI-FC-1PCOP-A | QLogic QLA2000/66 | 66-MHz / 64-bit | Single-port 1-Gb Fibre Channel with copper cable | 100 MB/s |
| Networking | | | | |
| PCI-GIGENET-OR | Alteon ACEnic 1000-SX | 66-MHz / 64-bit | Single-port Gigabit Ethernet Fiber optic | 125 MB/s |
| PCI-GIGENET-C | Alteon ACEnic 10/100/1000BaseT | 66-MHz / 64-bit | Single-port Gigabit Ethernet card 10/100/1000BaseT Unshielded Twisted Pair | 125 MB/s |

Table 4-5 Supported PCI Cards

| Marketing Code | Vendor's P/N | PCI Bus Clock / Data | Description | Bandwidth (Peak) (per port) |
|----------------|-----------------------------|----------------------|----------------------|-----------------------------|
| PCI-ATMOC3-1P | Fore ForeRunnerHE 155 | 33-MHz/ 64-bit | Single-port ATM OC3 | 19.3 MB/s |
| PCI-ATMOC12-1P | Fore 622 Hot Links | 66-MHz / 64-bit | Single-port ATM OC12 | 78 MB/s |

Note: The number of peripherals required to saturate a controller depends on the transfer rate of the disk device.

Table 4-6 PCI Card - Brick and System Maximum Quantities

| Marketing Code | Minimum Required Linux Level | Maximum Cards per Brick | Maximum Cards per System (Fully Tested) | Maximum Cards per System (Goal - not tested) |
|---|------------------------------|--|---|--|
| Storage Controllers | | | | |
| <p>Note: The system maximum numbers for the following four cards must be taken as an aggregate. Each card is assigned a weighted value and the combined total weighted value of these four cards cannot exceed the (fully tested) maximum values listed in the table. Each card has a weighted value of 1 except PCI-SCSI-LVD-2P which has a weighted value of 2. The system FC controller, located in the I-brick, is not counted when determining the maximum number of cards.</p> | | | | |
| PCI-SCSI-DF-2P | N/A | I-brick - 2 per bus (04 total) P-brick -2 per bus (12 total) | N/A | N/A |
| | N/A | I-brick - 2 per bus (04 total) P-brick -2 per bus (12 total) | N/A | N/A |
| | N/A | I-brick - 2 per bus (04 total) P-brick -2 per bus (12 total) | N/A | N/A |
| PCI-SCSI-U3-2P | N/A | I-brick - 2 per bus (04 total) P-brick - 2 per bus (12 total) | N/A | N/A |
| | N/A | I-brick - 2 per bus (04 total) P-brick - 2 per bus (12 total) | N/A | N/A |
| | N/A | I-brick - 2 per bus (04 total) P-brick - 2 per bus (12 total) | N/A | N/A |
| PCI-FC-1POPT-A | N/A | I-brick - 2 per bus (04 total) P-brick -2 per bus (12 total) | N/A | N/A |
| | N/A | I-brick - 2 per bus (04 total) P-brick -2 per bus (12 total) | N/A | N/A |
| | N/A | I-brick - 2 per bus (04 total) P-brick -2 per bus (12 total) | N/A | N/A |
| PCI-FC-1PCOP-A | N/A | I-brick - 2 per bus (04 total) P-brick -2 per bus (12 total) | N/A | N/A |
| | N/A | I-brick - 2 per bus (04 total) P-brick -2 per bus (12 total) | N/A | N/A |
| | N/A | I-brick - 2 per bus (04 total) P-brick -2 per bus (12 total) | N/A | N/A |

Table 4-6 PCI Card - Brick and System Maximum Quantities

| Marketing Code | Minimum Required Linux Level | Maximum Cards per Brick | Maximum Cards per System (Fully Tested) | Maximum Cards per System (Goal - not tested) |
|--|------------------------------|--|---|--|
| Networking Controllers | | | | |
| <p>Note: The system maximum numbers for the two GIGENET cards must be taken as an aggregate. The sum of the combined total of these two cards must not exceed the (fully tested) maximum values listed in this table. Also the combined total of the ATM cards must be taken as an aggregate.</p> <p>ATM Note: The PCI-ATMOC12-1P can not share the bus with any form of 32-bit addressable card. The PCI-ATMOC3-1P can not share the bus with any 32 bit addressable card other than another PCI-ATMOC3-1P.</p> | | | | |
| PCI-GIGENET-OR | N/A | I-brick - 2 on bus 2 (02 total) P-brick -2 per bus (10 total) | N/A | N/A |
| | N/A | I-brick - 2 on bus 2 (02 total) P-brick -2 per bus (10 total) | N/A | N/A |
| PCI-GIGENET-C | N/A | I-brick - 2 on bus 2 (02 total) P-brick -2 per bus (10 total) | N/A | N/A |
| | N/A | I-brick - 2 on bus 2 (02 total) P-brick -2 per bus (10 total) | N/A | N/A |
| PCI-ATMOC3-1P (see ATM note above) | N/A | I-brick - 2 on bus 2 (02 total) P-brick -2 per bus (8 total) | N/A | N/A |
| | N/A | I-brick - 2 on bus 2 (02 total) P-brick -2 per bus (12 total) | N/A | N/A |
| PCI-ATMOC12-1P (see ATM note above) | N/A | I-brick - 1 on bus 2 (01 total) P-brick -1 per bus (06 total) | N/A | N/A |
| | N/A | I-brick - 1 on bus 2 (01 total) P-brick -1 per bus (06 total) | N/A | N/A |

4.6 Supported Tape Drives

The tape drives listed in Table 4-7 are available at first customer ship (FCS).

Table 4-7 Supported Tape Drive Subsystems

| Peripheral Product Designator | Interface Type | Storage Capacity (Mbytes) | Transfer Rate (MBytes) |
|-------------------------------|----------------|---------------------------|------------------------|
| DLT 7000 | SCSI only | 35,000 | 5.0 |
| SONY AIT-2 | SCSI only | 60,000 | 6.0 |

4.7 System Bisection Bandwidth

The minimum bandwidth of the interconnect network that separates the system into two equal halves is referred to as the bisection bandwidth. Refer to Table 4-8 for the bisection bandwidth for common system sizes.

Table 4-8 Bisection Bandwidths

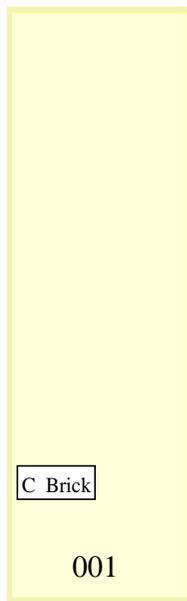
| System Models | Number of processors | System Bisection Bandwidth | Bisection Bandwidth per CPU |
|---------------|----------------------|----------------------------|-----------------------------|
| 3200 and 3400 | 8 | 3.2 Gbytes/s | 0.4 Gbytes/s |
| 3400 and 3800 | 16 | 6.4 Gbytes/s | 0.4 Gbytes/s |
| 3400 and 3800 | 32 ^{Note 1} | 6.4 Gbytes/s | 0.2 Gbytes/s |
| 3800 | 64 | 12.8 Gbytes/s | 0.2 Gbytes/s |
| 3800 | 128 | 25.6 Gbytes/s | 0.2 Gbytes/s |
| 3800 | 256 | 51.2 Gbytes/s | 0.2 Gbytes/s |
| 3800 | 512 | 102.4 Gbytes/s | 0.2 Gbytes/s |

Note 1: This configuration connects the two routers together with two NUMalink cables.

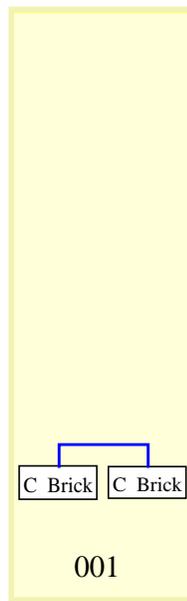
5. System Interconnection Fabric Drawings

The following figures illustrate the system interconnection fabric for various system configurations of the SGI SNIA 3000 series systems.

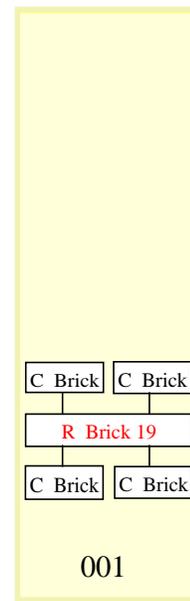
5.1 2-processor through 16-processor Systems



2 or 4 CPUs



4 to 8 CPUs



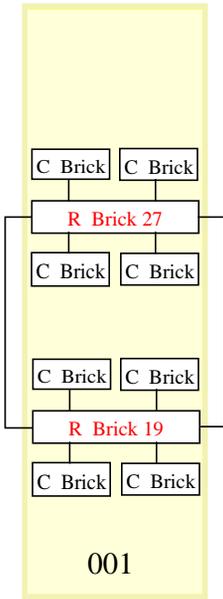
8 to 16 CPUs

| Quantities | |
|------------|------------|
| 1 Rack | 0 R bricks |
| 0 Cables | 1 C brick |

| Quantities | |
|------------|------------|
| 1 Rack | 0 R bricks |
| 1 Cables | 2 C bricks |

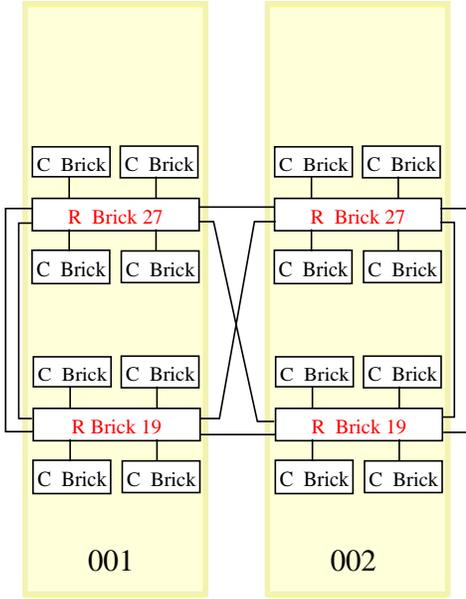
| Quantities | |
|------------|------------|
| 1 Rack | 1 R bricks |
| 4 Cables | 4 C bricks |

5.2 16-processor through 64-processor Systems



16 to 32 CPUs

| Quantities | |
|------------|------------|
| 1 Rack | 2 R bricks |
| 10 Cables | 8 C bricks |



64 CPUs

| Quantities | |
|------------|-------------|
| 2 Racks | 4 R bricks |
| 24 Cables | 16 C bricks |

6. System Partitioning

6.1 Partitioning Overview

Partitioning is defined as the ability to take a multinode distributed shared memory (DSM) system, such as an SGI SNIA 3000 series system, and divide it into a collection of smaller systems.

The two primary characteristics of partitioning are:

- The ability to run individual partitions. Each partition runs its own operating system kernel and behaves as a standalone system, without having to physically re-cable the system. A partition can be booted, powered down/up, and rebooted without effecting the normal operation of other partitions.
- The partitions are tightly coupled through the system's interconnection network (NUMALink 3) as a low latency/ high bandwidth interconnect. A failure that causes a kernel in one partition to crash will not cause a kernel in another partition to crash.

Partitioning can be thought of as "clustering in a box", or a tightly coupled cluster.

Table 6-1 Partitioning Support Timeline

| Partitioned System Size | Planned Availability | Max. Partition Size |
|-------------------------|----------------------|---------------------|
| 64 Processor | 1Q02 | 32P |
| 128 Processor | 2Q02 | 32P |
| 512 Processor | 3Q02 | 64P |

6.2 Definition of Terms

There has been much discussion on the terminology used to describe the different divisions of a system. The following definitions are not final, but are used as a starting point.

| TERM | DESCRIPTION |
|--------------|--|
| array | One or more "hosts" linked together using the array software. It may be used to link together one or more partitions, but this is TBD. |
| array node | An element known to the array management software. Currently an array node is a "host". |
| domain | A host which is a subset of a system. This is equivalent to what Amdahl and now Sun calls a domain. We may/probably want a different name. |
| host | A single Linux image running on a partition. |
| partition | A subset of a system that runs an Linux kernel and is protected by hardware barriers. A partition is identified by a Partition ID. The terms partition, host, and domain all have the same meaning, and can be used interchangeably in this context. |
| partitioning | Breaking a large physical machine into two or more smaller machines. |
| partition ID | A 16-bit quantity that identifies a particular partition. It is essentially a new name space. |
| system | A collection of hardware that supports shared memory running one or more partitions. |

6.3 Partitioning Rules

Partitioning rules define the set of valid configurations for a partitioned system. The partitioning rules and set of valid configurations are still subject to change.

Fault isolation is one of the major reasons for partitioning a system. A software or hardware failure in one partition should not cause a failure in another partition. To accomplish this, the following restrictions are placed on partitions:

- The minimum granularity for a partition is one C-brick (with its own power supply setup). On SGI SNIA 3000 series systems, this means four processors (assuming fully populated modules) is the minimum level of hardware isolation. There is no requirement for a module to be fully populated with processors.
- Each partition must have the infrastructure to run as a standalone system. This infrastructure includes a system disk and console connection.
- An I/O brick belongs to the partition that the attached C-brick belongs to. If an I/O brick is attached to two C-bricks, both C-bricks must be in the same partition. I/O bricks cannot be shared by two partitions. Peripherals, such as dual-ported disks, can be shared the same way two nodes in a cluster can share peripherals.
- Partitions must be contiguous in the topology (for example, the route between any two nodes in the same partition must be contained within that partition - and not route through any other partition). This allows intra-partition communication to be independent of other partitions.
- Partitions should not divide metarouters that are only connected to other metarouters.
- Partitions must be fully interconnected. That is to say, for any two partitions, there is a direct route between those partitions without passing through a third. This is required to fulfill true isolation of a hardware or software fault to the partition in which it occurs.
- When the full system is greater than 64 C-bricks (256 processors), it runs in coarse mode. In coarse mode the minimum partition size is four C bricks (16 processors).

6.4 Valid Configurations

Table 6-1 lists the valid system partitions for the SGI SNIA 3000 systems.

Table 6-1 Valid System Partitioning

| Number of C-bricks in the System | Number of Partitions | | Number of C-bricks in each Partition | Maximum Number of Processors in each Partition |
|----------------------------------|--|----------------------|---------------------------------------|--|
| 1 | 1 partition | with | 1 C-brick | 4 processors |
| 2 | 1 partition | with | 2 C-bricks | 8 processors |
| | 2 partitions | with | 1 C-brick | 4 processors |
| 4 | 1 partition | with | 4 C-bricks | 16 processors |
| | 1 partition | with | 2 C-bricks | 8 processors |
| | 2 partitions | with | 1 C-brick | 4 processors |
| | 1 partition 1 partition | with with | 1 C-brick 3 C-bricks | 4 processors 12 processors |
| 8 | 2 partitions | with | 2 C-bricks | 8 processors |
| | 1 partition | with | 8 C-bricks | 32 processors |
| | 2 partitions | with | 4 C-bricks | 16 processors |
| | 1 partition 1 partition | with with | 1 C-brick 7 C-bricks | 4 processors 28 processors |
| | 4 partitions | with | 2 C-bricks | 8 processors |
| | 8 partitions | with | 1 C-brick | 4 processors |
| | 1 partition 2 partitions | with with | 4 C-bricks 2 C-bricks | 16 processors 8 processors |
| | 1 partition 1 partition 2 partitions | with with with | 4 C-bricks 2 C-bricks 1 C-brick | 16 processors 8 processors 4 processors |
| | 1 partition 4 partitions | with with | 4 C-bricks 1 C-brick | 16 processors 4 processors |

Table 6-1 Valid System Partitioning

| Number of C-bricks in the System | Number of Partitions | | Number of C-bricks in each Partition | Maximum Number of Processors in each Partition |
|----------------------------------|----------------------|-------------|--------------------------------------|--|
| 16 | 1 partition | with | 16 C-bricks | 64 processors |
| | 2 partitions | with | 8 C-bricks | 32 processors |
| | 4 partitions | with | 4 C-bricks | 16 processors |
| | 1 partition | with | 1 C-brick | 4 processors |
| | 1 partition | with | 15 C-bricks | 60 processors |
| | 1 partition | with | 8 C-bricks | 32 processors |
| | 2 partitions | with | 4 C-bricks | 16 processors |
| | 1 partition | with | 8 C-bricks | 32 processors |
| 1 partition | with | 4 C-bricks | 16 processors | |
| 2 partitions | with | 2 C-bricks | 8 processors | |
| 1 partition | with | 8 C-bricks | 32 processors | |
| 4 partitions | with | 2 C-bricks | 8 processors | |
| 8 partitions | with | 2 C-bricks | 8 processors | |
| 32 | 2 partitions | with | 16 C-bricks | 64 processors |
| | 4 partitions | with | 8 C-bricks | 32 processors |
| | 1 partition | with | 16 C-bricks | 64 processors |
| | 2 partitions | with | 8 C-bricks | 32 processors |
| | 1 partition | with | 16 C-bricks | 64 processors |
| | 1 partition | with | 8 C-bricks | 32 processors |
| | 2 partitions | with | 4 C-bricks | 16 processors |
| 1 partition | with | 16 C-bricks | 64 processors | |
| 4 partitions | with | 4 C-bricks | 16 processors | |
| 8 partitions | with | 4 C-bricks | 16 processors | |

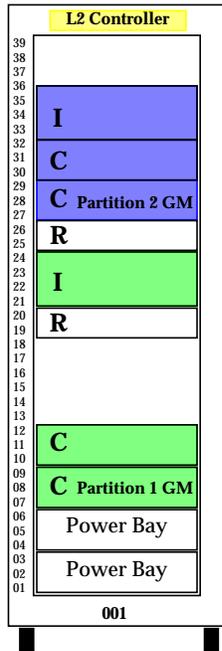
Table 6-1 Valid System Partitioning

| Number of C-bricks in the System | Number of Partitions | | Number of C-bricks in each Partition | Maximum Number of Processors in each Partition |
|----------------------------------|----------------------|------|--------------------------------------|--|
| 64 | 4 partitions | with | 16 C-bricks | 64 processors |
| | 8 partitions | with | 8 C-bricks | 32 processors |
| 128 | 8 partitions | with | 16 C-bricks | 64 processors |

Note:

1. A single entry in a table row defines a legal configuration where all partitions are the same size. Multiple entries in a table row defines a legal configuration of mixed partition sizes.
 2. Shaded partitions are not currently supported. As configurations are tested and verified, Table 6-1 will be updated to show supported partitioned configurations.
-

16 Processor System
 Partitioned as:
 Two 8 Processor Partitions



32 Processor System
 Partitioned as:
 One 4 Processor Partition
 and
 One 28 Processor Partition

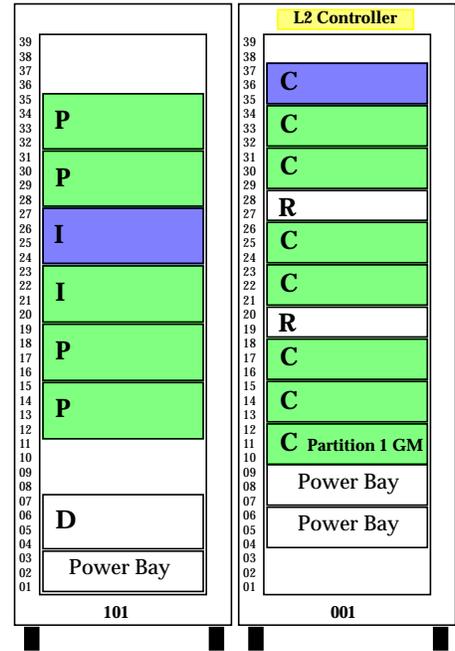


Figure 6-2 SGI SNIA 3400 Partitioned Systems

7. Brick Locations and Xtown2 Cable Configurations

This section defines the standard brick locations within the SGI Itanium™ series racks.

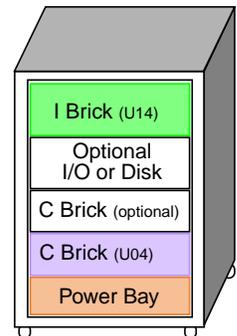
7.1 Brick Placement in Single-rack Systems

Brick placement and cabling for single-rack systems is straightforward. The first C-brick in the rack is called the global master or node 0. This C-brick is connected to an I-brick via a NUMAlink cable.

7.1.1 SGI SNIA 3200 System Brick Placement Guidelines

1. The SGI SNIA 3200 system contains:

- One or two C-bricks
 - The first C-brick (global master) is located at U04
 - The second C-brick (if configured) is located at U07
- One I-brick that is located at U14
- One optional I, P, or D-brick that is located at U10
- One power bay located at U01



17-U Rack

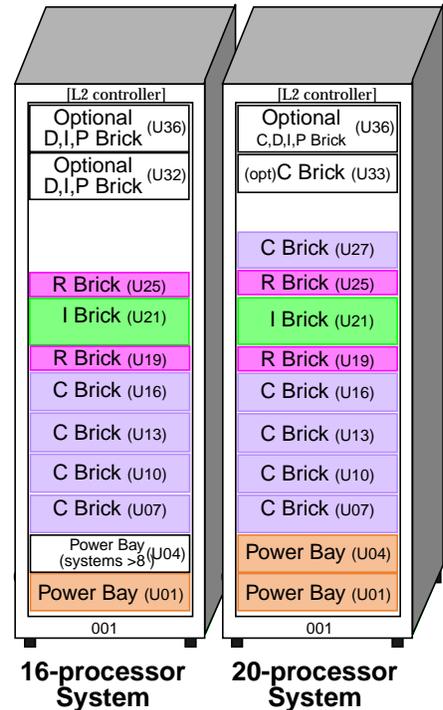
7.1.2 Single-rack SGI SNIA 3400 Brick Placement Guidelines

A single-rack SGI SNIA 3400 system contains multiple combinations of bricks and power bays. The following guidelines standardize the locations of the bricks within the rack.

The single tall rack system contains:

- A maximum of eight C-bricks that reside in the following locations:

| Qty C-brick | Location in Rack |
|-------------|--|
| One | U07 |
| Two | U07, U10 |
| Three | U07, U10, U13 |
| Four | U07, U10, U13, U16 |
| Five | U10, U13, U16, U21, U24 |
| Six | U10, U13, U16, U21, U24, U29 |
| Seven | U10, U13, U16, U21, U24, U29, U32 |
| Eight | U10, U13, U16, U21, U24, U29, U32, U35 |



- The first R-brick is located at U19 and the second R-brick is located at U25. Note: The R-brick in location U25 is not used until the number of C-bricks exceeds four.
- The first I-brick is located at U21. Up to three additional I-, or P-bricks are placed starting at U36 and working downward.
- D-bricks, if configured, are located directly above the highest I/O brick location.
- The locations of the power bays are fixed. The first power bay is located at U01. If a second power bay is required, it is located at U04.

Note: Two power bays are required for systems with more than eight 48-Vdc power connections.

7.2 Brick Placement in Multiple-rack Systems

The goal of this section is to define the brick placement and cabling for multiple-rack systems such that systems are configured in a uniform and consistent manner.

7.2.1 Multiple Rack SGI SNIA 3400 Brick Placement Guidelines

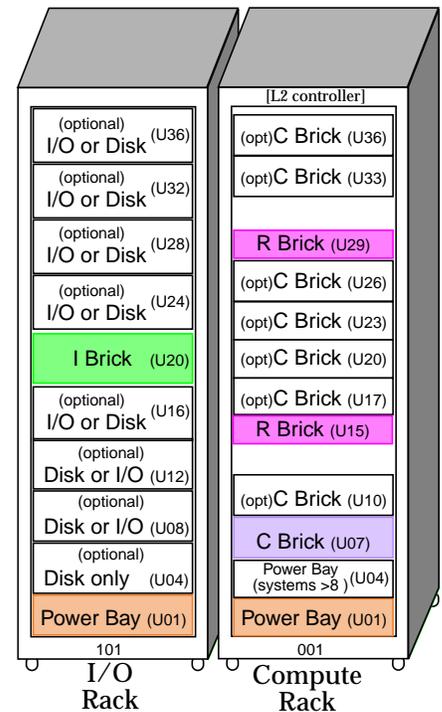
A multiple-rack SGI SNIA 3400 system contains multiple combinations of bricks and power bays. The C- and R-bricks are located in a compute rack and the I-, P-, and D-bricks are located in a separate I/O rack. In systems that require many D-bricks, the D-bricks can be placed in a separate rack(s) that contains only D-bricks. The following guidelines are provided to standardize the locations of the bricks within the rack.

The compute rack contains:

- A maximum of eight C-bricks that reside in the following locations:

| Qty C-brick | Location in Compute Rack |
|-------------|--|
| One | U07 |
| Two | U07, U10 |
| Three | U07, U10, U17 |
| Four | U07, U10, U17, U20 |
| Five | U07, U10, U17, U20, U23 |
| Six | U07, U10, U17, U20, U23, U26 |
| Seven | U07, U10, U17, U20, U23, U26, U33 |
| Eight | U07, U10, U17, U20, U23, U26, U33, U36 |

- The first R-brick is located at U15 and the second R-brick is located at U29.
- One or two power bays; the first power bay is located at U01. The second power bay is located at U04. Note: Compute racks that have more than eight 48 Vdc power connections require two power bays.



Front View

The I/O rack contains:

- One I-brick that is located at U20 of rack 101.
- Additional I-, or P-bricks that are located above or below the I-brick. The placement of the additional bricks depends on the location of the connecting C-brick. Note: In order to keep cable lengths as short as possible, the additional I-, or P-bricks are placed in the same horizontal plane as the connecting C-brick.
- Optional D-bricks; D-bricks are added to the I/O rack starting in the first location above the power bay (U04). When the lower section of the rack is filled, install any remaining D-bricks in the upper section of the rack starting at the first available location above the highest slotted I/O brick. The exact location of the additional D-bricks is dependent on the location of the I- and P-bricks.
Note: In a disk-only rack, D-bricks are added starting at U01 and grow upward.
- The I/O rack contains a maximum of one power bay which is located at U01.

7.2.2 SGI SNIA 3800 Brick Placement Guidelines

The SGI SNIA 3800 system contain multiple combinations of bricks and power bays. The C and R-bricks are located in compute racks and the I-, P-, and D-bricks are located in separate I/O racks. In systems that require many D-bricks, the D-bricks can be placed in separate D-brick rack(s). The following guidelines are provided to standardize the locations of the bricks within the rack.

The compute rack contains:

- A maximum of eight C-bricks. One compute rack must be filled before you add additional compute racks. The C-bricks use the following locations:

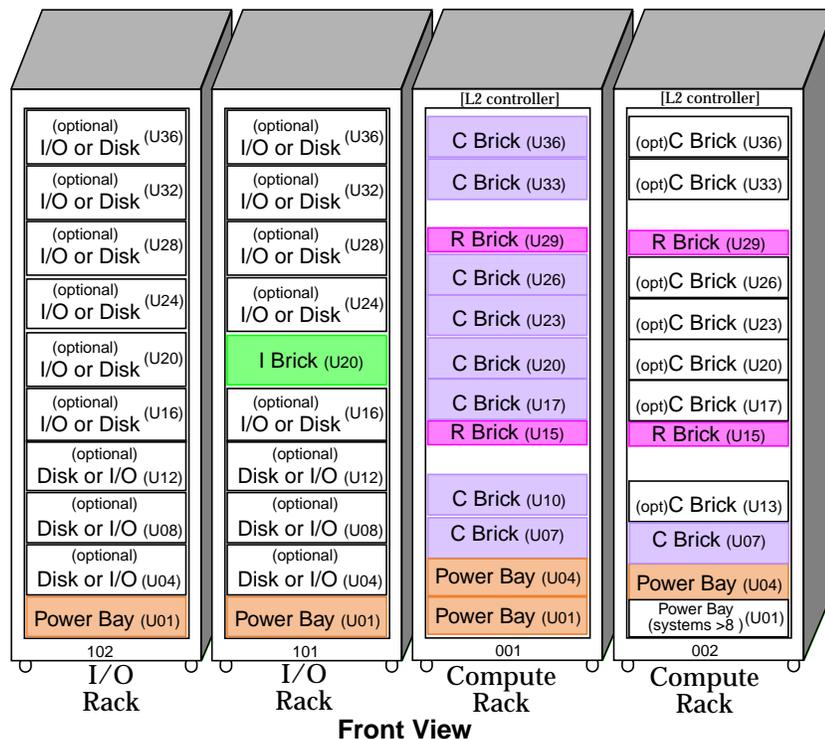
| Qty C-brick | Location in Compute Rack |
|-------------|--|
| One | U07 |
| Two | U07, U10 |
| Three | U07, U10, U17 |
| Four | U07, U10, U17, U20 |
| Five | U07, U10, U17, U20, U23 |
| Six | U07, U10, U17, U20, U23, U26 |
| Seven | U07, U10, U17, U20, U23, U26, U33 |
| Eight | U07, U10, U17, U20, U23, U26, U33, U36 |

- The first R-brick is located at U15 and the second R-brick is located at U29.
- One or two power bays; the first power bay is located at U01. The second power bay is located at U04.

Note: Compute racks that have more than eight 48-Vdc power connections require two power bays.

The I/O rack contains:

- One I-brick that is located at U20 of rack 101.
- Additional I-, or P-bricks that are located above or below the I-brick. The placement of the additional bricks depends on the location of the connecting C-brick. Note: In order to keep cable lengths as short as possible, the additional I-, or P-bricks should be placed in the same horizontal plane as the connecting C-brick.
- Optional D-bricks; D-bricks are added to the I/O rack starting in the first location above the power bay (U04). When the lower section of the rack is filled, install any remaining D-bricks in the upper section of the rack starting at the first available location above the highest slotted I/O brick. The exact location of the additional D-bricks is dependent on the location of the I- and P-bricks.
Note: In a disk-only rack, D-bricks are added starting at U01 and grow upward.
- The I/O rack contains a maximum of one power bay which is located at U01.



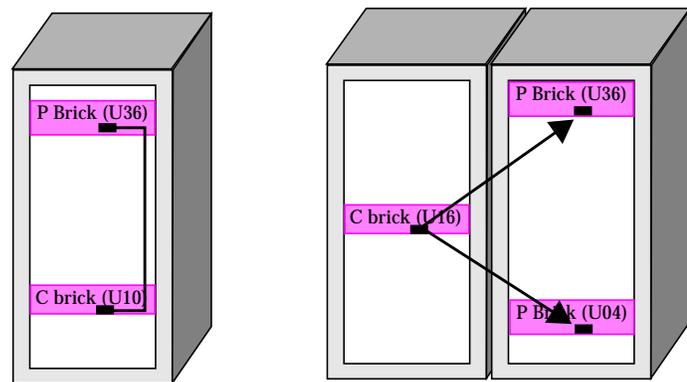
7.3 Xtown2 Cable Configurations

The NUMAlink cables are available in four lengths:

- 1 meter
- 2 meter
- 3 meter
- 4 meter

One-meter cables connect C-bricks-to-C-bricks or C-bricks-to-R-bricks within the same rack.

Two-meter cables connect SNIA 3000 C-bricks-to-I/O bricks within a rack and C bricks-to-I/O bricks between adjacent racks. A two-meter cable also connects a C-brick in any location to an I/O brick within the same rack. When connecting a C-brick to an I/O brick in an adjacent rack, the I/O brick must be within 24 U of the C-brick.



Rear View

Figure 7-1 NUMAlink Cable Connecting within a Rack and to an Adjacent Rack

Three-meter cables connect C-bricks-to-I/O bricks that span one or two full racks. If the span is a single rack, the I/O brick location must be within 24 U of the C-brick. If the span is two racks, the I/O brick location in the rack must be within 24 U of the C-brick.

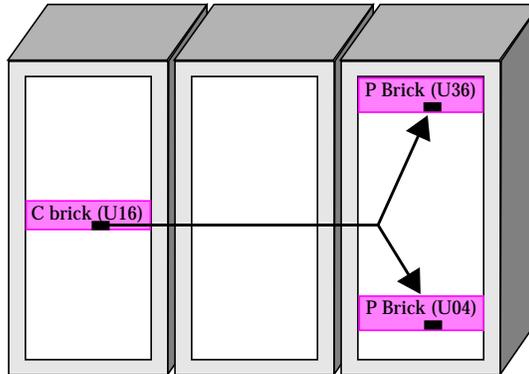


Figure 7-2 NUMAlink Cable Spanning One Full Rack

Four-meter cables connect C-bricks-to-I/O bricks that span two full racks. If the span is two racks, the I/O brick location in the rack must be within 24 U of the C-brick.

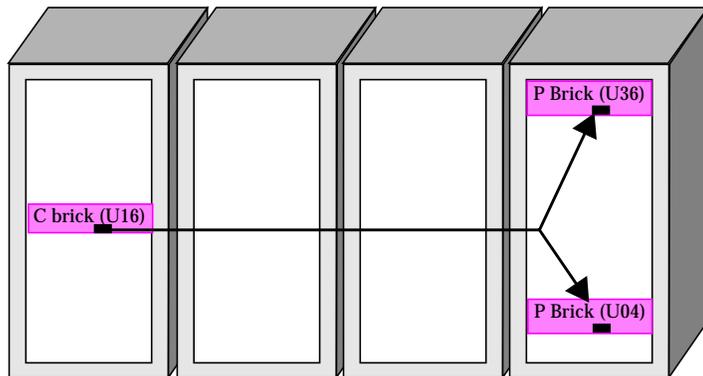
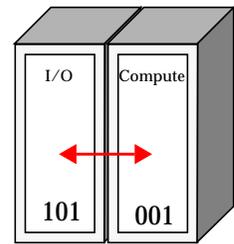


Figure 7-3 NUMAlink Cable Spanning Two Full Racks

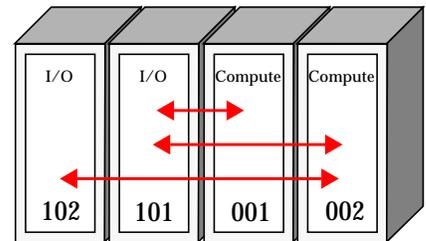
The following guidelines describe the interconnection between C-bricks and I/O bricks in multi-rack systems.

- In multiple rack systems the C- and R-bricks reside in the compute rack and the I/O bricks reside in the I/O rack. Use two-meter cables to connect the C-brick to the I/O bricks.

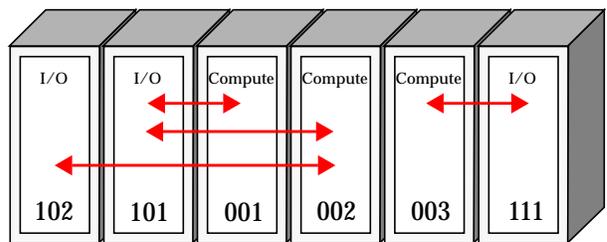


- In systems with two compute racks and two I/O racks:
 - Rack 001 connects to Rack 101; use two-meter cables
 - Rack 002 connects to Rack 101; use three-meter cables
 - Rack 002 connects to Rack 102; use four-meter cables.

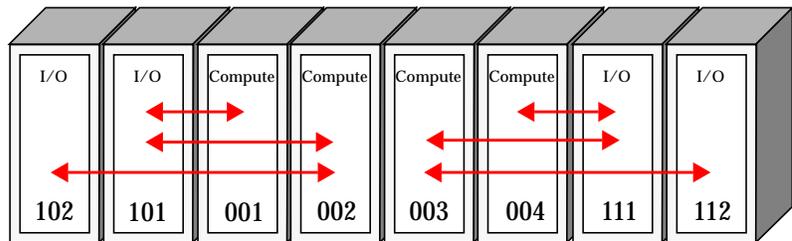
Note: If the I/O brick is within 2 U of the same location as the C-brick, then a three-meter cable can be used. For example, if the C-brick is at location U10 and the I/O brick is at location U12, a three-meter cable can be used.



- In systems with three compute racks and three I/O racks:
 - Rack 001 connects only to Rack 101; use two-meter cables
 - Rack 002 connects to Rack 101; use three-meter cables
 - Rack 002 connects to Rack 102; use four-meter cables
 - Rack 003 connects to Rack 111; use two-meter cable.



- In systems with four compute racks and four I/O racks:
 - Rack 001 connects only to Rack 101; use two-meter cables
 - Rack 002 connects to Rack 101; use three-meter cables
 - Rack 002 connects to Rack 102; use four-meter cables
 - Rack 003 connects to Rack 111; use three-meter cables
 - Rack 003 connects to Rack 112; use four-meter cables
 - Rack 004 connects to Rack 111; use two-meter cables



8. Configuration Guidelines

This section provides guidelines for configuring SGI SNIA 3000 series systems. These systems have a PCI-based I/O subsystem.

8.1 System Configuration Process

To configure a system within the SGI SNIA 3000 series, follow these five steps:

1. Determine the processor requirements.
2. Determine the memory requirements.
3. Determine the system partitioning requirements.
4. Determine the I/O subsystem requirements.
 - a. Select the disk subsystems.
 - b. Select the tape subsystems.
 - c. Determine the network bandwidth requirements.
 - Does the customer have any network protocols/interfaces that are required for interoperability with other equipment (either already owned or intended for purchase)?
 - What are the network bandwidth requirements for this customer's applications?
5. Select the physical configuration to meet the present requirements and also the future needs.

Some of the items that are determined or selected in this process are not included in the base system price. The following items may be required for a system, but must be purchased separately.

- Disk drives
- Tape drives and other storage devices
- Network communication devices
- Peripheral controller for external devices
- Additional I/O subsystems
- Additional racks
- Compilers and programming environments
- Other layered software
- A table and chair for the system console and remote support equipment

8.2 Guidelines for Configuring the C-brick

To configure a C-brick, follow these guidelines:

Select The Number Of Processors

- SGI SNIA 3200 systems can be configured with 2, 4, 6, or 8 processors.
Note: the 4-processor configuration is offered as:
One 4 processor C-brick or two C-bricks, each with two processors (2×2).
 - Only the 733 MHz (2MB cache) processor is offered with the SGI SNIA 3200 series.
 - SGI SNIA 3200 system sizes can be increased in 2-or 4-processor increments.
- SGI SNIA 3400 systems can be configured with 4, 8, 16, or 32 processors; upgrades are in 4-processor increments.
 - The 800 MHz (2MB cache) processor and the 800 MHz (4MB cache) are offered with the SGI SNIA 3400 system.
- SGI SNIA 3800 systems can be configured with 16 to 64 processors; upgrades are as follows:
 - The system sizes increase in 4-processor increments up to 64 processors.
 - The 800 MHz (2MB cache) processor and the 800 MHz (4MB cache) are offered with the SGI SNIA 3800 system.

Select the Memory

- Each DIMM pair has two banks of memory. Each of the two memory banks are split between the two DIMMs, half of each bank is located on the first DIMM and the other half of each bank is located on the other DIMM of the DIMM pair.
- Memory is added in whole-bank increments (two DIMMs).
- The DIMMs that compose a single bank of memory must be the same size.
- All DIMM pairs do not have to be the same size. For example, one DIMM pair could be 512-MB and the next DIMM pair could be 1-GB.
- Three DIMM sizes (per two DIMMs) are available:
 - 512 MB (standard directory memory)
 - 1 GB (premium directory memory)
 - 2 GB (premium directory memory)

Premium DIMMS are required for systems that have more than 128 processors.

Note: Refer to Table 9-1 for memory DIMM marketing codes.

Follow these guidelines when you install memory DIMMs in a C-brick:

- Install memory DIMMs in pairs (refer to Figure 8-1 for pair locations).
- Install the first DIMM pair in pair 0 locations (Required).
- Install the second DIMM pair in DIMM pair 1 locations (optional).
- Install the third DIMM pair in DIMM pair 2 locations (optional).
- Install the fourth DIMM pair in DIMM pair 3 locations (optional).

Refer to Figure 8-1 for DIMM pair locations.

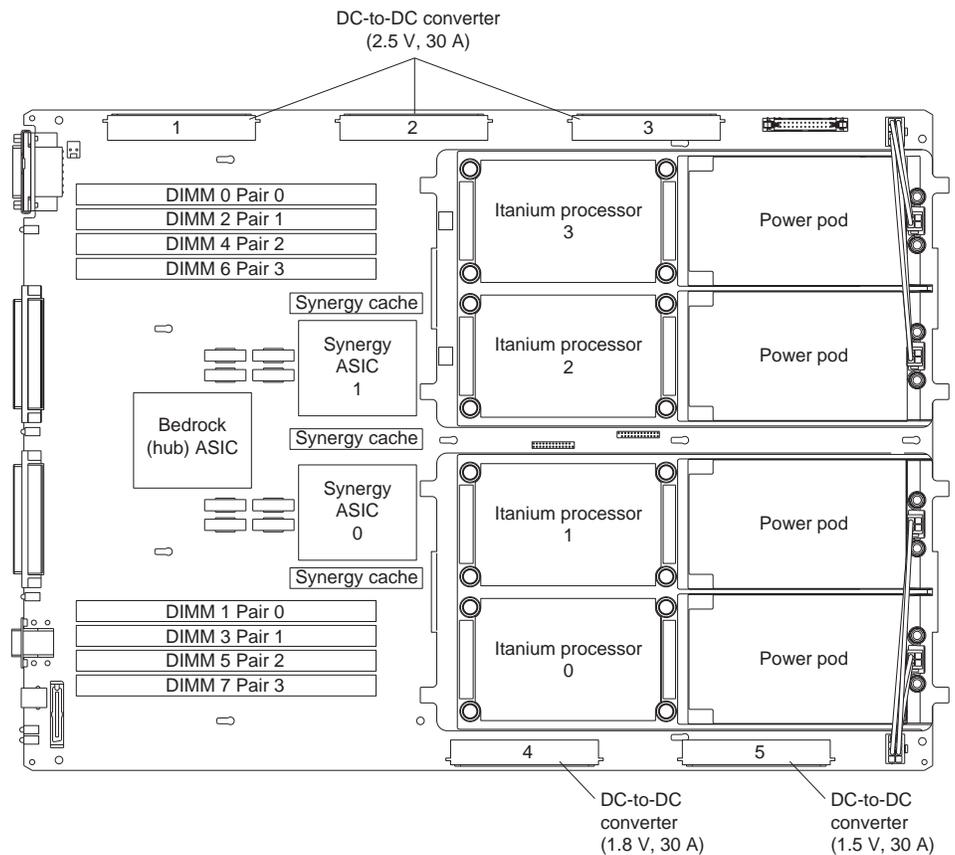


Figure 8-1 Memory DIMM and Processor Heat Sink Assy locations

Table 8-1 lists memory sizes of various DIMM combinations for a single C-brick. The number of C-bricks within a system determines the range of system memory. For example, a system with 4 C-bricks has a memory size range of 2 Gbytes to 32 Gbytes in either 512-Mbyte or 1-Gbyte increments. To ensure the best system performance, distribute the memory uniformly across all C-bricks within a system.

Table 8-1 C-brick Memory Configurations

| C Brick Memory Size | Memory per Processor (4-P C-brick) | Configuration Possibility 1 | Configuration Possibility 2 | Configuration Possibility 3 |
|------------------------------|---|------------------------------------|---|------------------------------------|
| STANDARD MEMORY SIZES | | | | |
| 512 Mbytes | 128 MBs | One MEM-512 | | |
| 1.0 Gbyte | 256 MBs | One MEM-1G | Two MEM-512 | |
| 1.5 Gbytes | 375 MBs | One MEM-1G (plus) One MEM-512 | Three MEM-512 | |
| 2.0 Gbytes | 512 MBs | Two MEM-1G | One MEM-1G (plus) Two MEM-512 | Four MEM-512 |
| 3.0 Gbytes | 750 MBs | One MEM-2G (plus) One MEM-1G | One MEM-2G (plus) Two MEM-512 | |
| 4.0 Gbytes | 1.0 GBs | Two MEM-2G | Four MEM-1G | |
| 6.0 Gbytes | 1.5 GBs | Three MEM-2G | Two MEM-2G (plus) Two MEM-1G | |
| 8.0 Gbytes | 2.0 GBs | Four MEM-2G | | |
| CUSTOM MEMORY SIZES | | | | |
| 2.5 Gbytes | 625 MBs | One MEM-2G (plus) One MEM-512 | Two MEM-1G (plus) One MEM-512 | |
| 3.5 Gbytes | 875 MBs | One MEM-2G (plus) Three MEM-512 | One MEM-2G (plus) One MEM-1G (plus) One MEM-512 | |
| 4.5 Gbytes | 1.125 GBs | Two MEM-2G (plus) One MEM-512 | | |
| 5.0 Gbytes | 1.25 GBs | Two MEM-2G (plus) One MEM-1G | Two MEM-2G (plus) Two MEM-512 | |
| 7.0 Gbytes | 1.75 GBs | Three MEM-2G (plus) One MEM-1G | | |

8.3 Guidelines for Configuring D-bricks

To configure Fibre Channel disks follow these guidelines:

There are two methods of configuring Fibre Channel disks: JBOD and RAID.

D-bricks that reside within an SGI SNIA 3000 series rack are configured as JBOD. D-bricks configured as JBOD are configured at the factory and shipped as part of the system to the customer.

Features of a D-brick:

- The D-brick is a dual-ported 4-U high brick.
- A maximum of nine D-bricks can be placed in a 39-U disk rack.
- The D-brick may be configured with 2-to-12 Fibre Channel (3.5" media) disk drives.
- Each Fibre Channel loop configured as JBOD supports up to 84 disk drives.
- The disk system performance that the customer requires determines the number of required D-bricks and the number of disk drives per D-brick. Follow these guidelines to determine the quantity of D-bricks and disk drives for a system:
 - Performance system - if the customer requires maximum bandwidth; configure the disk system with smaller capacity disks by using more D-brick enclosures. This creates more I/O channels, which increases the I/O bandwidth of the disk system.
 - Capacity system - if the customer wants large amounts of storage and bandwidth is not an issue, configure the disk system with large capacity disks that need fewer I/O channels.

The D-brick can also be ordered as part of the TP-9100 storage system. The TP-9100 is a separate 38-U high rack that holds a maximum of nine D-bricks and can be configured as JBOD and/or RAID.

When you install disk drives use the following guidelines:

- All disk drive bays must be filled with either a disk drive or a filler plate; no drive bay should be left empty.
- The disk drives labeled “1” and “2” in Figure 8-2 are required for enclosure management; these disk drives must always be present.
- Install disk drives in the order shown in Figure 8-2. The number on the disk drive indicates the order in which they are installed. For example, if four disk drives are to be installed, they would fill the bottom row of the D-brick.

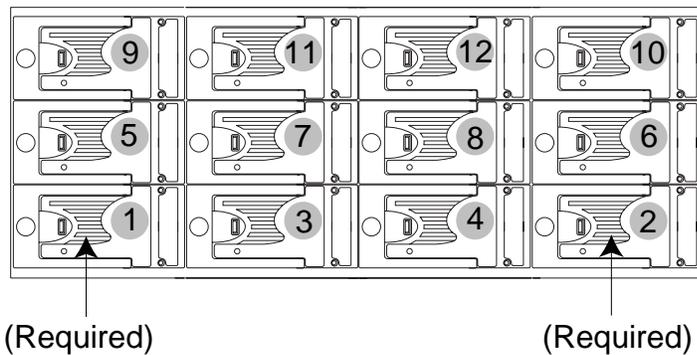


Figure 8-2 Disk Drive Installation Order

8.4 Guidelines for Determining I/O Brick Locations

Due to the flexibility that has been designed into the SGI SNIA 3000 series systems, configuring the I/O in a consistent manner offers a formidable challenge. The goals of configuring the I/O are as follows:

- Maximum I/O performance.
- Consistency; that is two systems that have the same type and quantity of parts should be configured the same way.
- Ease of manufacturing; cable routing is simple and direct.
- Ease of upgrading; bricks should be located such that it is not necessary to move them to new locations when the system is upgraded.

The first step in the I/O configuration process is to determine the type and quantity of PCI controllers required by the customer. Once you determine this, you can determine the type and quantity of I/O bricks.

The goal of placing the I/O bricks within the racks is to distribute the I/O bandwidth over the hypercube to prevent I/O traffic bottlenecks. Use the following guidelines to place the I/O bricks within the racks and connect I/O bricks to the C-bricks:

- Refer to Chapter 7 of this guide for standard I/O brick locations.
- Place the I/O brick in the same horizontal plane as the connecting C-brick.
Note: Exception is the Global master C-brick connecting to an I-brick.
- In multi-rack systems, the first I/O connections are made from rack 001 to 101.
- Refer to Chapter 6 for system partitioning rules.
- Determine which C-brick connects to which I/O brick.
 - The system I-brick connects to node 0 (global master) C-brick.
Note: There is an I-brick and an associated C-brick in each partition of a system.

Note: A router group is an R-brick and up to four C-bricks that connect to the R-brick.

- If there is more than one router group in the system, then the second I/O brick should connect to the C-brick in the lowest position of the second router group.
- Continue alternating I/O bricks between the router groups within a partition.

8.4.1 Order for Installing I/O Bricks in a Single-rack SGI SNIA 3400 System.

1. Install the I-brick in U21. Then install any additional I-bricks next.
2. Install P-bricks next.
3. Install D-bricks in the first available location above the I-brick.

8.4.2 Order for Installing I/O Bricks in an I/O Rack.

The first I-brick of the system is installed in U20 of rack 101.

If this is a partitioned system, install the additional I-bricks as follows:

- If there is only one I/O rack then, install I/O bricks above and below the I-brick in U20. Start by installing the second I-brick directly below U20.
- If there are multiple I/O racks then, install the second I-brick in U20 of I/O rack 102.

If the system is not partitioned:

- Install the additional I-bricks above and below the I-brick in U20. Start by installing the second I-brick below U20.

The goal when installing I-bricks is to position them in the middle of the rack to allow easy access to the CD-ROM.

After all the I-bricks are installed then:

4. Install the P-bricks. The P-bricks are alternately installed above and below the I-bricks.
5. Install D-bricks starting in first available position above the I/O bricks. Any remaining D-bricks are placed in the lower section of the rack starting at U04.

8.5 Guidelines for Configuring Controller Cards within I/O Bricks

The goal of this section is to provide guidelines that support consistent configuration of I/O bricks. The PCI bus slots of an I- or P-brick can support all of the various types of 3.3-volt PCI cards. No single slot is dedicated to a specific type of controller. Therefore, the guidelines serve to ensure an even distribution of bandwidth and to provide consistency from one configuration to the next.

8.5.1 General Guidelines

There are two ways to configure a system: performance and capacity. The following guidelines are directed toward performance configuration:

DO NOT

- Do not intermix different types of SCSI controllers on a single PCI bus.
- Do not intermix Fibre Channel and SCSI controllers on a single PCI bus.
- Do not mix 33-MHz and 66-MHz cards on the same 66 MHz bus (if possible).
- Do not configure PCI cards so that the bandwidth of the brick is exceeded. This guideline can be violated for capacity configurations.

DO

- When possible, intermix high and low bandwidth controllers within the same brick. This balances the bandwidth demands across bricks.
- Install cards starting with the lowest numbered bus/slot.
- Distribute I/O cards within the brick; ensure that all buses have one controller card installed before you place two controller cards on one bus.

To install I/O cards in systems with multiple I/O bricks, follow these guidelines:

- Organize the I/O adapter cards by card type. Distribute the types of adapter cards evenly among I/O bricks. For example, if there are eight Fibre Channel controllers and two P-bricks, place 4 Fibre Channel controllers in each P-brick.

8.5.2 Configuring the I-brick

To configure PCI cards in the I-brick follow these guidelines:

Installing PCI cards in small systems that contain only one I-brick is straightforward. Install only 66-MHz cards in the two 66-MHz slots, then fill the remaining slots.

- Distribute the PCI cards across the two PCI buses, by placing the high bandwidth 66-MHz cards in Bus 2.
- For partitioned systems with multiple I-bricks refer to “Partitioning Rules” on page 67.

The following are basic configuration characteristics of the I-brick:

- The I-brick has two PCI buses.
 - Bus 1 is a 33-MHz bus with three slots
 - Bus 2 is a 66-MHz bus with two slots
- If one of the two 66-MHz PCI slots contains a 33-MHz PCI card, then both slots of that 66-MHz bus will run at 33 MHz.

Note: When power is applied to a system that has a 33-MHz card and a 66-MHz card on the same 66-MHz bus, the system automatically changes the bus frequency from 66-MHz to 33-MHz. Reducing the frequency by one half reduces the bus bandwidth by one half. The actual bandwidth of the 66-MHz card does not change, however the ability of the PCI bus to support the bandwidth of the card is limited. It is important to note that a PCI bus that runs at 66 MHz would support the combined bandwidth of two high-bandwidth PCI cards, but the same two PCI cards would be bandwidth starved on a 33-MHz bus.

- The PCI bus supports 32-bit and 64-bit PCI cards on the same bus.
- The system Fiber Channel controller is standard and uses one of the five PCI slots. The standard location for the Fibre Channel controller is Bus 1 Slot 1.
- PCI cards are carrier mounted. Non-standard PCI cards that do not fit in the carrier are not supported.
- The I-brick can contain one or two disk drives; one 18-GB disk drive is standard and the second 18-GB disk drive is optional.

8.5.3 Configuring the P-brick

The following are guidelines for configuring PCI cards in the P-brick:

- Distribute PCI cards evenly among the P-bricks.
- Distribute PCI cards evenly among the slots of each P-brick.
- For maximum bandwidth, distribute the PCI cards across all six PCI buses.
- For partitioned systems with multiple P-bricks, refer to “Partitioning Rules” on page 67

The following are basic configuration characteristics of the P-brick:

- The P-brick has six 66-MHz PCI buses; each 66-MHz bus has two slots (total 12 slots)
- If one of the two 66-MHz PCI slots contains a 33-MHz PCI card, then both slots of that 66-MHz bus will run at 33 MHz.

Note: When power is applied to a system that has a 33-MHz card and a 66-MHz card on the same 66-MHz bus, the system automatically changes the bus frequency from 66-MHz to 33-MHz. Reducing the frequency by one half reduces the bus bandwidth by one half. The actual bandwidth of the 66-MHz card does not change, however the ability of the PCI bus to support the bandwidth of the card is limited. It is important to note that a PCI bus that runs at 66 MHz would support the combined bandwidth of two high-bandwidth PCI cards, but the same two PCI cards would be bandwidth starved on a 33-MHz bus.

- The PCI bus supports 32-bit and 64-bit PCI cards on the same bus.
- PCI cards are mounted on a carrier. Non-standard PCI cards that do not fit in the carrier are not supported.

8.6 Configuring the Power Bay

The power bay is 3 U high and provides eight 48-Vdc power connections. Each power bay contains a minimum of two distributed power supplies (DPSs). Each DPS provides 950 watts of continuous power. The number of power supplies required depends on the number and type of bricks in the rack.

Use the following formula to determine the number of power bays required for a single rack:

$$\text{TPB} = \text{NumCon} / 8$$

Where:

TPB is the total power bays required.

NumCon is the total number of 48-Vdc connections required in the rack.

Note: The result is always rounded up to the next integer.

Use the following formula to determine the number of DPSs required in each power bay:

$$\#\text{DPS} = 1 + (\text{Totalwatt} / 950)$$

Where:

#DPS is the number of DPSs required in a single power bay.

Totalwatt is the sum of the wattage from each device connected to the power bay.

Note: The result is always rounded up to the next integer.

Refer to Appendix H for brick power requirements.

The following components require 48V connection from the power bay:

- C-brick
- I-brick
- P-brick
- R-brick
- L2 controller

The power bay connects to the bricks with either a 1.0-meter or a 2.0-meter power cord. The power bay located in U04 of a rack uses a 1.0-meter power cord to connect to bricks in locations U04 through U14. Bricks in locations above U14 require a 2.0-meter power cord.

The L2 controller connect to the power bay at location U01 with a specially designed power cord.

The power bay located in U07 of a rack uses a 1.0-meter power cord for connecting to bricks in location U14 and below, and a 2.0-meter power cord to connect to bricks in location U14 and above.

8.6.1 Installing Distributed Power Supplies (DPS)

The DPSs are installed in the power bay starting from the left side as you face the front of the power bay. Refer to Figure 8-3.

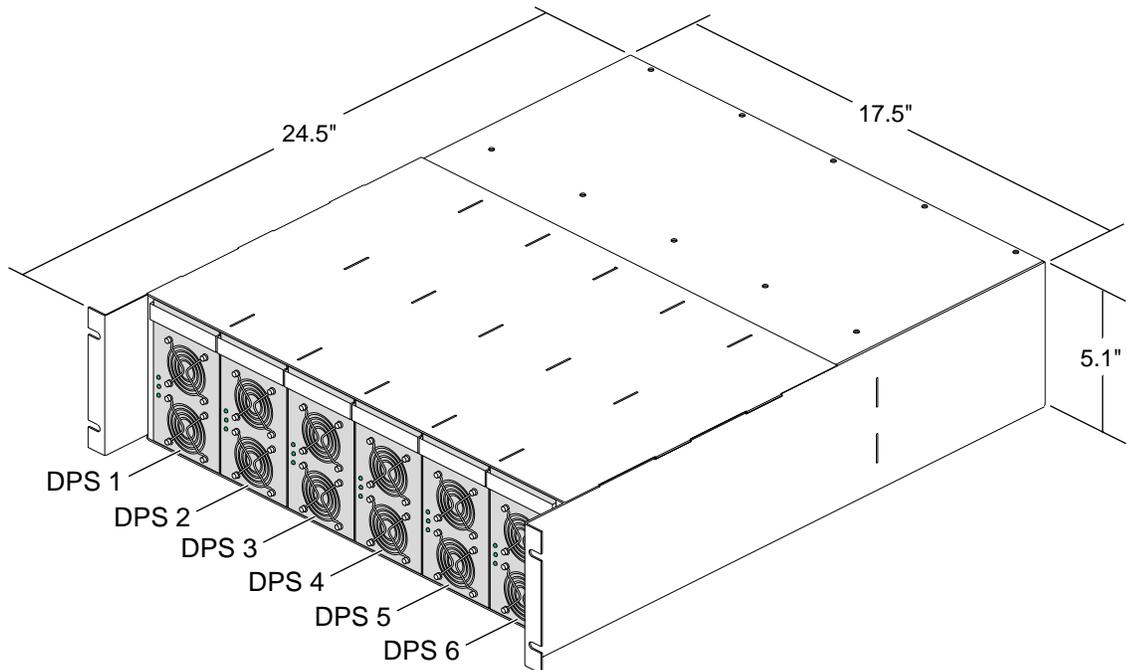


Figure 8-3 DPS Power Bay Locations

8.7 Guidelines for Configuring the System Control Network

The system control network can have the following features/enhancements added:

- Remote dumb terminal - The dumb terminal connects to the C-brick by using an RS-232 connection.
Note: If an L2 controller is present the terminal connects to it.
- L2 controller - The L2 system controller connects to the C-brick via the USB connection on the C-brick.
- Optional modem - for a remote connection, connect an optional modem to the L2 controller by using an RS-232 connection.
- Remote dumb terminal - connect an optional remote dumb terminal to the L2 controller by using an RS-232 connection.
- Optional L3 controller - connect an optional L3 controller to the L2 controller via an Ethernet connection. If an L2 controller is not present, connect the L3 controller to the C-brick via an USB connection.

Note: The Ethernet connection on an L2 or the Ethernet hub that connects L2s in a multi-rack system should never be connected to a network. Only a direct cable connection to an L3 is allowed to connect to this Ethernet port (this is a private Ethernet connection that uses a non-standard protocol).

8.8 Base System Components

8.8.1 SGI SNIA 3200 System

Four basic marketing codes define the SGI SNIA 3200 system. These marketing codes are referred to as bundles.

S-I3200-2P2C-A

- One 17-U deskside rack with trim
- One power bay with four power supplies
- One single-phase power distribution strip
- One 2-processor (733 MHz 2 MB) C-brick without memory
- One I-brick with SVGA PCI card, FC PCI card, cable, and 18-GB system disk
- One 1-meter NUMAlink cable
- System software (Linux)

S-I3200-4P2C-A

- One 17-U deskside rack with trim
- One power bay with four power supplies
- One single-phase power distribution strip
- One 4-processor (733 MHz 2 MB) C-brick without memory
- One I-brick with SVGA PCI card, FC PCI card, cable, and 18-GB system disk
- One 1-meter NUMAlink cable
- System software (Linux)

S-I3200-2X2P2C-A

- One 17-U deskside rack with trim
- One power bay with four power supplies
- One single-phase power distribution strip
- Two 2-processor (733 MHz 2 MB) C-bricks without memory
- One I-brick with SVGA PCI card, FC PCI card, cable, and 18-GB system disk
- Two 1-meter NUMALink cables
- System software (Linux)

S-I3200-8P2C-A

- One 17-U deskside rack with trim
- One power bay with four power supplies
- One single-phase power distribution strip
- Two 4-processor (733 MHz 2 MB) C-bricks without memory
- One I-brick with SVGA PCI card, FC PCI card, cable, and 18-GB system disk
- Two 1-meter NUMALink cables
- System software (Linux)

Each of the four SGI SNIA 3200 system bundles require that you select from following items to complete the system configuration:

- Memory DIMMs for the C-bricks

Note: the following items are customer configurable options:

- Customer configurable PCI cards and optional disk drive for I-brick
- One additional I, or P-brick
- Customer configurable PCI cards, if a P-brick is configured
- A D-brick and quantity of disk drives for additional disk storage
- Additional software packages
- Optional L2 or L3 controller

8.8.2 SGI SNIA 3400 System

Four basic marketing codes define the SGI SNIA 3400 system.

S-I3400-4P4C-B

- One 39-U rack with trim
- One power bay with five power supplies
- One single-phase or three-phase PDU
- One 4-processor (800 MHz 4 MB) C-brick without memory
- One R-brick
- One I-brick with SVGA PCI card, FC PCI card, cable, and 18-GB system disk
- L2 controller
- One 2-meter NUMAlink cable
- System software (Linux)

S-I3400-8P4C-B

- One 39-U rack with trim
- One power bay with five power supplies
- One single-phase or three-phase PDU
- Two 4-processor (800 MHz 4 MB) C-bricks without memory
- One R-brick
- One I-brick with SVGA PCI card, FC PCI card, cable, and 18-GB system disk
- L2 controller
- Two 1-meter NUMAlink cables
- One 2-meter NUMAlink cable
- System software (Linux)

S-I3400-16P4C-B

- One 39-U rack with trim
- One power bay with five power supplies
- One single-phase or three-phase PDU
- Four 4-processor (800 MHz 4 MB) C-bricks without memory
- One R-brick
- One I-brick with SVGA PCI card, FC PCI card, cable, and 18-GB system disk
- L2 controller
- Four 1-meter NUMAlink cables
- One 2-meter NUMAlink cable
- System software (Linux)

S-I3400-32P4C-B

- One 39-U rack with trim
- Two power bays, each power bay with five power supplies
- Two single-phase PDUs or one three-phase PDU
- Eight 4-processor (800 MHz 4 MB) C-bricks without memory
- Two R-bricks
- One I-brick with SVGA PCI card, FC PCI card, cable, and 18-GB system disk
- L2 controller
- Ten 1-meter NUMAlink cables
- One 2-meter NUMAlink cable
- System software (Linux)

Each of the four SGI SNIA 3400 system bundles require that you select from the following items to complete the system configuration:

- Select the memory DIMMs for the C-bricks

Note: the following items are customer configurable options:

- Customer configurable PCI cards for I-brick
- Optional disk drive for the I-bricks
- Additional I, or P-bricks
- Customer configurable PCI cards, if P-brick is configured
- Additional I/O racks if required (eight I/O bricks per rack)
- A D-brick and the quantity of disk drives for additional disk storage
- An additional 39-U rack if configurable space in the first rack is full
- An additional 39-U rack for disk storage
- Additional software packages
- L3 controller

8.8.3 SGI SNIA 3800 Systems

The SGI SNIA 3800 system is configured with a minimum of two 39-U racks. One rack contains the C-bricks and R-bricks (compute rack) and the other rack contains the I, P, and D-bricks (I/O rack). The SGI SNIA 3800 system can have up to sixteen compute racks; the number of racks depends on the number of C-bricks in the system. The system size ranges from 4 to 128 C-bricks. The minimum system upgrade is done in one C-brick increments.

The base compute rack has the following basic components:

SB-I3800-16P4C-B

- One 39-U rack with trim
- One power bay with five power supplies
- One single-phase or three-phase PDU
- Four 4-processor (800 MHz 4 MB) C-bricks without memory
- One R-brick
- L2 controller
- Four 1-meter NUMAlink cables
- System software (Linux)

UPG-I3800-32P4C-B

- One 39-U rack with trim
- Two power bays, each power bay with five power supplies
- Two single-phase PDUs or one three-phase PDU
- Eight 4-processor (800 MHz 4 MB) C-bricks without memory
- Two R-bricks
- L2 controller
- Twelve 1-meter NUMAlink cables
- System software (Linux)

The base I/O rack has the following components:

FTO-OS-Linux-SN

- One 39-U rack with trim
- One power bay with four power supplies
- One single-phase or three-phase PDU
- One I-brick with SVGA PCI card, FC PCI card, cable, and 18-GB system disk
- One 2-meter NUMALink cable

To complete the SGI SNIA 3800 system configuration select from the following items:

- memory DIMMs for the C-bricks

Note: the following items are customer configurable options:

- Additional 4-processor C-bricks
- Additional compute racks if required (eight C-bricks per rack)
- Additional I- or P-brick
- Customer configurable PCI cards, if P-brick is configured
- Select customer configurable PCI cards for I-brick
- Optional disk drive for the I-brick
- D-bricks and the quantity of disk drives for additional disk storage
- An additional 39-U rack if configurable space in the first rack is full
- An additional 39-U rack for disk storage
- Additional software packages
- L3 controller

9. Expanding the SGI SNIA 3000 Series System

There are three areas in which an SGI SNIA 3000 series system can be expanded: processor capacity, memory capacity, and peripheral capacity. Note that expansion in one area often affects another area.

9.1 Processor Capacity

Each system size has a maximum number of C-bricks that the system supports.

SGI 3200 systems can have C-bricks that contains two or four-processors. C-bricks that contain two processors can be upgraded by adding an two additional processor heat sink assemblies (a processor heat sink assembly contains a single processor) to each two-processor C-brick. A system with a single four-processor C-brick can be upgraded by adding an additional four-processor C-brick.

SGI 3400 systems are configured with four-processor C-bricks. SGI 3400 systems can be expanded to a maximum of 8 four-processor C-bricks. Systems with eight C-bricks may require a separate rack for the I/O components.

SGI 3800 systems are configured with four-processor C-bricks. SGI 3800 systems can be expanded to a maximum of 16 four-processor C-bricks in a SSI. However, in the future utilizing partitioning, the total system size will increase to 512 processors. Each compute rack can hold a maximum of eight C-bricks; therefore, additional racks must be configured when the number of C-bricks exceed the capacity of existing racks.

9.2 Main Memory Capacity

Memory can be expanded in single-bank increments. There are three bank sizes to choose from: 512 Mbytes, 1 Gbytes, and 2 Gbytes. Each C-brick has four memory banks that can contain any combination of DIMM pairs; however, each DIMM in a DIMM pair must be the same memory size. Table 9-1 lists the marketing codes for the three memory bank sizes.:

Table 9-1Memory Marketing Codes

| Memory Size | Marketing Code |
|-------------|----------------|
| 512 Mbytes | MEM-512 |
| 1 Gbytes | MEM-1G |
| 2 Gbytes | MEM-2G |

If a customer requires more memory, but the memory in the existing C-bricks are filled; additional memory can be added by increasing the number of C-bricks, up to the system maximum (refer to “Processor Capacity” on page 107).

9.3 Peripheral Capacity

When expanding peripheral capacity, the number of I/O bricks in the system cannot exceed the number of C-bricks in the system. If the number of I/O bricks in the system equals the number of C-bricks, the number of I/O bricks can be increased by adding an additional C-brick for each new I/O brick.

Note: Each system type has a maximum number of C-bricks that it can support.

10. SGI 3000 Series Software Requirements

10.1 Required Software

SGI SNIA 3000 series systems use the Linux operating system.

More information needed.

10.2 Optional Software

Refer to the online price book for a list of optional software. The URL for the online price book is:

<http://nafois.corp.sgi.com/psd/psdUSOnlinePb.html>

10.3 Selecting Third-party Applications

There are more than 3600 third-party software applications. Refer to the SGI external Web page for a current list of applications:

<http://www.sgi.com/Products/appsdirectory.html>

SGI SNIA 3200 System Configuration

System Definition:

- Single short rack configuration; no additional racks for processors or I/O.

Main Memory:

- Minimum memory size: 512 MB (one C-brick with one 512-MB bank)
- Maximum memory size: 16 GB (two C-bricks; each C-brick has four 2-GB banks)

Processors:

- Minimum of one C-brick (2 processors)
- Maximum of two C-bricks (8 processors)

The C-brick has the following processor configuration options:

| Processor Base Configurations | Processor Upgrade Possibilities |
|--|---|
| (qty 1) 2-P C-brick | Add (2) processor heat sink assy Add (1) 4-P C-brick Add (2) processor heat sink assy and (1) 4-P C-brick |
| (qty 1) 4-P C-brick | Add (1) 4-P C-brick |
| (qty 2) 2-P C-bricks | Add (2) processor heat sink assy Add (4) processor heat sink assy |
| (qty 1) 4-P C-brick (qty 1) 2-P C-brick | Add (2) processor heat sink assy |
| (qty 2) 4-P C-bricks | None |

Input/output:

- Maximum of two I/O bricks; one of the I/O bricks must be an I-brick

Disk:

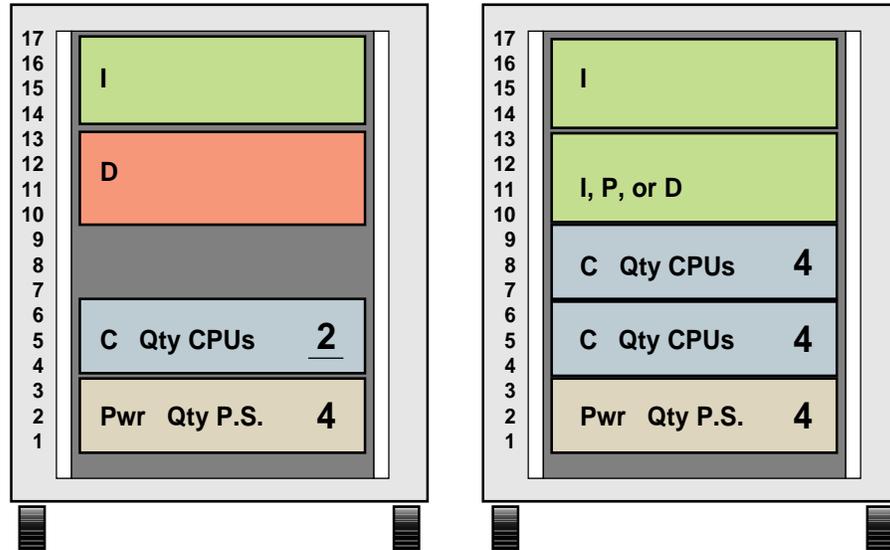
- Maximum of one D-brick in the base rack
- 1 to 2 additional D-brick racks as required

Miscellaneous:

- One power bay
- One L2 controller (optional)

Note: The R-brick is not used in SGI SNIA 3200 systems.

SGI SNIA 3200 Systems



Note: Disks can be placed in a separate SGI or third-party disk rack.

SGI SNIA 3400 System Configuration

System Definition:

- Single tall rack configuration; additional racks for I/O and disks as required

Main Memory:

- Minimum memory size 512 MB: (one C-brick with one 512-MB bank)
- Maximum memory size 64 GB: (eight C-bricks, each C-brick has four 2-GB banks)

Processors:

- Minimum of one C-brick (4 processors)
- Maximum of eight C-bricks (32 processors)

The C-brick has the following processor configuration options:

| Processor Base Configurations | Processor Upgrade Possibilities |
|---|---|
| One 4-processor C-brick and one R-brick | Add one 4-processor C-brick Add Two 4-processor C-bricks Add three 4-processor C-bricks |
| Two 4-processor C-bricks and one R-brick | Add two 4-processor C-bricks |
| Four 4-processor C-bricks and one R-brick | N/A |
| Eight 4-processor C-bricks and two R-bricks | N.A |

Input/output:

- Minimum of one I-brick; maximum of eight I/O bricks

Disk:

- D-bricks may be included in the base rack
- Additional D-brick rack(s) for disks as required

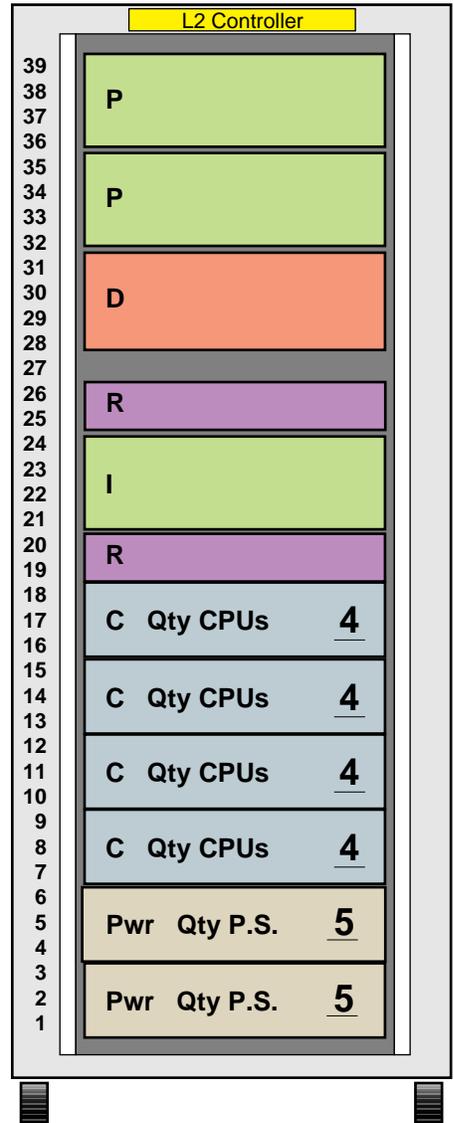
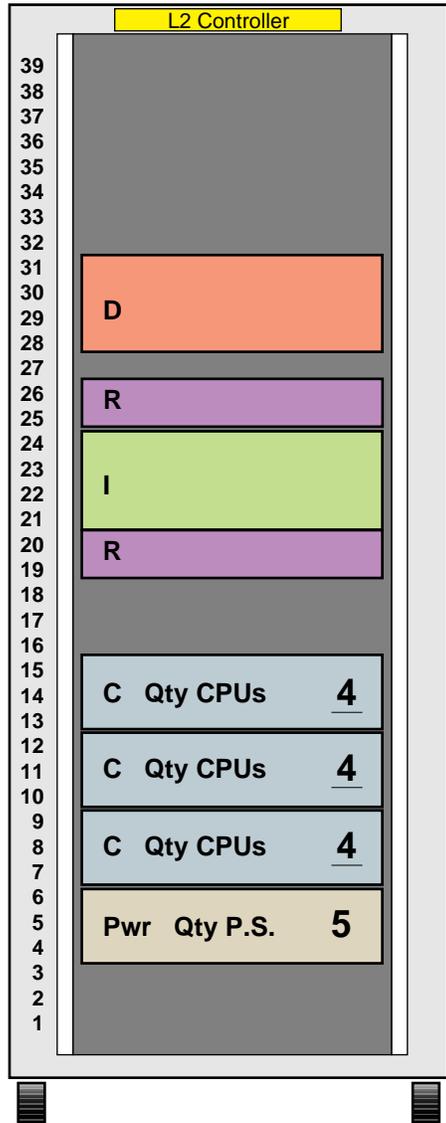
Miscellaneous:

- Maximum of two R-brick-16s
- Minimum of 1 power bay, maximum of 2 power bays

12-processor System

16-processor System

SNIA 3400
Systems



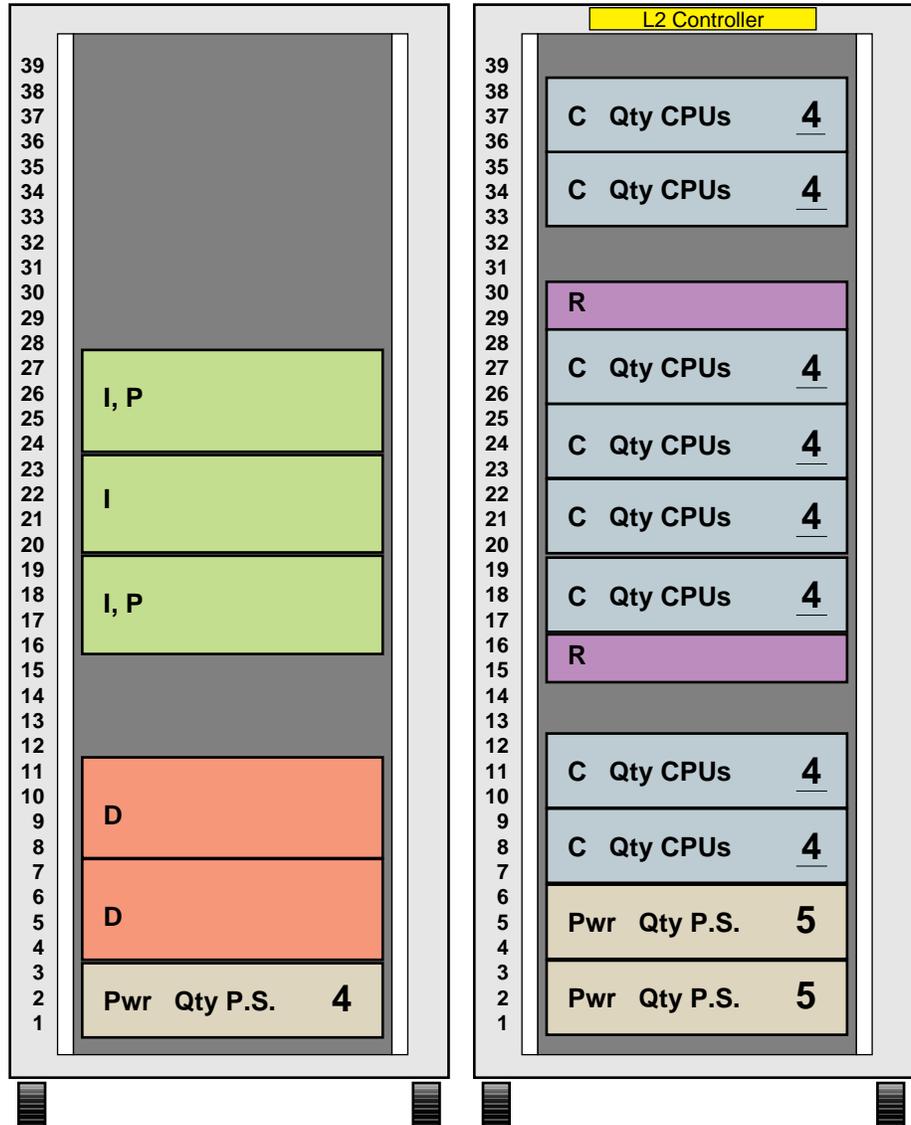
Note:
Additional D-brick racks
not shown

FRONT VIEW

32-processor System

SGI SNIA 3400 System

Note:
Additional racks with
D bricks not shown



FRONT VIEW

SGI SNIA 3800 System Configuration

System definition:

Main Memory:

- Minimum SSI memory size: 2 GB (four C-bricks, each with one 512-MB bank)
- Maximum SSI memory size: 64 GB (16 C-bricks, each with four 2-GB banks)
- Maximum system memory size: 1024 GB (128 C-bricks, each with four 2-GB banks)
- System sizes above 128 processors require premium DIMMs

Processors:

- 16 to 512 processors (16P to 64P increased in increments of 4 processors)

Rack Configurations:

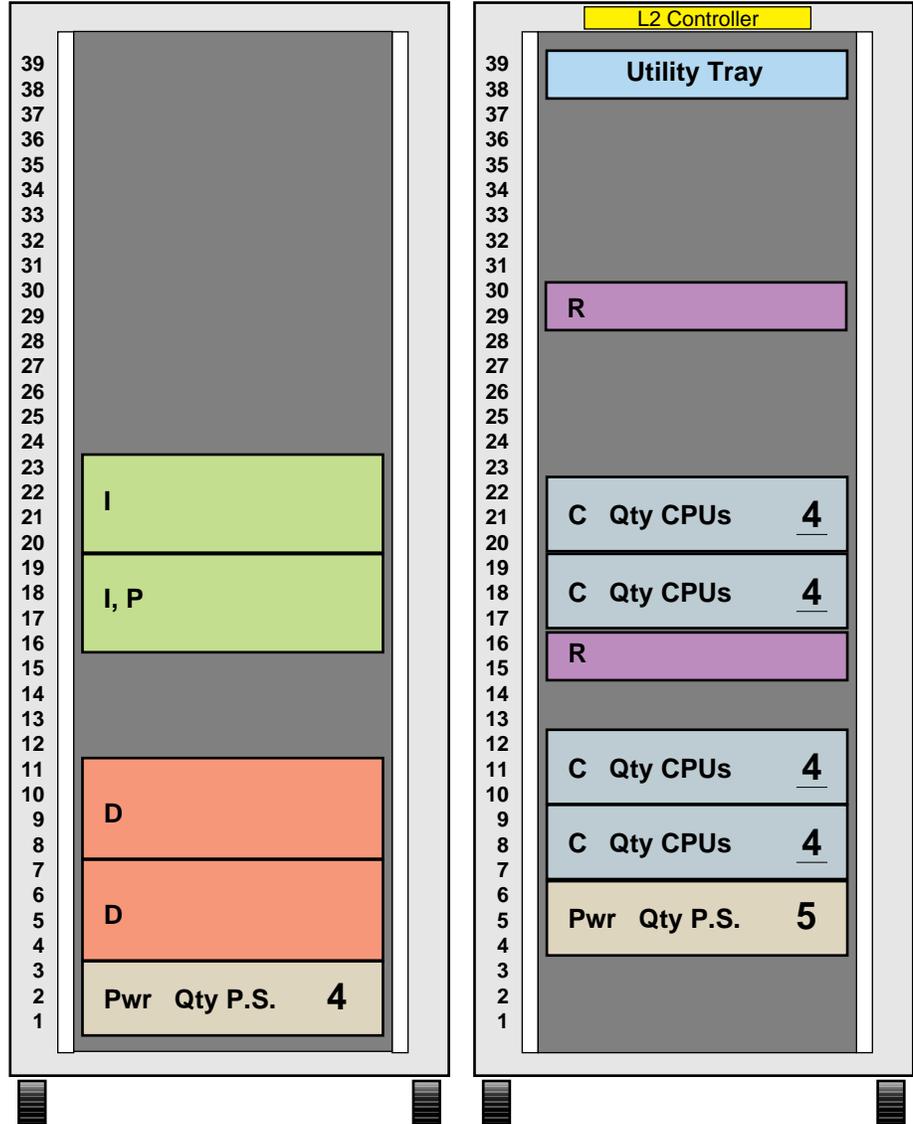
- Compute racks contain only C-bricks and R-bricks
- Maximum of two power bays in compute racks
- Maximum of sixteen compute racks
- Maximum of eight C-bricks per rack
- Maximum of four R-bricks per rack
- I/O racks can contain both D-bricks and I/O bricks
- Maximum of eight I/O bricks per I/O rack
- Maximum of one power bay in I/O racks
- Maximum of nine D-bricks per rack (no power bays in rack)

Miscellaneous:

- One L2 controller is required for all compute racks

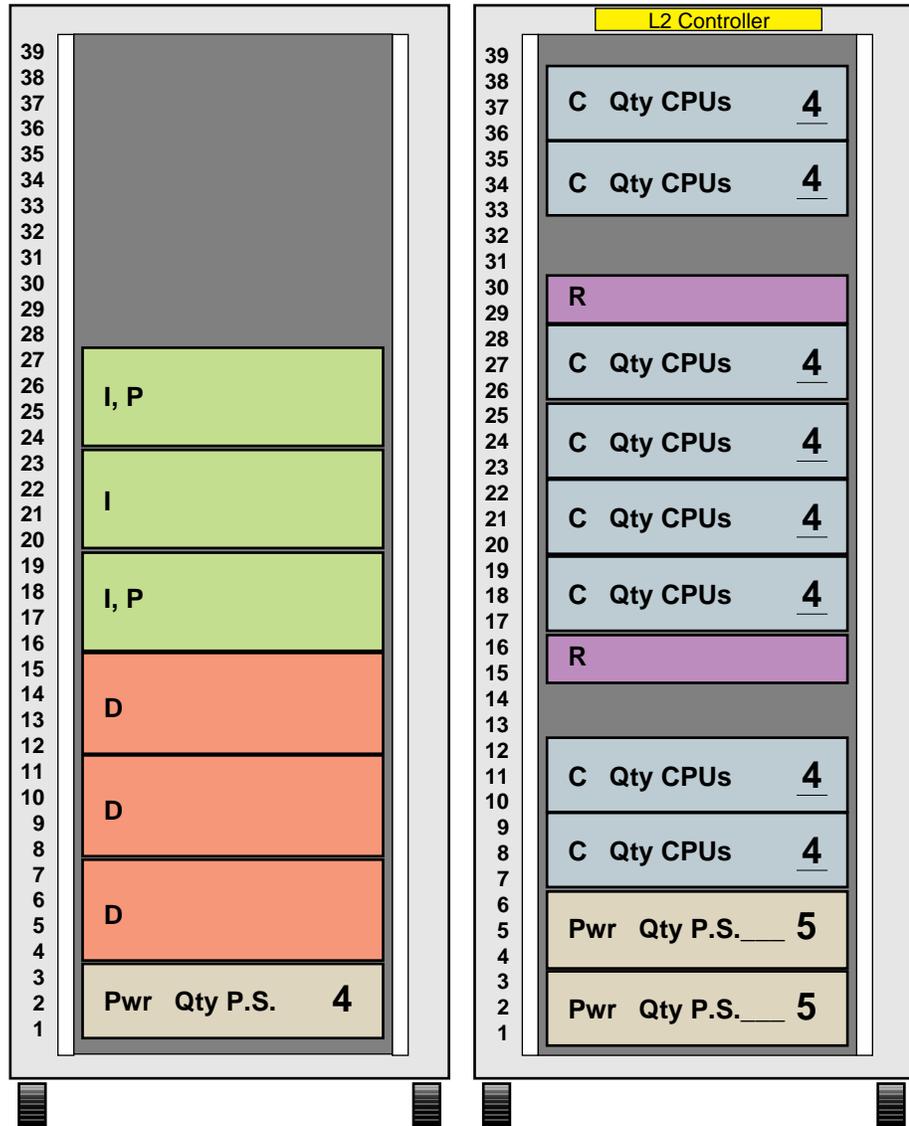
16-processor (Minimum System)

SNIA 3800
System

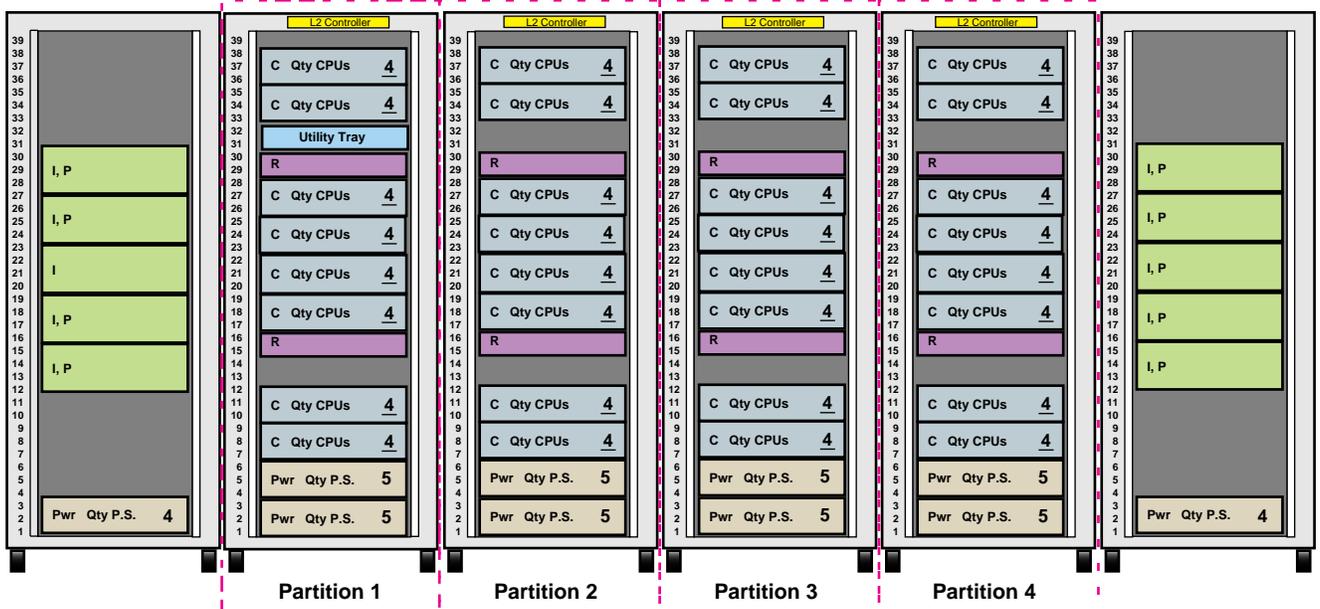


32-processor System

SNIA 3800
System



**SGI SNIA 3800 System
128-processor system
(Configured as four 32-processor systems)**



Note: Racks that contain the D bricks are not shown.

Mechanical and Electrical Specifications

Short Rack Mechanical and Electrical Specifications

| Characteristic | SGI 3000 Series | |
|---|---|--|
| Short Rack Mechanical Requirements | | |
| Height | 35.50 in. | (902 mm) |
| Width | 25.38 in. | (645 mm) |
| Depth (less system display) | 40.63 in. | (1032 mm) |
| Weight (maximum) Short Rack | 600 lbs. | (272 kg) |
| Shipping Weight (Maximum) Short Rack | 750 lbs. | (340 kg) |
| Access Requirements: Front Rear Side | 36.0 in. min. 36.0 in. None | (914 mm) (914 mm) (48.0 in. recommended) |
| Short Rack Electrical Requirements | | |
| Voltage Nominal Tolerance | North America / Japan 200-240 Vac 180-254 Vac | International 230 Vac 180-254 Vac |
| Frequency Nominal Tolerance | North America / Japan 50/60 Hz 47-63 Hz | International 50 Hz 47-63 Hz |
| Phases | Single-Phase | |
| Power Requirements (maximum) Short Rack | 2.32 kVA | (2.27 kW) |
| Hold-up Time | 20 ms | |
| Power Cable | 8 ft. (2.4 m) pluggable deep cords | |

| Characteristic | SGI 3000 Series | |
|--|---|-----------------------------------|
| Power Receptacle: Short rack (Single-phase Option) | North America / Japan NEMA 6-15R | International Country specific |
| Wall Breaker Size: Single-phase Single-phase (Europe) | 30 A 32 A | |
| Short Rack Environmental Requirements | | |
| Non-Operating Environment: Temperature Humidity Altitude | -40 to 140 °F (-40 to +60 °C) 10% to 95% non-condensing 40,000 ft. max. | |
| Operating Environment Air Temp (0 to 5000 ft). Air Temp (5000 ft to 10,000 ft) Humidity Dewpoint Altitude | 41 to 95 °F (+5 to +35 °C) 41 to 86 °F (+5 to +30 °C) 10% to 95% non-condensing TBD 0 to 10,000 ft. (0 to 3048 m) | |
| Acoustical Noise Level (maximum) | Less than 65 dBa | |
| Heat Dissipation to Air (maximum) | 7.75 Kbtu/hr (Based on 2.27 kW) | |
| Cooling Requirement | Ambient air | |
| Airflow: (Intake, Front; Exhaust, Rear) | Less than 2000 CFM | |

Tall Rack Mechanical and Electrical Specifications

| Characteristic | SGI 3000 Series | |
|--|---|--|
| Tall Rack Mechanical Requirements | | |
| Height | 74.25 in. | (1886 mm) |
| Width | 30.00 in. | (762 mm) |
| Depth (less system display) | 51.50 in. | (1308 mm) |
| Weight (maximum) | Compute Rack I/O Rack Disk Rack | 970 lbs. (440 kg) 1050 lbs. (478 kg) 1230 lbs. (558 kg) |
| Shipping Weight (Maximum) | Compute Rack I/O Rack Disk Rack | 1335 lbs. (605 kg) 1415 lbs. (642 kg) 1595 lbs. (728 kg) |
| Access Requirements: | Front Rear Side | 36.0 in. min. (914 mm) (48.0 in. recommended) 36.0 in. (914 mm) None |
| Tall Rack Electrical Requirements | | |
| Voltage | North America / Japan Nominal Tolerance | International 230 Vac 180-254 Vac |
| Frequency | North America / Japan Nominal Tolerance | International 50 Hz 47-63 Hz |
| Phases | Single-Phase or Optional Three-Phase | |
| Power Requirements (maximum) | Compute Rack I/O Rack Disk Rack | 7.14 kVA (7.00 kW) 2.18 kVA (2.14 kW) 3.16 kVA (2.97 kW) |

| Characteristic | SGI 3000 Series | |
|---|--|-------------------------------|
| Hold-up Time | 20 ms | |
| Power Cable | 8 ft. (2.4 m) pluggable deep cords | |
| Power Receptacle: | North America / Japan | International |
| Compute Rack (Three-phase Option) | (1) 60 Amp, IEC60309 (Hubbell 460C9W or Equiv.) | (1) 32 Amp, IEC60309 |
| I/O or Disk Rack (Three-phase Option) | (1) 60 Amp, IEC60309 (Hubbell 460C9W or Equiv.) | (1) 32 Amp, IEC60309 |
| Compute Rack (Single-phase Option) | (2 or 4) 30 Amp, NEMA L6-30R | (2 or 4) 32 Amp, IEC60309 |
| I/O or Disk Rack (Single-phase Option) | (2) 30 Amp, NEMA L6-30R | (2) 32 Amp, IEC60309 |
| Wall Breaker Size: | Single-phase Single-phase (Europe) Three-phase | Multiple 30 A 32 A 60 A |
| Tall Rack Environmental Requirements | | |
| Non-Operating Environment: Temperature Humidity Altitude | -40 to 140 °F (-40 to +60 °C) 10% to 95% non-condensing 40,000 ft. max. | |
| Operating Environment Air Temp (0 to 5000 ft.) Air Temp (5000 ft to 10,000 ft) Humidity Dewpoint Altitude Facilities Water Temp (for systems with heat exchanger) | 41 to 95 °F (+5 to +35 °C) 41 to 86 °F (+5 to +30 °C) 10% to 95% non-condensing TBD 0 to 10,000 ft. (0 to 3048 m) 4.4 °C to 15.6 °C | |
| Acoustical Noise Level (maximum) | Less than 65 dBa | |
| Heat Dissipation to Air (maximum) | 23.88 Kbtu/hr (Based on 7.00 kW) | |
| Cooling Requirement | Ambient air | |
| Airflow: (Intake, Front; Exhaust, Rear) | Less than 3200 CFM | |

System / Channel Bandwidths

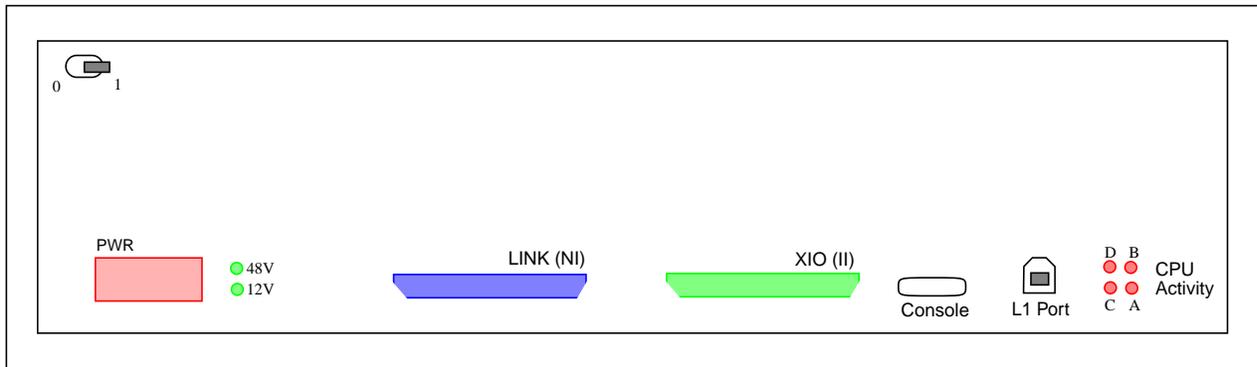
| Description | Peak Bandwidth | Sustained Bandwidth | Clock Frequency |
|------------------------------|--|--|-----------------|
| Main memory bandwidth | 3200 MB/s | 3200 MB/s | 200 MHz |
| NUMALink 3 channel bandwidth | 3.2 GB/s (full-duplex) 1.6 GB/s each direction | 1420 MB/s each direction | 800 MHz |
| Xtown2 channel bandwidth | 2.4 GB/s (full-duplex) 1.2 GB/s each direction | 1066 MB/s (half-duplex) 872 MB/s each direction | 600 MHz |
| Xtown channel bandwidth | 1.6 GB/s (full-duplex) 800 MB/s each direction | N/A | 400 MHz |
| PCI channel bandwidth | 128 MB/s in 32-bit mode 256 MB/s in 64-bit mode | N/A | 33 MHz |
| | 256 MB/s in 32-bit mode 512 MB/s in 64-bit mode | N/A | 66 MHz |

List of Acronyms / Abbreviations

| | |
|-------|---|
| ASIC | Application specific integrated circuit |
| ATAPI | AT attachment packet interface |
| ATM | Asynchronous Transfer Mode |
| CDROM | Compact disk read only memory |
| CFM | Cubic feet per minute |
| CPU | Central processing unit |
| DIMM | Dual inline memory module |
| DVD | Digital video disk |
| EIA | Electronics Industry Association |
| FCS | First customer ship |
| FDDI | Fiber distributed device interface |
| FC | Fibre channel |
| FRU | Field replaceable unit |
| HIPPI | High performance parallel interface |
| IOS | Input / output subsystem |
| NUMA | Non-uniform memory access |
| PCB | Printed circuit board |
| PCI | Peripheral Component Interconnect (an industry standard for connecting peripherals to a CPU) |
| RAID | Redundant array of independent disks |

| | |
|------|---|
| SCSI | Small Computer System Interface |
| U | Unit (one U = 1.75 inches) |
| USB | Universal serial bus (12 mbits/s transfer rate) |
| VRM | Voltage regulator module |

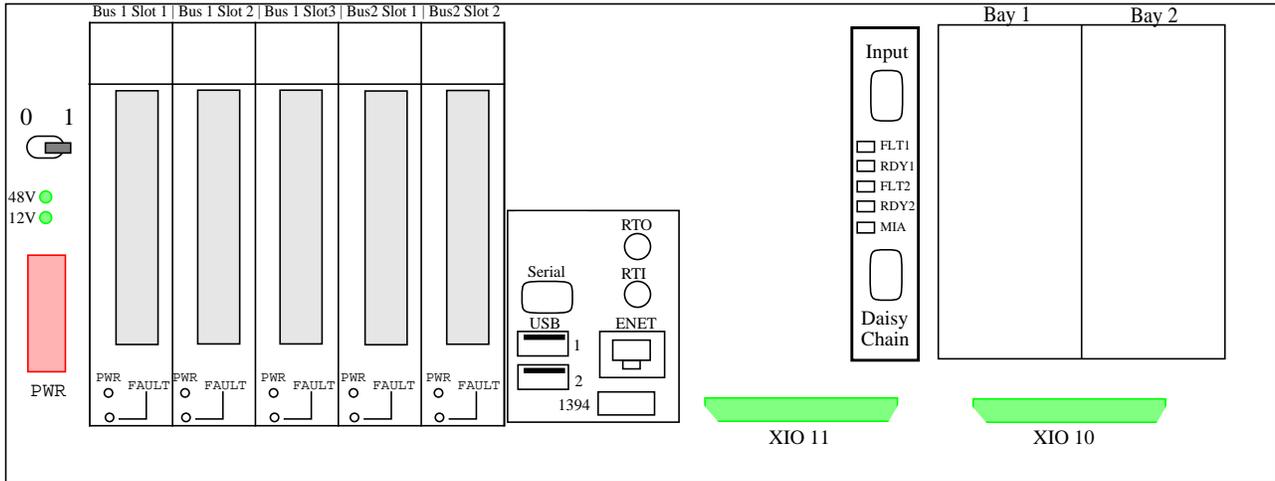
Rear View of Bricks



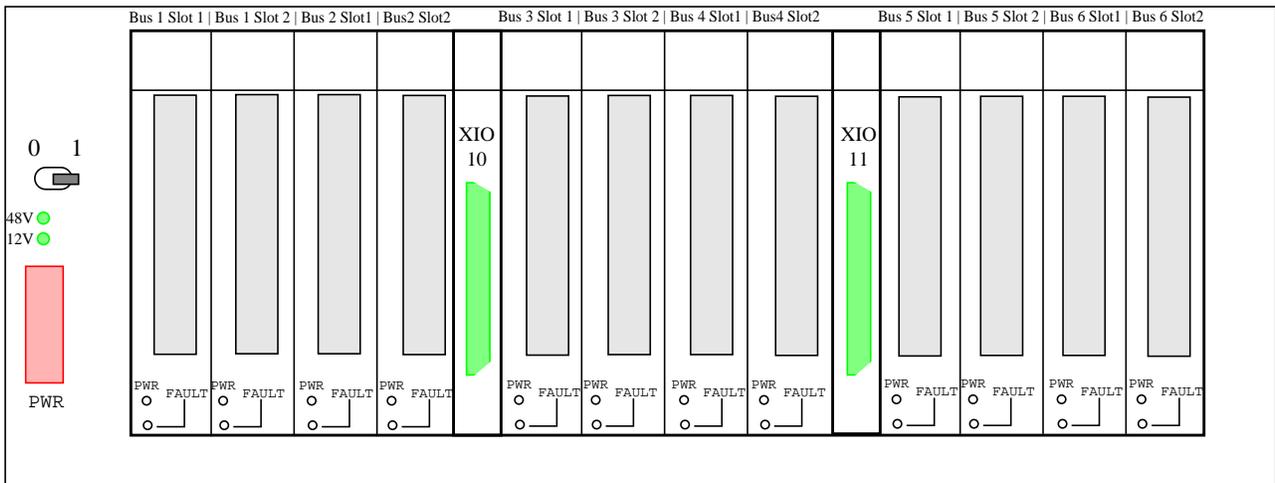
C Brick
(Rear View)



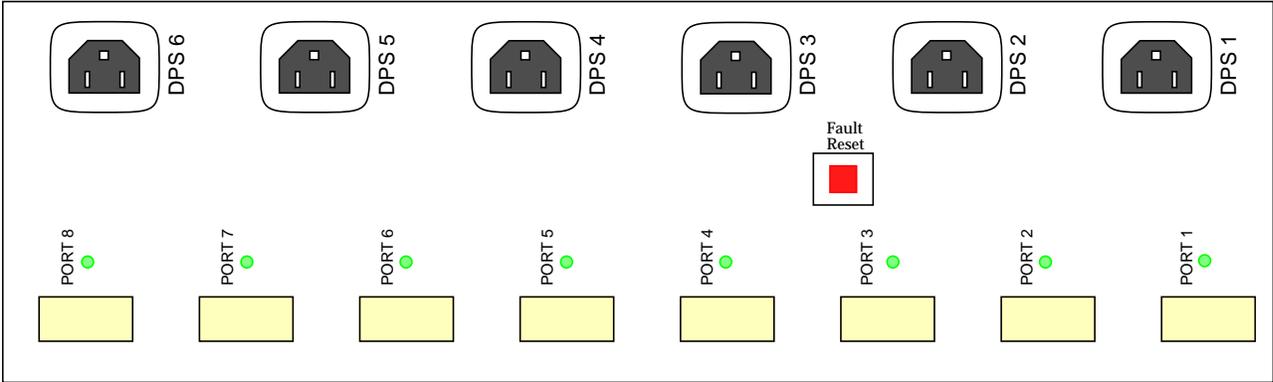
R Brick
(Rear View)



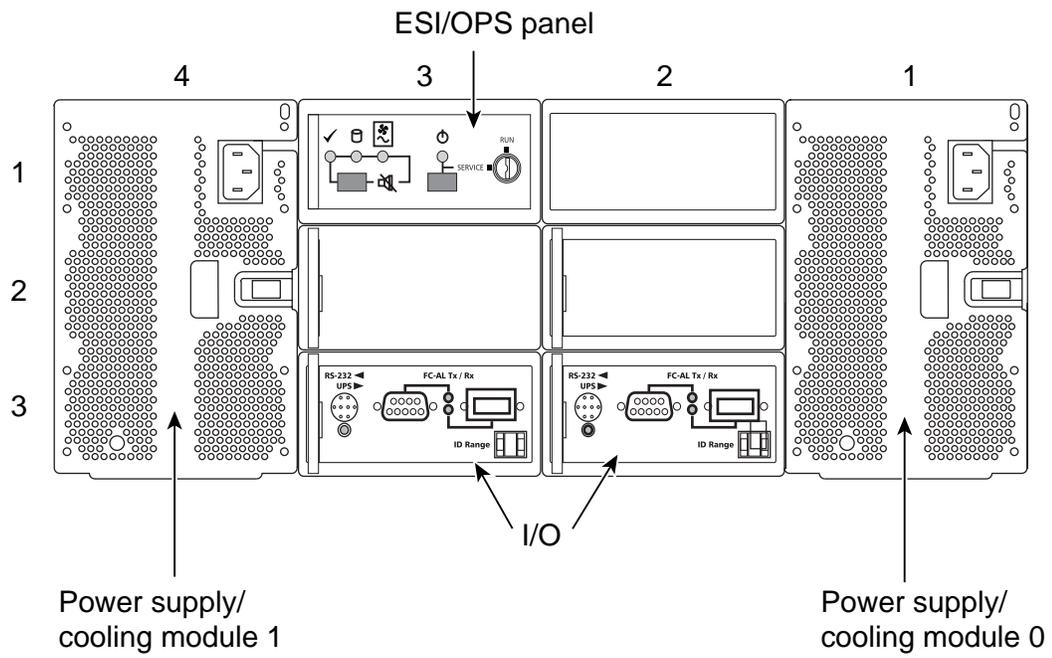
I Brick
(Rear View)



P Brick
(Rear View)



Power Bay
(Rear View)



D Brick
(Rear View)

Component Dimensions, Weights, and Power

| Product Name | Dimensions | Weight | Power | Number Of U |
|---------------|--|-----------------------|----------------------------|-------------|
| Tall Rack | 74.0 in. H × 32.0 in. W × 48.0 in. D | 400 lbs (empty) | N/A | N/A |
| Short Rack | 34.0 in. H × 24.0 in. W × 42.0 in. D | 200 lbs (empty) | N/A | N/A |
| C-brick | 5.06 in. H × 17.19 in. W × 27.80 in. D | 55 lbs (fully loaded) | ~825 watts | 3 |
| D-brick | 6.64 in. H × 17.50 in. W × 27.74 in. D | 94 lbs (fully loaded) | ~400 VA | 4 |
| I-brick | 6.64 in. H × 17.50 in. W × 27.74 in. D | 69 lbs (fully loaded) | ~190 watts (3.3 Vdc cards) | 4 |
| R-brick | 3.35 in. H × 17.38 in. W × 27.5 in. D | 18 lbs | ~60 watts | 2 |
| P-brick | 6.64 in. H × 17.50 in. W × 27.74 in. D | 70 lbs (fully loaded) | ~225 watts (3.3 Vdc cards) | 4 |
| Power Bay | 5.01 in. H × 17.5 in. W × 24.5 in. D | 72 lbs (fully loaded) | N/A | 3 |
| L2 Controller | 1.75 in. H × 14.0 in. W × 6.5 in. D | 4 lbs | 30 watts | N/A |

Index

Numbers

1394, 32, 33, 34

A

acronyms, 129-130

B

bandwidth, 53
 bisection, 61
 system/channel, 127

brick
 defined, 3
 placement, 73
 rear view, 131-134
 See also C-brick, D-brick, and so on.

C

cabling, 79-82
card, defined, 3
C-brick, 14, 15, 26-28
 and expanding system, 107-108
 bandwidth and performance, 53
 cabling, 79
 configuration, 85-87
 defined, 3

 rear view, 131
channel defined, 3
connectors, 131-134
Crosstalk
 defined, 4

D

D-brick, 30-31
 bandwidth and performance, 54
 defined, 4
DIMM, 29
 configuring, 86
 defined, 4
 not compatible with Origin2000 and Origin200
 DIMMs, 29
 placement, 87
 See also memory.
 sizes, 29
disk drives
 installation guidelines, 90
distributed power supply. *See* DPS.
DPS, defined, 4
DSM, defined, 9
dumb terminal, configuring, 98

E

Ethernet, 32, 33, 34, 46

hub, 47-50
L2 controller, 43

F

fibre channel, 33
and SGI 3200 system control, 44
disk drives, in I-brick, 32, 33, 35
loops, in D-brick, 31

H

hot plug
defined, 4
disk drives, 31
PCI cards, 32
hot swap
defined, 4
DPSs, 19
fans, 28, 35, 37, 39
hub, 9, 14, 15
Ethernet, 47-50

I

I-brick, 32-35
bandwidth and performance, 54
configuration, 94
defined, 4
drawing, 34
rear view, 132
See also I/O.
IEEE 1394. *See* 1394.
interconnection
C-brick and I/O brick, 81-82
fabric, defined, 5
interconnection fabric, 63

002 through 16 processor, 63
016 through 64 processor, 64

I/O, 9

brick, cabling to C-brick, 79-82
expansion. *See* P-brick.
See also I-brick, PCI,

Itanium microprocessor, 14, 15

L

L1, L2, L3. *See* system controller.

link

channels and R-brick, 38
defined, 5

M

memory, 29
adding, 29, 108
clock speed, 29
removing, 29
See also DIMM.

Merced, defined, 5

MetaRouter, 46
defined, 5

MIPS

defined, 5

module, defined, 5

N

network, defined, 6

node, defined, 6

NUMALink 3, 14, 15, 46
cable, 44, 45
defined, 6

O

Origin200, DIMMs not compatible, 29
Origin2000
DIMMs not compatible, 29

P

partition, defined, 6
Partitioning rules, 67
P-brick, 35-37
bandwidth and performance, 55
block diagram, 36
configuration, 95
defined, 6
drawing, 37
rear view, 132
PCI
buses
in I-brick, 32
in P-brick, 35
cards, slots
I-brick, 32
cards—bandwidth and performance, 56
defined, 6
hot plugging, 32
in SGI 3000 series system, 41
peripheral
adding capacity, 108
component interconnect. *See* PCI.
POD, defined, 6
port, defined, 6
power bay, 19-20
configuration, 96
defined, 6
power board, defined, 6
power-on diagnostics. *See* POD.
processor, 49

adding, 107
and DIMMs, 29
defined, 6
See also Itanium

R

rack, 6
cooling, 51
defined, 6-7
deskside, 17
full-height, 16
L2 controller required, 43
specifications, 123
R-brick, 38-39
block diagram, 39
cabling, 79
defined, 7
drawing, 40
in rack, and L2 controller, 43
rear view, 131
real-time sync I/O, 34
remote connection, configuring, 98
router, 9, 15
defined, 7
RTI, 34
RTO, 34

S

serial port, 33, 34
cable, 44, 45
L2 controller, 43
SGI 3000 series system
configuration, 99-105
cooling, 51
SGI 3000 system
component dimensions and weights, 135

SGI 3200
brick placement, 73, 74, 75, 77
components, 99-100
configuration, 111-113
L3 controller, 44
PCI slots, 41
remote maintenance requirement, 43

SGI 3400
components, 101
configuration, 115-118
L2 controller, 43
L3 controller, 44, 45
PCI slots, 41

SGI 3800
components, 104-105
configuration, 119-??
Ethernet hub, 48, 49
L2 controller, 43
L3 controller, 44
L3 system controller, 46
PCI slots, 41

software, 109

system controller
L1
defined, 5
L2, 43
configuration, 98
defined, 5
R-brick link, 38
required in rack, 43
L3, ??-46
configuration, 98
defined, 5
network configuration, 98

T

terminology, 3-8
acronyms, 129-130

third-party applications, 110

U

Universal Serial Bus. *See* USB.
USB, 32, 33, 34, 46
and R-brick, 38, 39
cable, 45
defined, 7
L2 controller, 43

V

voltage regulator module. *See* VRM.
VRM, defined, 7

W

wrapped drive, defined, 7
wrapper, defined, 7

X

Xbridge, 33, 36
XIO
defined, 8
Xtown2, 14, 15, 33, 36
cable, 44, 45, 79-82
defined, 8
ports on P-brick, 32, 35