SGI™ Origin® 3000 Series

Internal Technical Configuration Manual

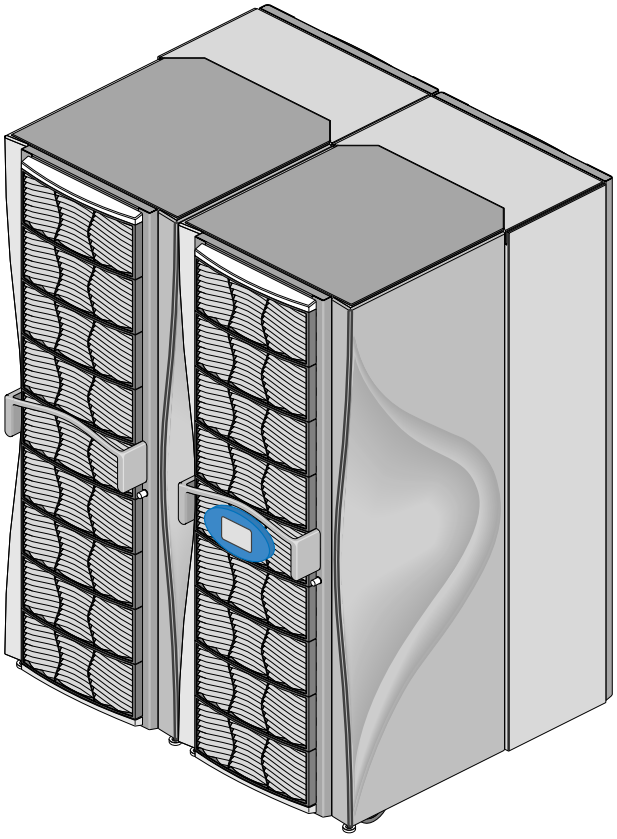

Document Number 108-0266-001

Last updated 03-21-01

CONTRIBUTORS

Written by Dick Brownell

Illustrated by Dan Young and Kelly Begley

Edited by Allison Gosbin

Production by Karen Jacobson

Engineering contributions by SGI Hardware Engineering Dept.

Cover Design By Sarah Bolles, Sarah Bolles Design, and Dany Galgani, SGI Technical Publications.

Record of Revision

Version	Description
001	October, 2000 - Original Release.

Contents

List of Figures

List of Tables

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 Origin 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 Origin 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 help create a specific SGI Origin 3000 series solution.

1.2 Related Documentation

For more information, consult the following documents:

- SGI Origin 3000 Server Series Site Planning 007-3601-001
- SGI Origin 3000 System Architecture 108-0240-001

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 Origin 3000 Server Family
  Overview of the SGI Origin 3000 series.
- Section 3 System Overview
  Technical overview of the SGI Origin 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 Brick Locations and Xtown2 Cable Configurations
  Defines brick locations for each system type and NUMAlink cable restrictions.
- 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:
  - Required software
  - optional software
  - 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	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: • **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 Origin 3000 series. Also referred to as the *compute node*.
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	*See* rack
compute node	*See* C-brick
CPU	*See* 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	*See* 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 *Level 1* 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 *Level 2* 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 *Level 3* 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.
PIMM (processor integrated memory module)	A printed circuit assembly that contains two processors and memory. The C-brick can be configured with either one or two PIMMs.
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 *cabinet, chassis*. • **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.

Term	Description
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.
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 Origin 3200	An entry-level system that is contained within a single cabinet and has a maximum of 8 processors.
SGI Origin 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 Origin 3400 system.
SGI Origin 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 Origin 3800 system
sled	*See* wrapper
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.

Term	Description
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 *sled.*
X-brick	A 4-U high rackmount enclosure that supports 4 half-height XIO cards.
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 *Crosstown two*.

2. Introducing the SGI Origin 3000 Series Family

The SGI Origin 3000 series comprises a family of multiprocessor distributed shared memory (DSM) computer systems. The SGI Origin 3000 series uses a global-address-space cache-coherent multiprocessor that scales to 512 processors (~410 GFLOPS peak) in a cache-coherent domain. Four processors, each with 8 Mbytes of private secondary cache, are connected at a Bedrock ASIC. This Bedrock ASIC acts as a crossbar between the processors, local SDRAM memory, the network interface, and the I/O interface. Four Bedrocks, each supporting four processors, are connected to an 8-ported router that can connect up to 32 router chips in a maximum $4 \times 4 \times 32$ extended hypercube topology. The modularity of the DSM approach combines the advantages of low entry-level cost with global scalability in processors, memory, and I/O.

Initial versions of the SGI Origin 3000 series are based upon the MIPS R12000A processor.The MIPS R12000A processor is a four-way, 64-bit superscalar RISC processor. “Four-way” means that it fetches and decodes four instructions per clock cycle and issues them to five fully pipelined execution units. “Superscalar” means that it has enough independent, pipelined execution units that it can complete more than one instruction per clock cycle. The MIPS R12000A also has speculative branching and out-of-order execution and is initially designed with a 400-MHz clock speed.

The SGI Origin 3000 series uses a PCI-based I/O subsystem as its primary I/O protocol. A Crosstalk I/O system, that is compatible with some XIO boards from Origin 2000 and Octane systems, is available to support legacy I/O devices and certain specialized devices such as HIPPI-6400.

Table 2-1 describes ranges of system capacities and performance.

Table 2-1 System Configuration Ranges

CATEGORY	MINIMUM	MAXIMUM
Number of Processors	2	512
Peak Performance with R12K	~1600 MFLOPS (one 2-processor node brick)	~410 GFLOPS (128 four-processor node bricks)
Peak Performancd with R14K	~2000 MFLOPS (one 2-processor node brick)	~512 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	0.768 GB/s (one I-brick)	~ 152 GB/s (one I-brick and 63 P-bricks)

2.1 SGI Origin 3000 Series Model Naming Conventions

Table 2-2 describes the three classes of the SGI Origin 3000 series systems.

Table 2-2 SGI Origin 3000 Series Systems

System	Description
SGI Origin 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 Origin 3200C	Contains 8 processors per node with the ability to scale to thousands of processors. A node is made up of two C-bricks and one I-brick. Three nodes are configured into a single tall rack. Nodes are interconnected using a 16 port Myrinet switch.
SGI Origin 3400	Contains 4 to 32 processors. The standard processor and I/O configuration is contained in a single tall rack. However, the configuration can be expanded to include a second tall rack for larger I/O configurations. Additional disk racks can be added as required.
SGI Origin 3800	Contains from 16 processors to 512 processors (from one to sixteen tall 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.

The product marketing codes for the SGI Origin 3000 series are structured as follows:

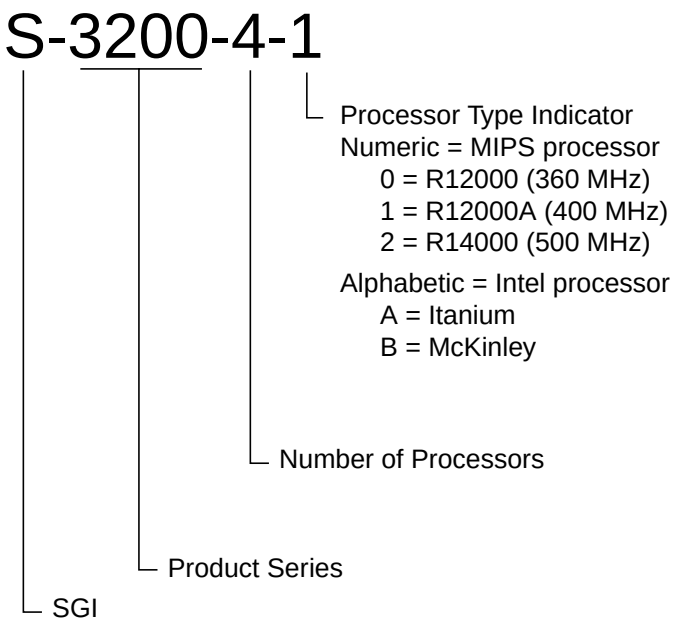

3. SGI Origin 3000 Series System Overview

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

The SGI Origin 3000 series systems are distributed shared memory (DSM) computer systems that scale from 2 to 512 processors. In a DSM system, each processor contains memory that it shares with the other processors in the system. Cache coherence is maintained through a directory-based scheme.

The modularity of the DSM systems combine the advantages of low entry-level cost with global scalability in processors, memory, and I/O. Table 2-1 lists the system configuration ranges for the SGI Origin 3000 series system.

Initial SGI Origin 3000 series systems use the MIPS R12000A processor. The R12000A processor is a 64-bit RISC, superscalar processor with speculative branching, out-of-order execution, and a 400-MHz operating clock speed.

A new R14000 processor will be available in the 1st quarter of 2001. The R14000 processor is the next generation processor from MIPS which has a 500 MHz operating clock speed.

Four processors, each with 8 Mbytes of private secondary cache, are connected at a Bedrock (refer to Figure 3-1). This Bedrock ASIC acts as a crossbar between the processor interface, local memory interface, the network interface, and the I/O interface.

Four Bedrocks, each supporting four processors, connect to an eight-port router that can connect up to 32 routers in a maximum 4 x 4 x 32 extended hypercube 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".

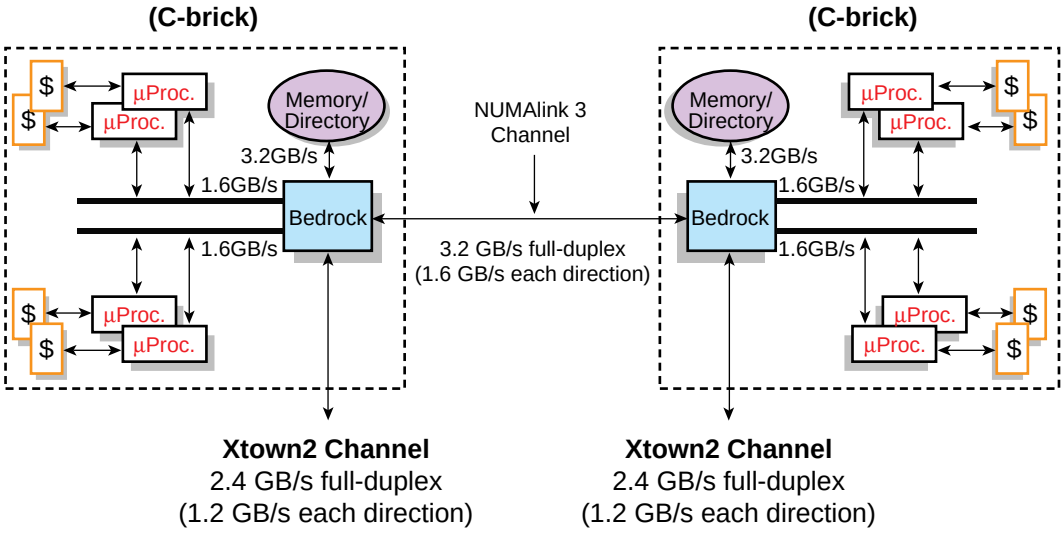

Figure 3-1 Eight-processor SGI Origin 3000 Series Block Diagram

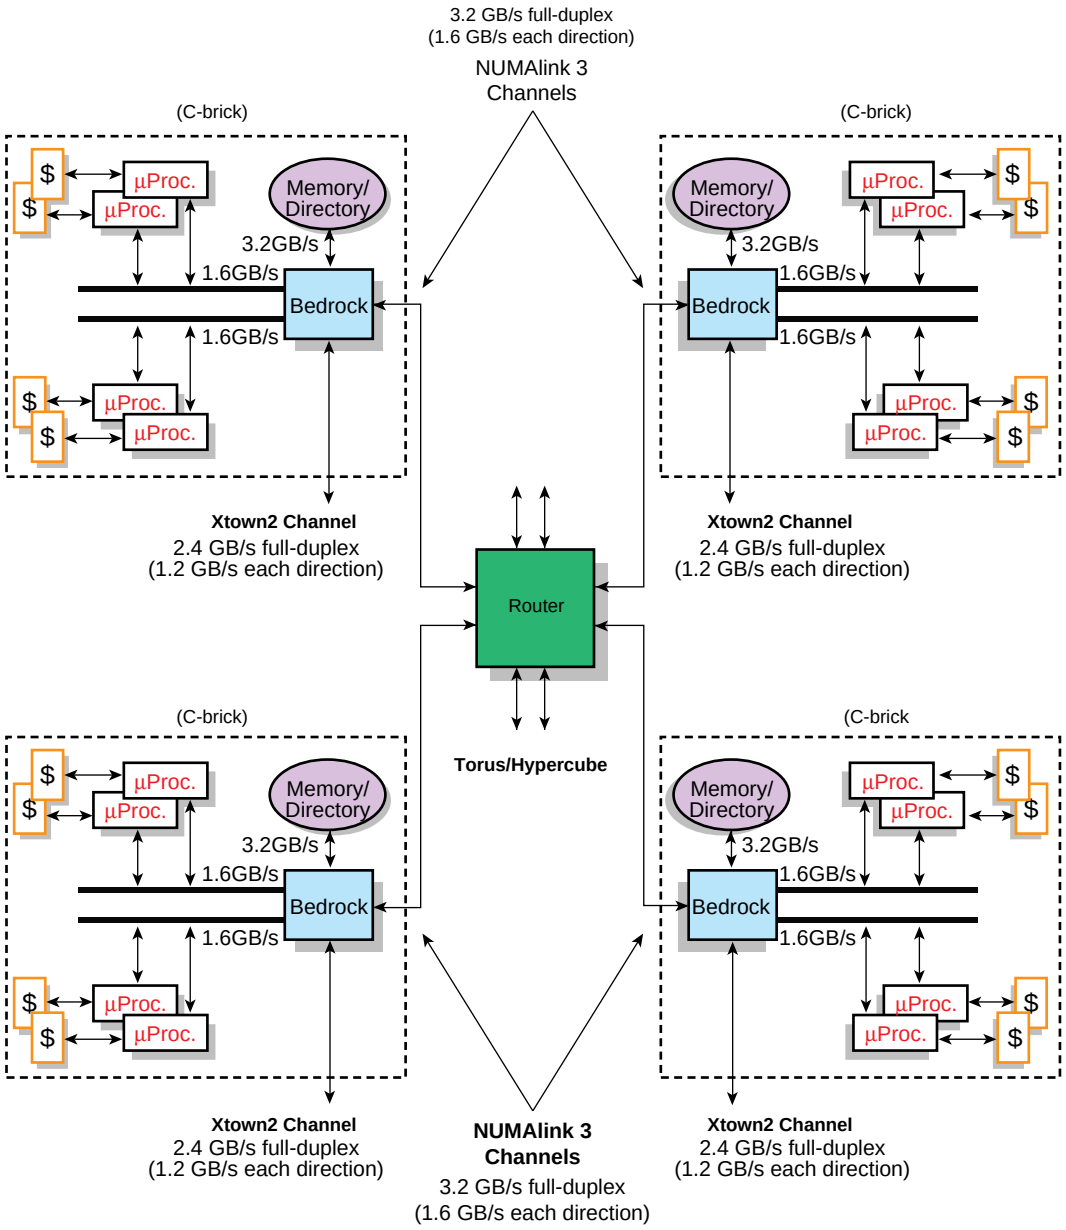

Figure 3-2 Sixteen-processor SGI Origin 3000 Series Block Diagram

3.1 Tall Rack

The SGI Origin 3000 series tall rack can house the following standard 19-inch rackmounted subassemblies: C-brick, D-brick, G-brick, I-brick, P-brick, R-brick, X-brick and power bay. The tall rack is used by SGI Origin 3200C, SGI Origin 3400 and SGI Origin 3800 systems. The outer dimensions of the tall rack with casters, side panels, and decorative doors does not exceed 74 in. high × 30 in. wide × 50 in. deep.

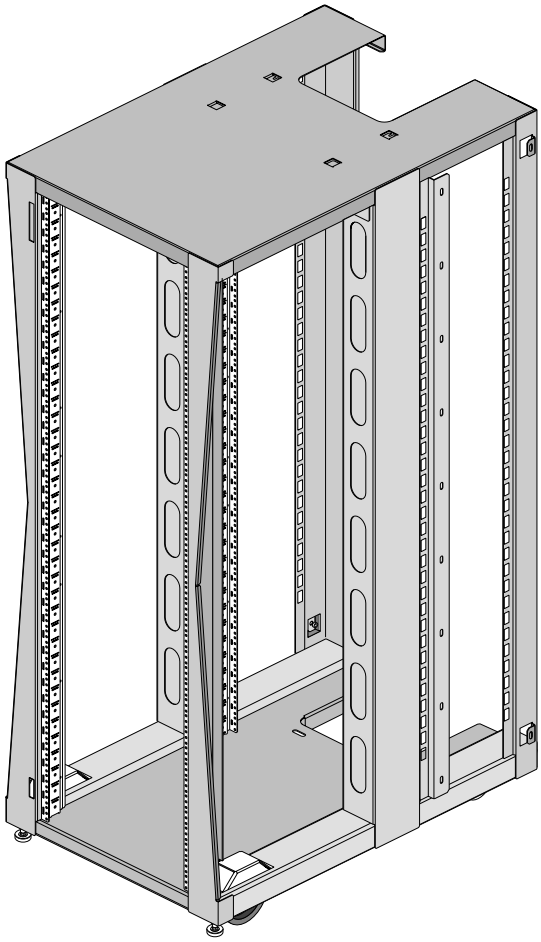

Figure 3-3 Tall Rack (Front View)

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

3.2 Short Rack

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

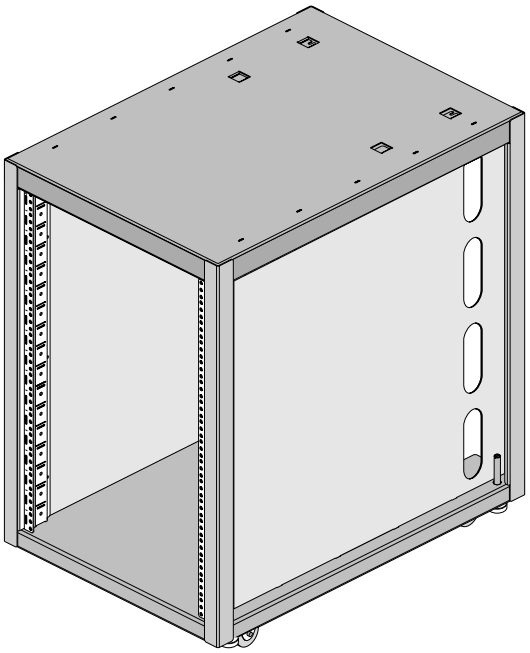

Main features of the short rack:

- 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 the power bay, D-bricks, and third-party equipment.

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) in an SGI Origin 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.

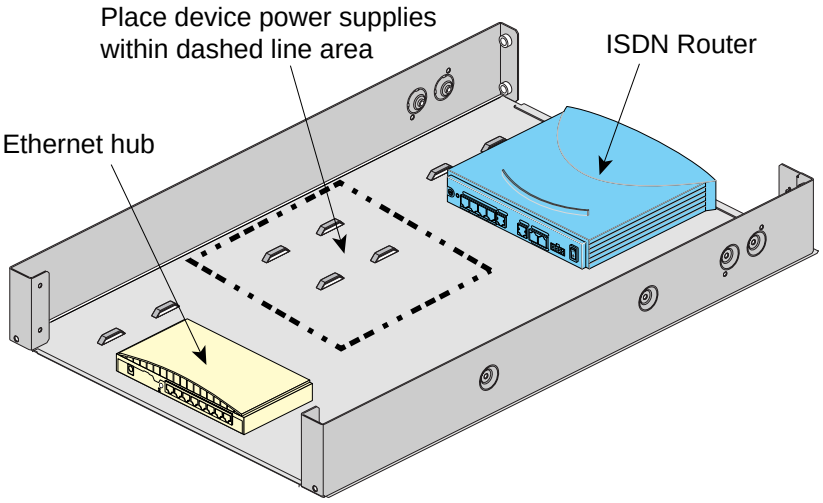

Figure 3-5 Utility Tray (Rear View)

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 bricks are added to the configuration.

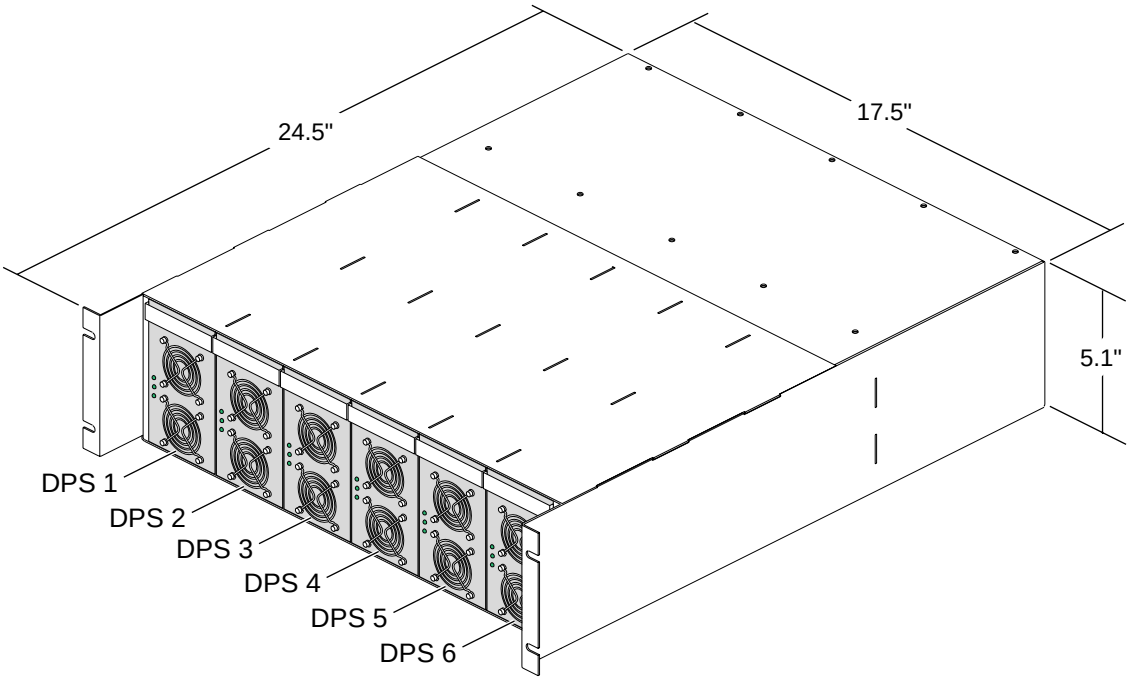

Figure 3-6 Power Bay (Front View)

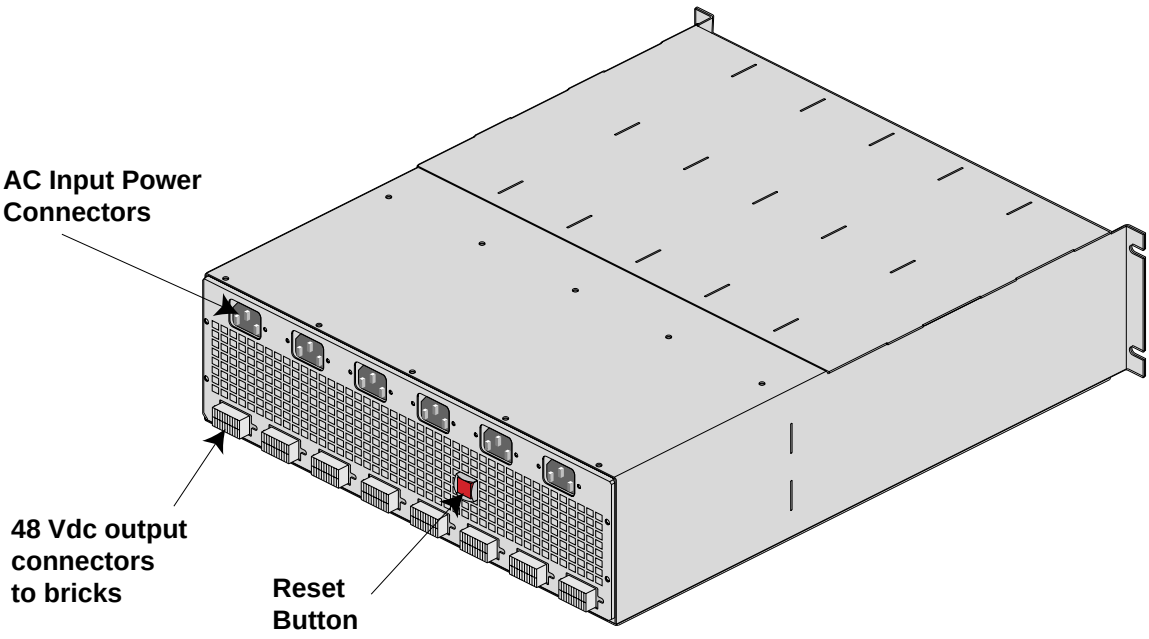

Figure 3-7 Power Bay (Rear View)

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 Origin 3200C, SGI Origin 3400 and SGI Origin 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 Origin 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 tall racks.

Table 3-1 PDU 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

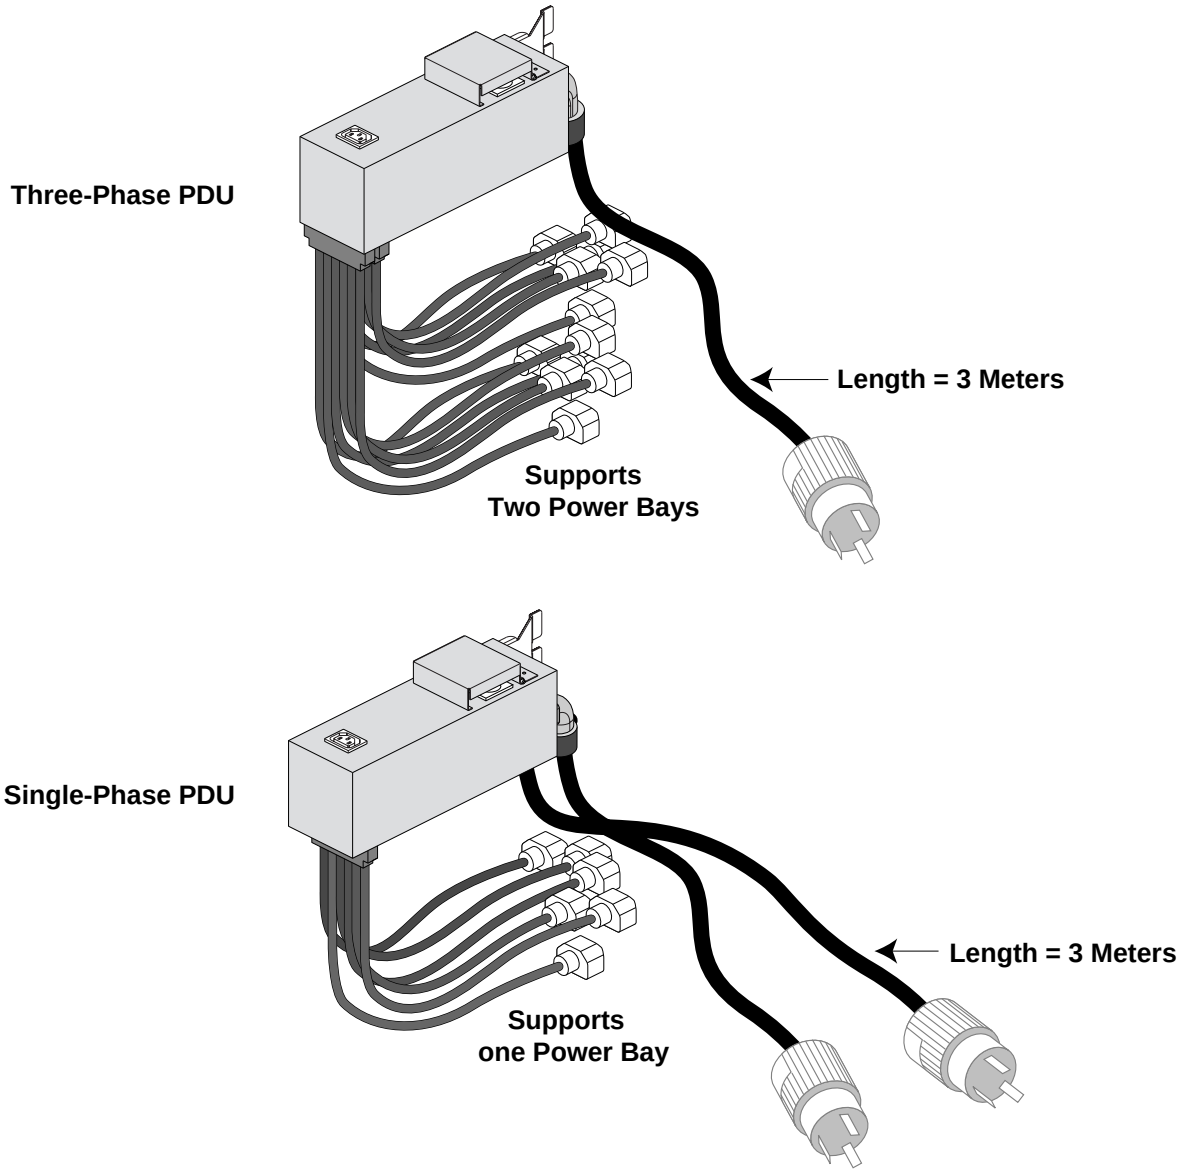

Figure 3-8 Single-Phase and Three-Phase PDUs

3.6 Power Distribution Strip (PDS)

All SGI Origin 3000 series systems use power distribution strips. In the SGI Origin 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 Origin 3200C, SGI Origin 3400 and SGI Origin 3800 tall 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 under-floor outlet. In a tall rack, a 2.0-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

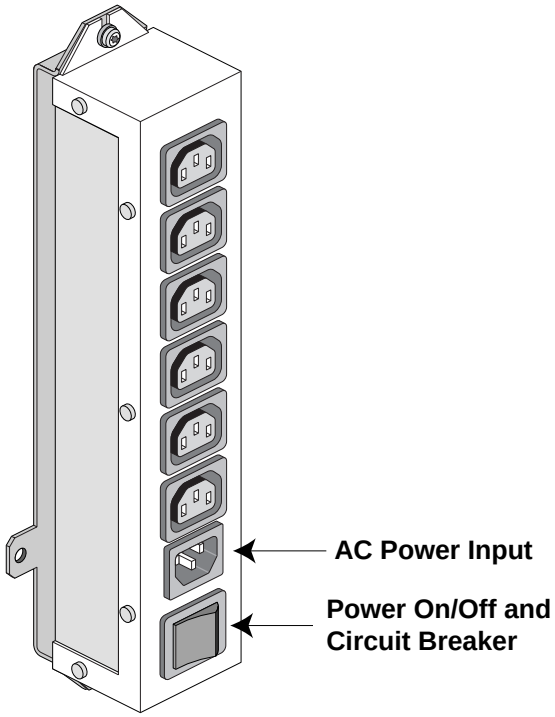

Figure 3-9 Power Distribution Strip

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 DK-N1P-003	(1) kit per power bay
		DK-N3P-001	(1) kit per rack
	Europe and other	DK-N1P-002	(1) kit per power bay
		DK-N3P-002	(1) kit per rack
G-brick in graphics rack	U.S.A, Canada Mexico, Japan	DK-GM3-003	(1) kit per G-brick
	Australia	DK-GM3-007	
	Europe and other	DK-GM3-011	
Power distribution strip	U.S.A, Canada Mexico, Japan	DK-S-PWRCBL-001	(1) kit per graphics rack (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 RISC processors with an 8-MB secondary cache
- Eight DIMM slots; each DIMM pair has two banks 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 and cache are housed on separate PIMM boards. Each PIMM contains two processors and secondary cache. Figure 3-10 shows the block diagram of a C-brick.

Note: All transfer rates in Figure 3-10 are peak rates.

Note: The "$" in Figure 3-10 means "cache."

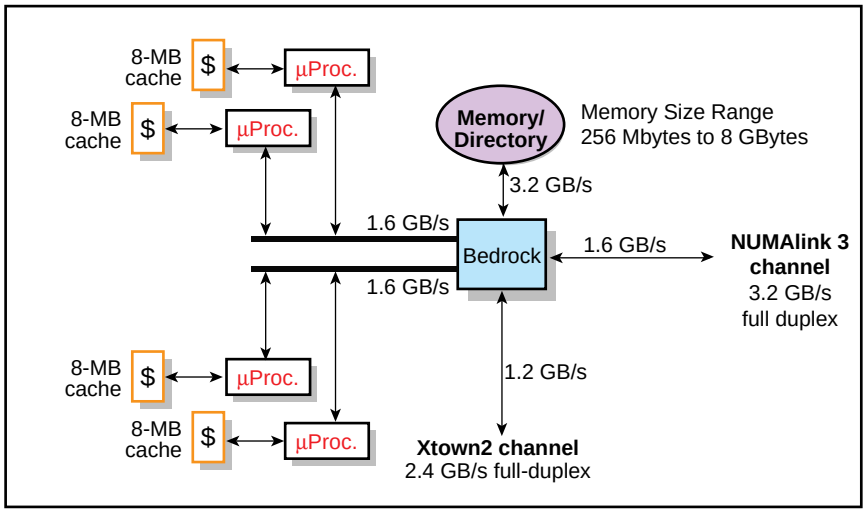

Figure 3-10 C-brick Block Diagram

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 8-Mbyte secondary cache per processor
- Contains one 1.6-GB/s (each direction) NUMAlink 3 channel
- Contains one 1.2-GB/s (each direction) Xtown2 channel
- Contains one console port with DB9 connector
- Contains one L1 USB port

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

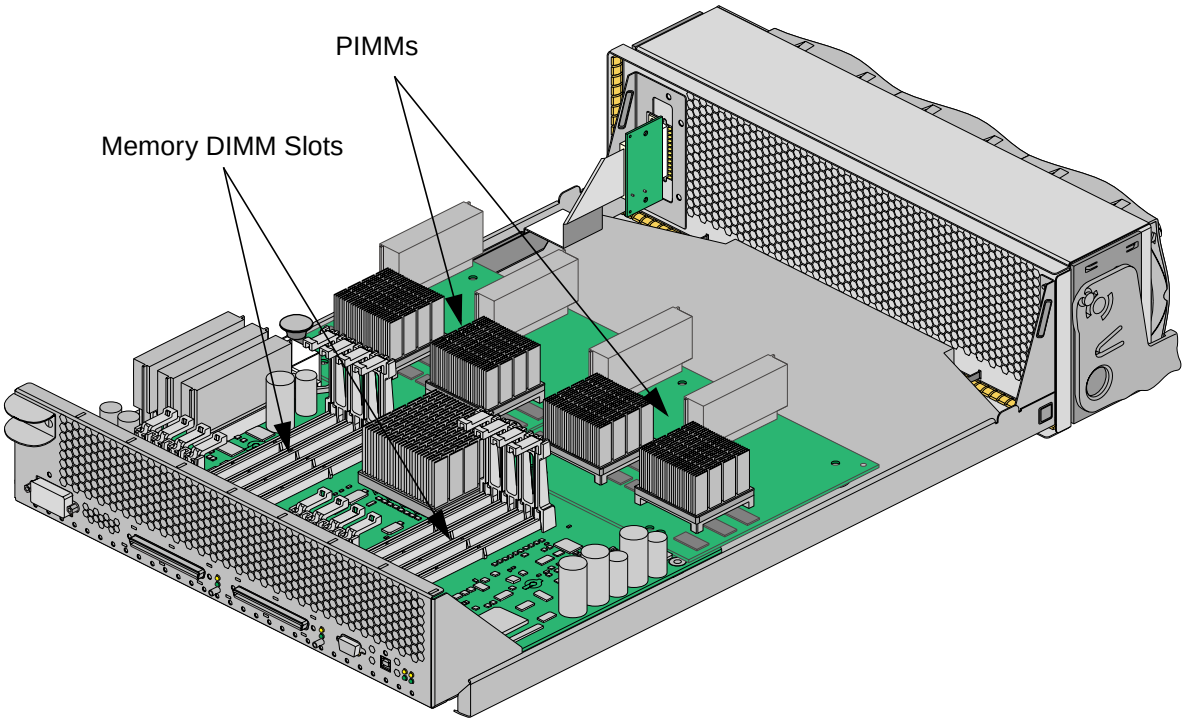

Figure 3-11 C-brick Logic Carrier Assembly (Rear View)

The C-brick has the following mechanical features:

- Height: 3 U
- Weight: 65 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 (~308 Watts)

3.8.1 Processor

Two 400-MHz MIPS R12000A processors or two 500-MHz MIPS R14000 processors and 8-MB of secondary cache are mounted on a PCB. This assembly is called the processor-integrated memory module (PIMM). The C-brick can be configured with one or two PIMMs.

The MIPS R12000A processor has the following features:

- 64-bit RISC design, 0.25-micron CMOS process
- Single-chip superscalar RISC dataflow architecture
- 8-MB secondary cache
- 32-KB 2-way set-associative data cache
- 32-KB 2-way set-associative instruction cache
- 2,048-entry branch prediction table
- 48-entry active list
- 32-entry two-way set-associative branch target address cache (BTAC)
- Doubled secondary cache prediction table for improved hit rate
- Improved branch prediction by using global history mechanism
- Maintains code and instruction set compatibility with the MIPS R10000

3.8.2 Memory

Main memory consists of up to four memory bank pairs 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 the memory size must be increased or decreased in whole banks, and a single bank of memory is contained on two DIMMs. 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 112 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 DIMM Pair Installed	2 DIMM Pairs Installed	3 DIMM Pairs Installed	4 DIMM Pairs 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 Origin 3000 series DIMMs are not compatible with the DIMMs used in SGI Origin 200,SGI Origin 2000, or SGI Octane systems.

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) within an SGI Origin 3000 series rack. The D-bricks configured as JBOD will be factory installed and ship as part of the system.

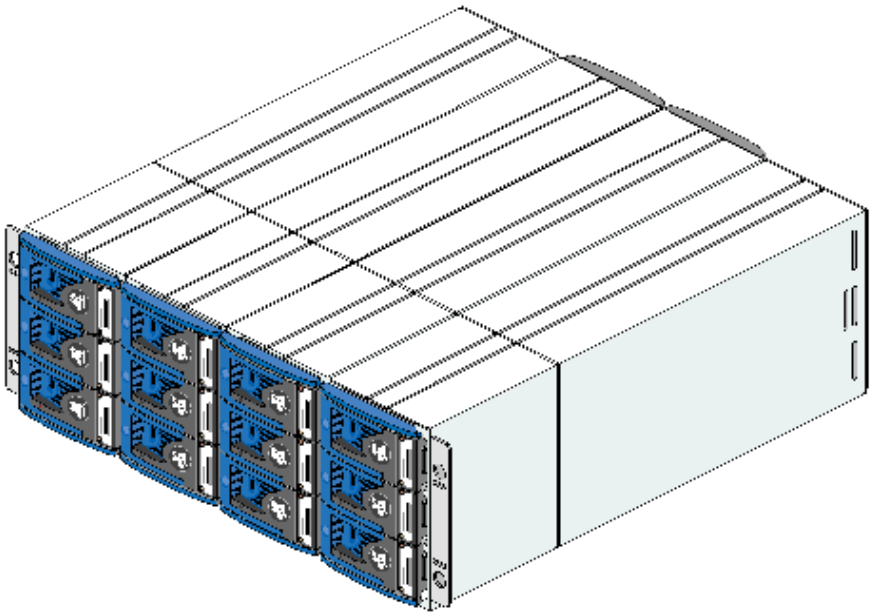

Figure 3-12 D (Disk) Brick Assembly

The SGI TP-9100 storage system is another storage solution for SGI Origin 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 Origin 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 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 4-8 on page 62 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.

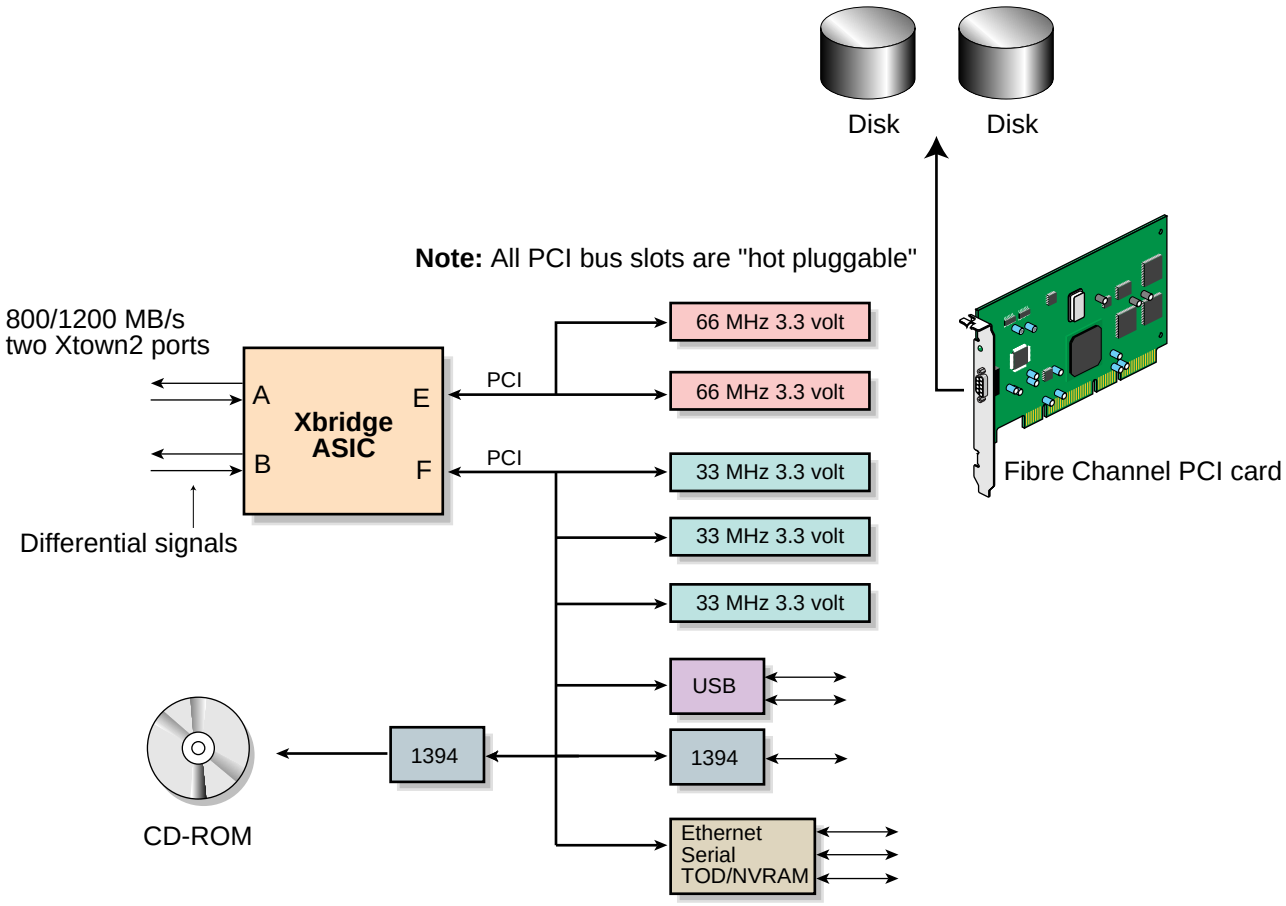

Figure 3-13 I-brick Block Diagram

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

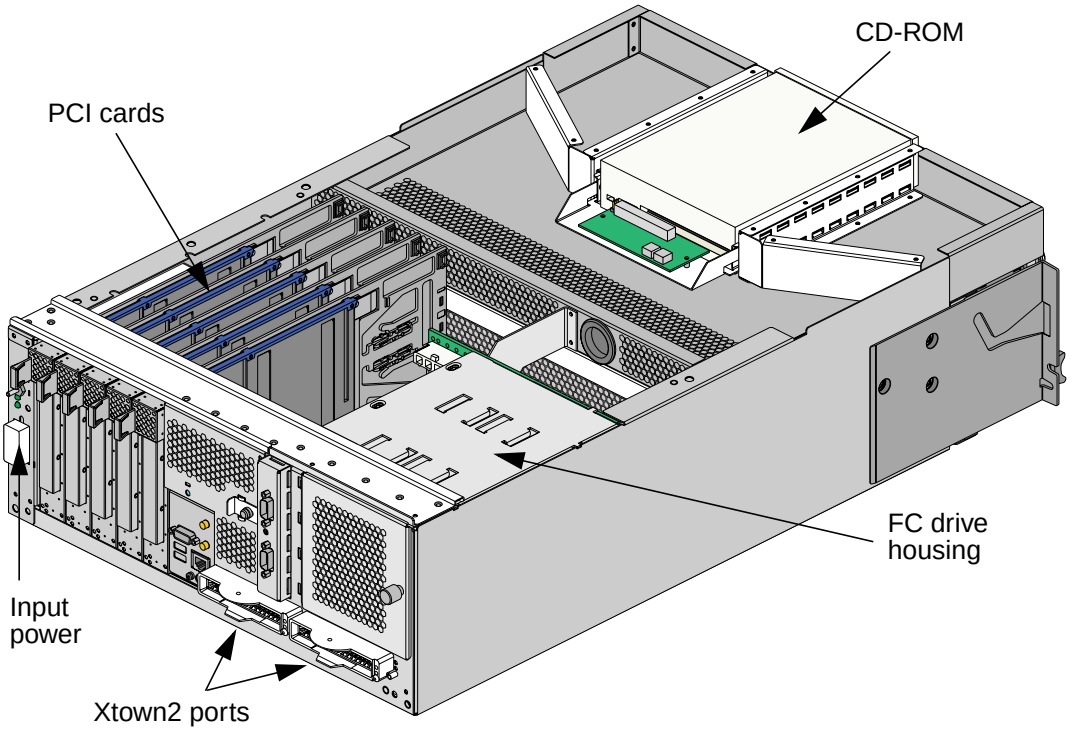

Figure 3-14 I-brick Assembly (Rear View)

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 4-8 on page 62 for a list of supported PCI cards and to Section 7 for configuration guidelines.

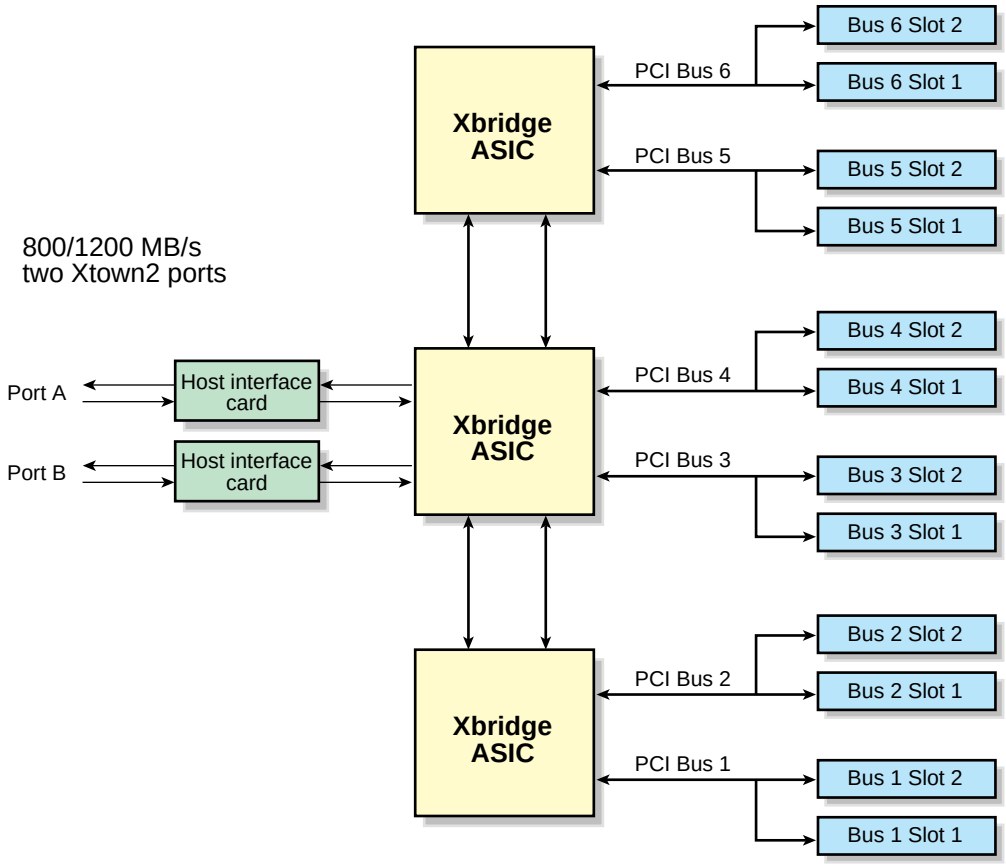

Figure 3-15 P-brick Block Diagram

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

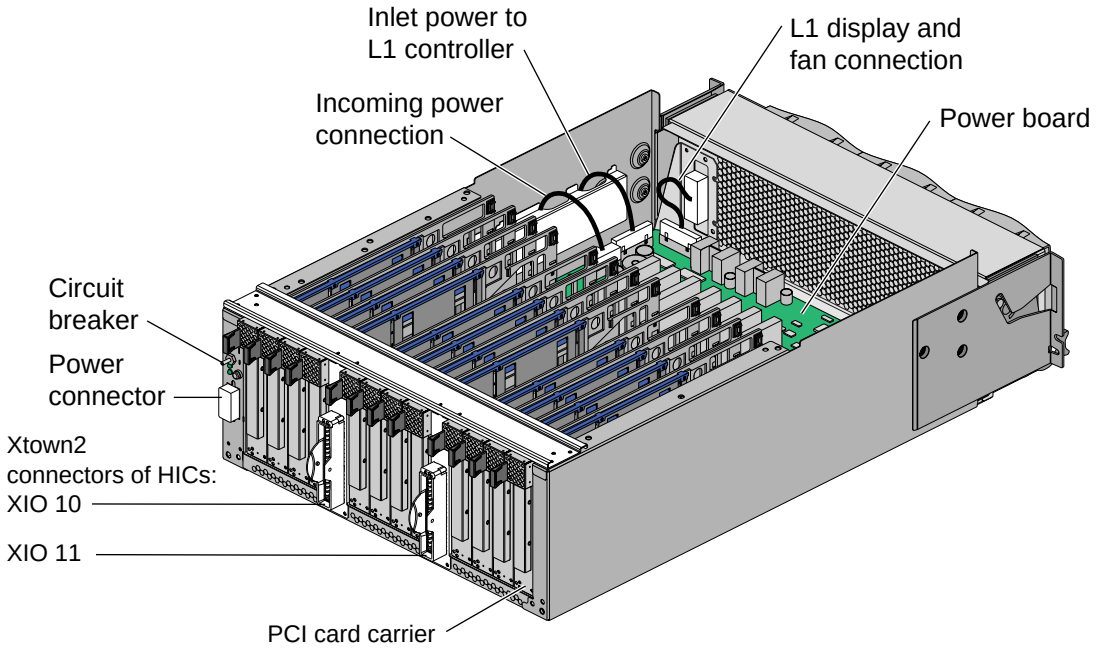

Figure 3-16 P-brick Assembly (Rear View)

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 XIO Expansion (X-brick)

The X-brick is an XIO expansion brick; it contains four half-height XIO slots that are fully compatible with the XIO slots in Origin 2000 systems. This enables existing Origin customers to migrate their XIO cards to the SGI Origin 3000 series systems. Each slot is capable of supporting cards with maximum power levels of 50 watts. Refer to Table 4-6 on page 60 for a list of supported XIO cards.

The X-brick 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 X-brick is 800 MB/s. Refer to Figure 3-17 for the X-brick block diagram.

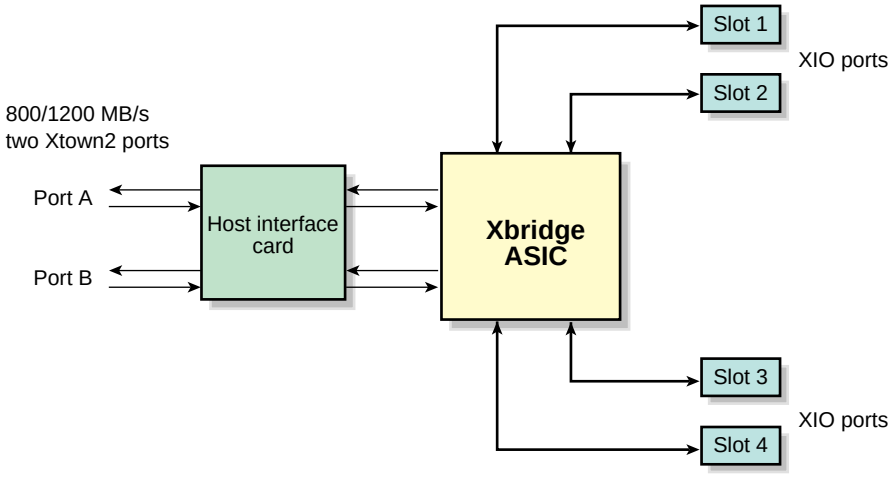

Figure 3-17 X-brick Block Diagram

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

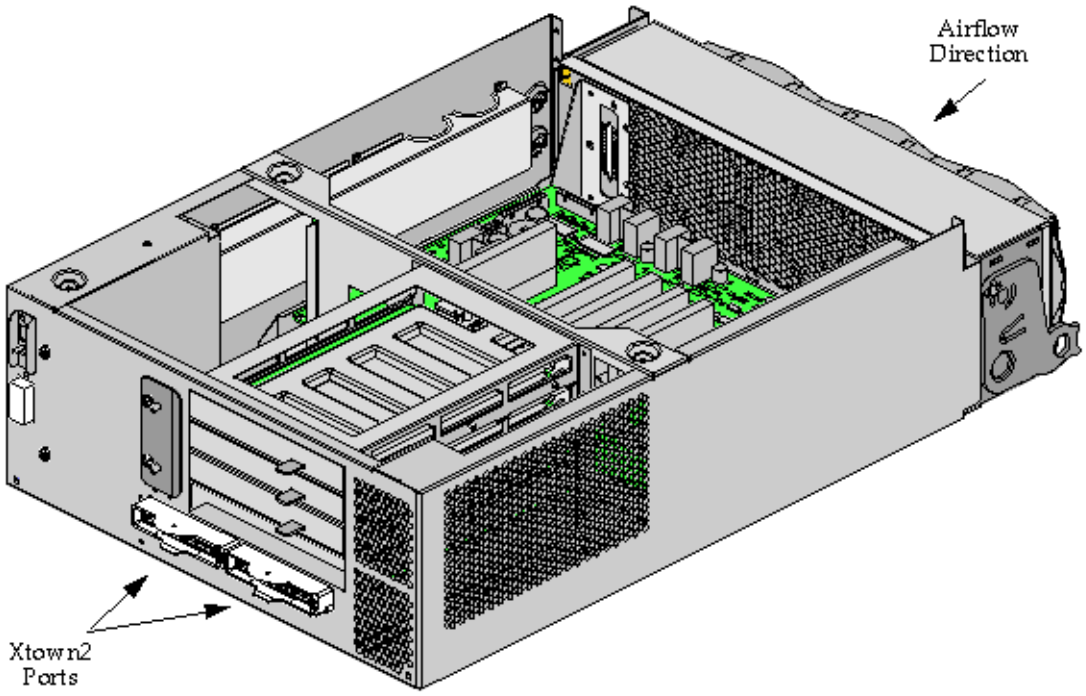

Figure 3-18 X-brick Assembly (Rear View)

The X-brick has the following electrical feature:

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

The X-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 four Crosstalk half-height cards
- Cooling fans are hot-swappable
- Input power is +48 Vdc (~225 watts)

3.13 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 R-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 3-19. Metarouters and repeat routers use all eight ports to connect to other R-bricks.

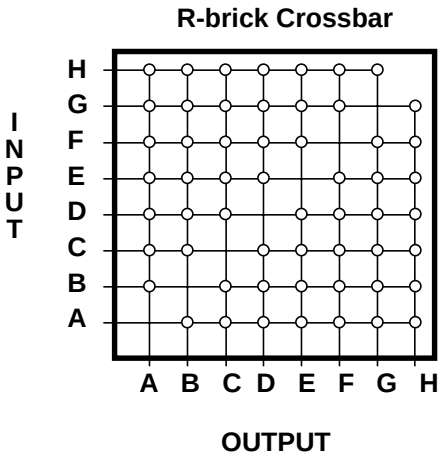

When an R-brick-to-R-brick connection is made through ports that carry USB signals, the USB signals are ignored. USB signals to the C-bricks are distributed over the network cables. Because an R-brick can have a maximum of four C-bricks attached to it, only four of the R-brick's 100-pin network connectors 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.

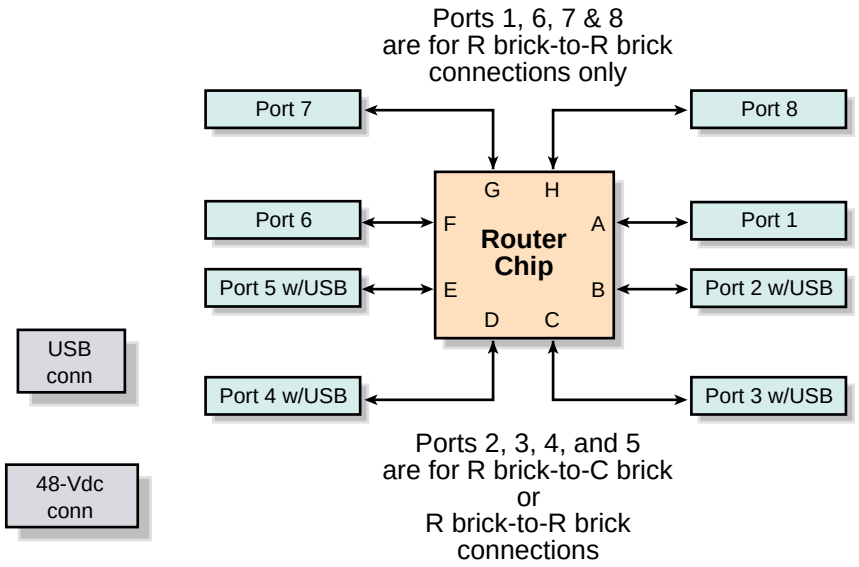

Figure 3-19 R-brick Block Diagram

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 3-20 shows a rear view of the R-brick enclosure with the cover removed.

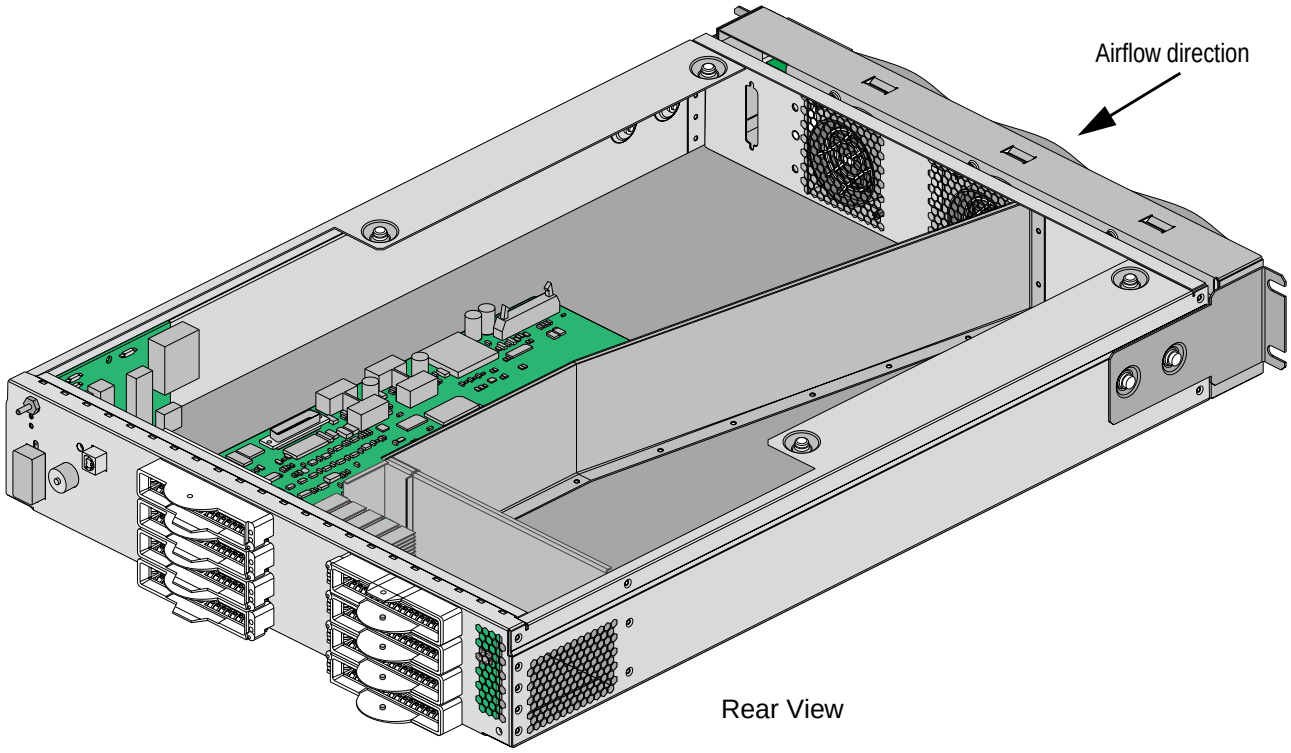

Figure 3-20 R-brick Assembly (Rear View)

3.14 Graphics Subsystem (G-brick)

The G-brick is a second generation (Onyx3) graphics subsystem that can be scaled from 1 to 16 pipes (eight G-bricks in a system). The G-brick is 18 U high and is rackmounted in an SGI Origin 3000 series rack. Each G-brick has one 2RM port and one 4RM port. The higher the number of RMs (raster managers) per port the higher the performance of the pipe. Refer to the *SGI Onyx 3000 Internal Technical configuration Manual* for detailed information on configuring the G-brick.

Each tall rack can hold from one to two G-bricks. Each pipe of a G-brick connects to a Xtown2 channel of an I-, or X-brick via a NUMAlink cable. The G-brick's input power cord plugs into the AC wall outlet and the L1 controller is connected to a L2 controller via a USB cable. If a G-brick is in a rack that does not have an L2 controller, a USB hub can be used to pass the system controller information to an L2 in an adjacent rack.

Refer to Figure 3-21 for an illustration of the G-brick.

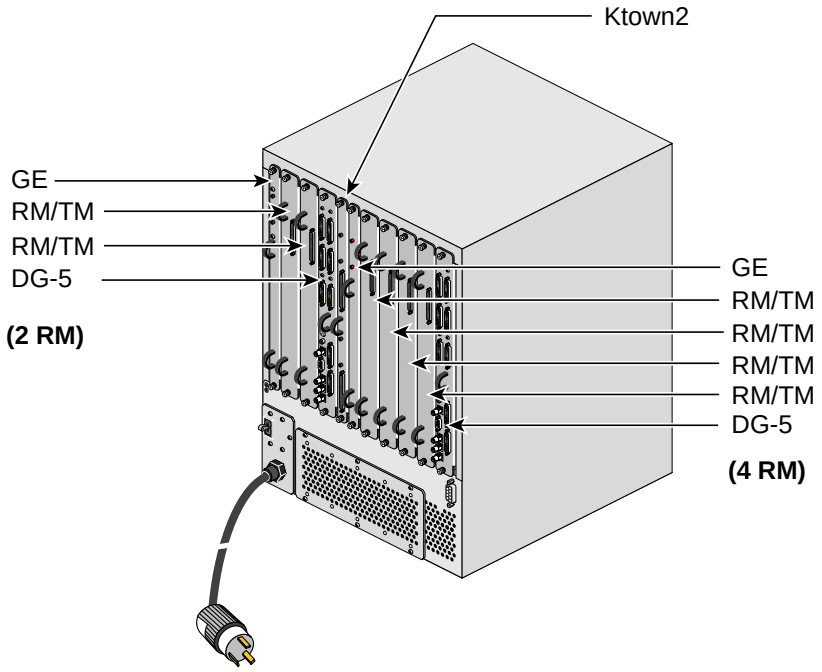

Figure 3-21 G-brick (Rear View)

3.15 Peripherals

The SGI Origin 3000 server series supports a wide range of peripherals to deliver unmatched performance for large supercomputing workloads. Refer to Table 4-6 and Table 4-8 for a list of the supported XIO and PCI cards.

3.16 PCI Based I/O Channels

The I-brick and P-brick are the subsystems that provide PCI support for the SGI Origin 3000 server series. The I-brick provides five PCI slots: four slots for customer-defined PCI cards and one slot for a Fibre Channel disk controller 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 that are available in SGI 3200, SGI Origin 3400 and SGI Origin 3800 systems.

Table 3-5 Number of Available PCI Slots

System	I/O Bricks	Maximum Number of PCI Slots Available
SGI Origin 3200	one I-brick and one P-brick	17
SGI Origin 3400	one I-brick and seven P-bricks	89
SGI Origin 3800	one I-brick and 63 P-bricks	761

3.17 Sixteen Port Myrinet Switch

The 16-port Myrinet switch is used in the SGI Origin 3200C systems. This switch provides the network fabric to cluster the 8-processor nodes together. The Myrinet sixteen-port switch requires one Myrinet LAN PCI card to be installed in each node (I-Brick).

Figure 3-22 Sixteen-port Myrinet Switch (Rear View)

The Sixteen-port Myrinet switch has the following Features:

- Weight: 15 lbs
- Dimensions: 3.4 H × 17.0 W × 12.0 D
- Uses convection cooling - requires one U of clearance above unit
- Input power is 220 Vac (~40 watts)

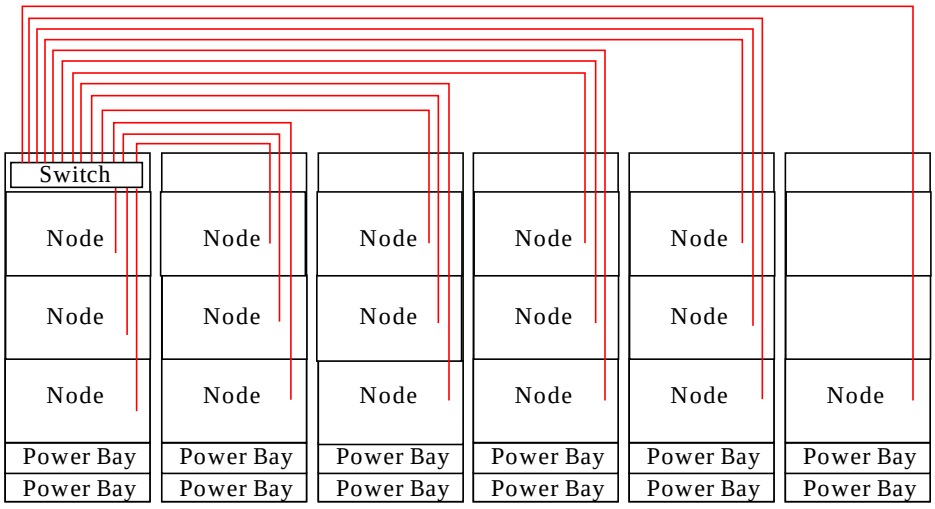

Figure 3-23 SGI Origin 3200C 128-processor Cluster cabling

3.18 System Control

The purpose of the SGI Origin 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 Origin 3000 series control system is composed of three levels:

- L1 controller - brick-level system controller
- L2 controller - rack-level system controller
- L3 controller - system-level controller

3.18.1 L1 Controller

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

3.18.2 L2 Controller

The L2 controller is optional in SGI Origin 3200 systems, however, it is required with all SGI 3400 and SGI Origin 3800 systems. The dimensions of the L2 controller are 1.75 in. H × 14.0 in. W × 6.5 in. D. The L2 controller is mounted in the top of the rack and therefore 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 3-24):

- 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 Origin 3200 systems)
- A rack display is desired

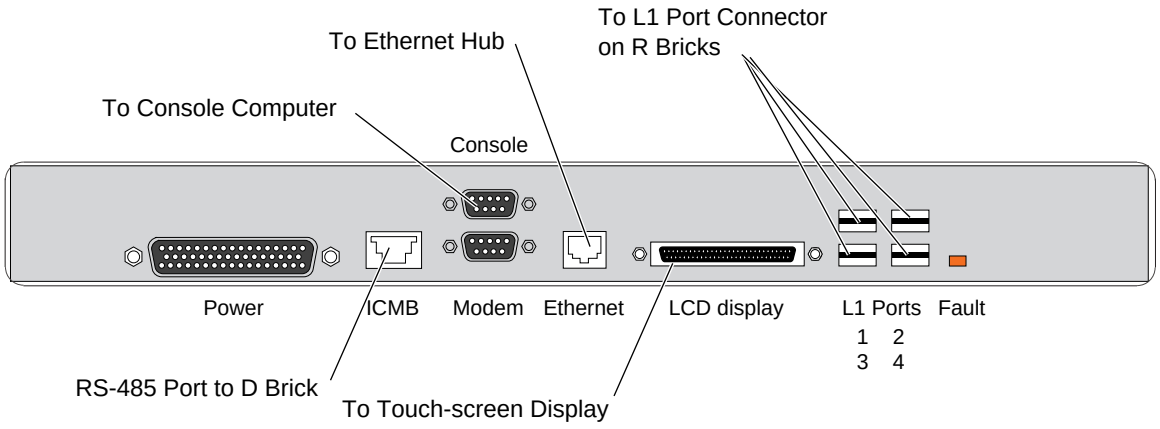

Figure 3-24 L2 Controller (Rear View)

3.18.3 L3 Controller

The L3 controller is a system-level controller. The L3 controller is software that runs on a stand-alone workstation or laptop computer. The L3 connects to the L2 controllers via a 10BaseT Ethernet hub. In an SGI Origin 3200 system, the L3 can connect directly to an L1 controller in a C-brick via a USB port. The L3 controller is optional in all system sizes.

The system control network configuration depends on the class of the SGI Origin 3000 series system. Figure 3-25 through Figure 3-27 show typical SGI Origin 3200, SGI Origin 3400, and SGI Origin 3800 system control configurations.

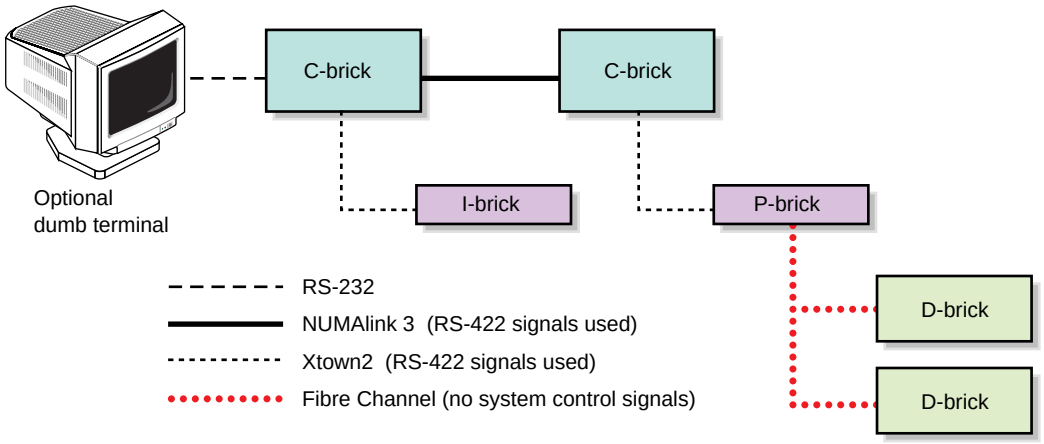

Figure 3-25 SGI Origin 3200 System Control Network (Typical)

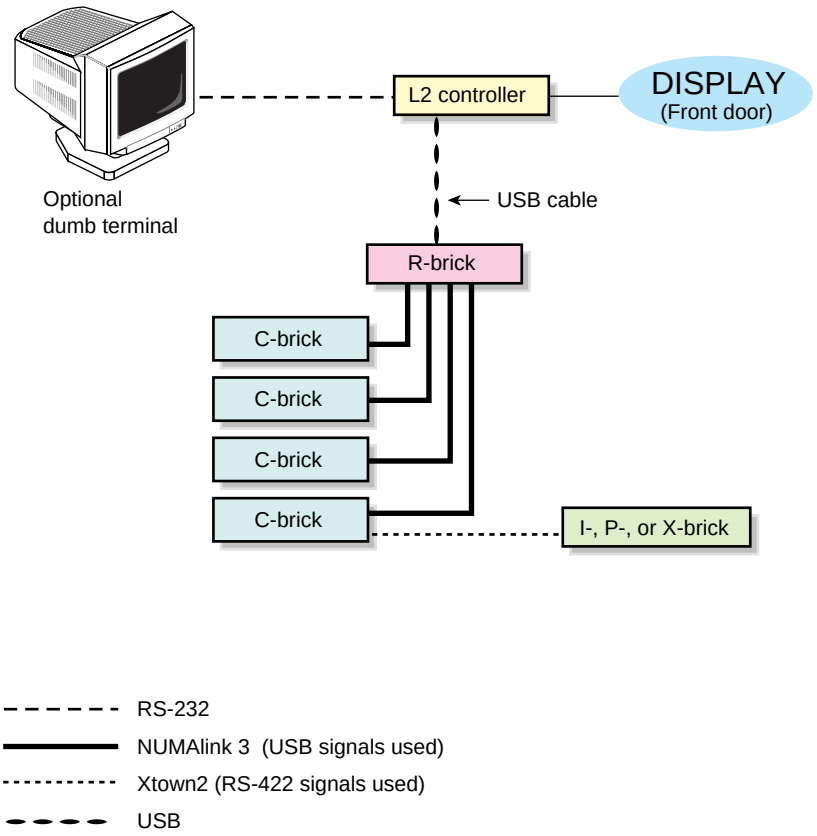

Figure 3-26 SGI Origin 3400 System Control Network (Typical)

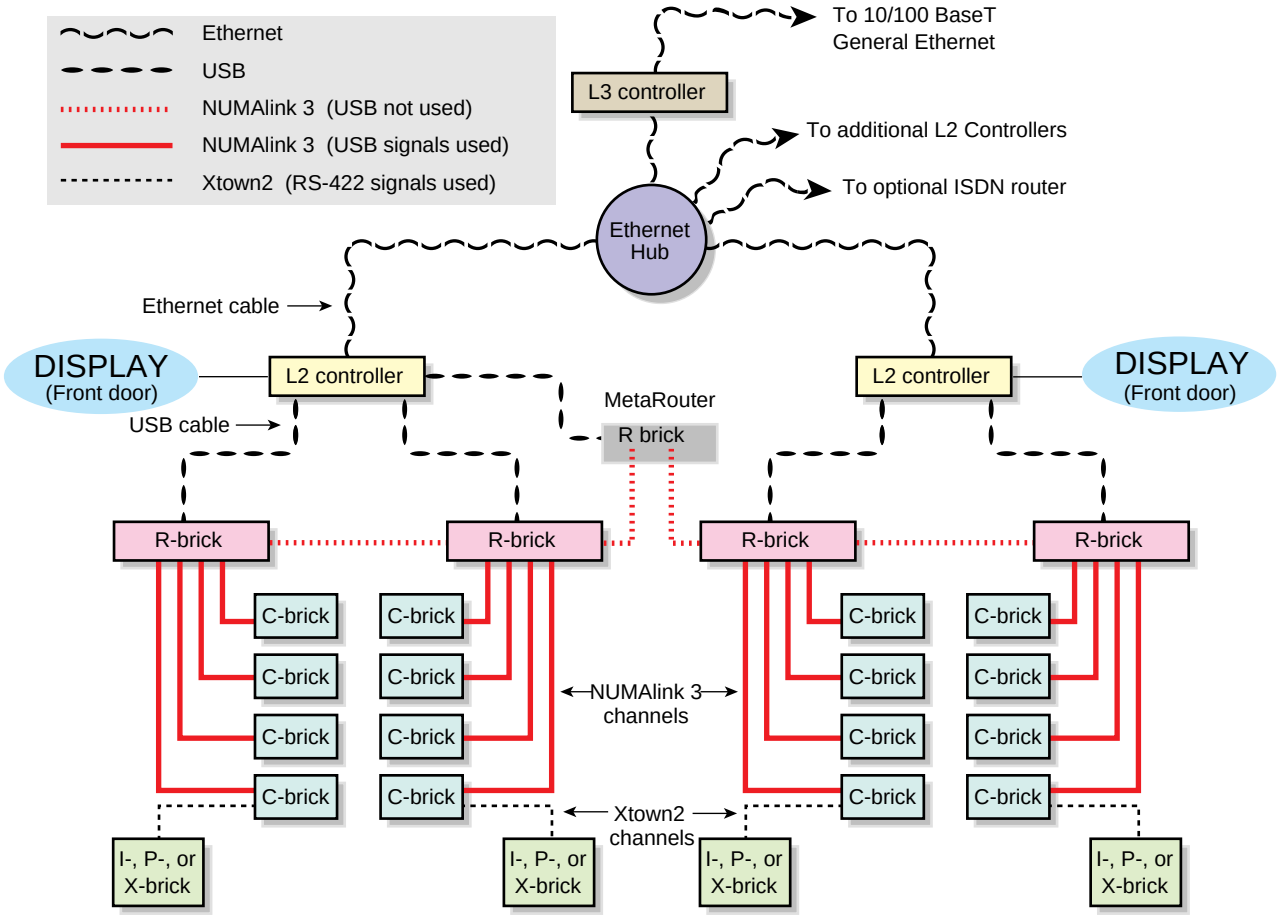

Figure 3-27 SGI Origin 3800 System Control Network (Typical)

3.18.4 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 processors 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.

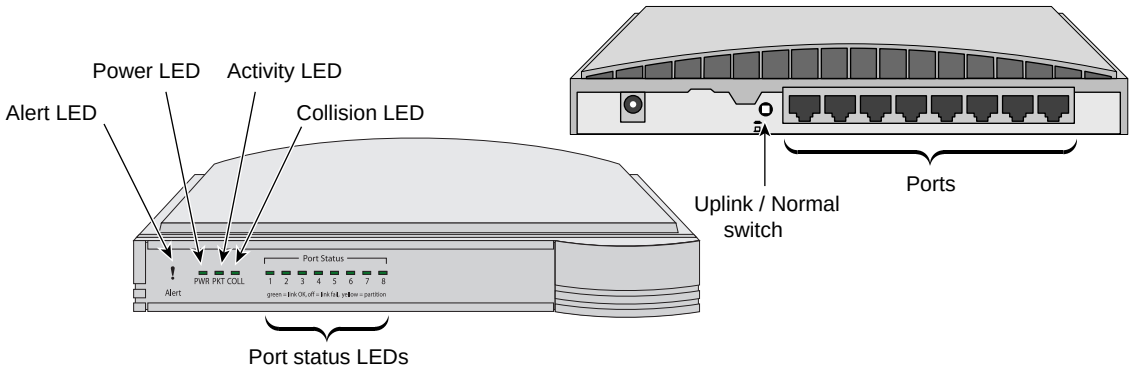

Figure 3-28 Eight-Port Ethernet Hub

The Ethernet hub has the following features:

- Eight RJ45 ports (port 8 is controlled by Uplink / Normal switch)
- Maximum 100 meter (328 ft.) 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 3-29, Figure 3-30, and Figure 3-31 illustrate how Ethernet hubs are cabled to the L2 controllers in various system sizes.

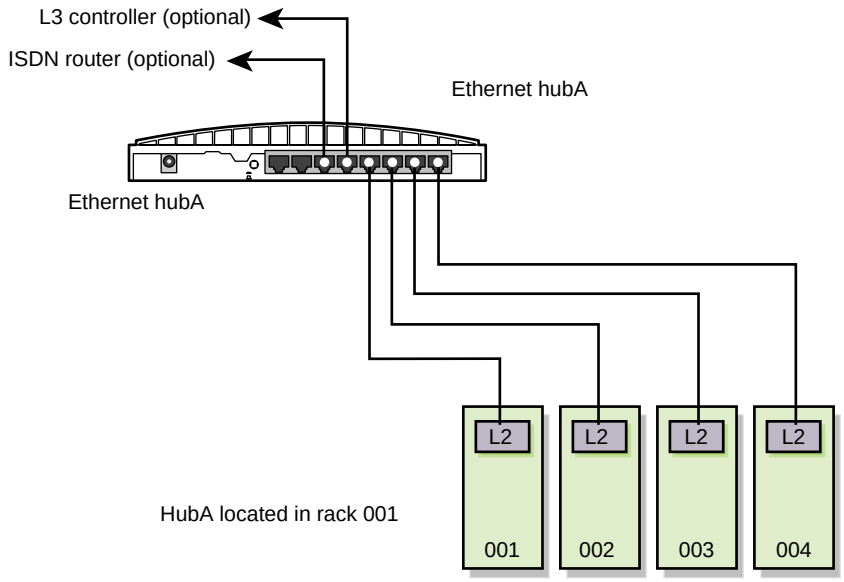

Figure 3-29 Single-hub Network in 128-processor System

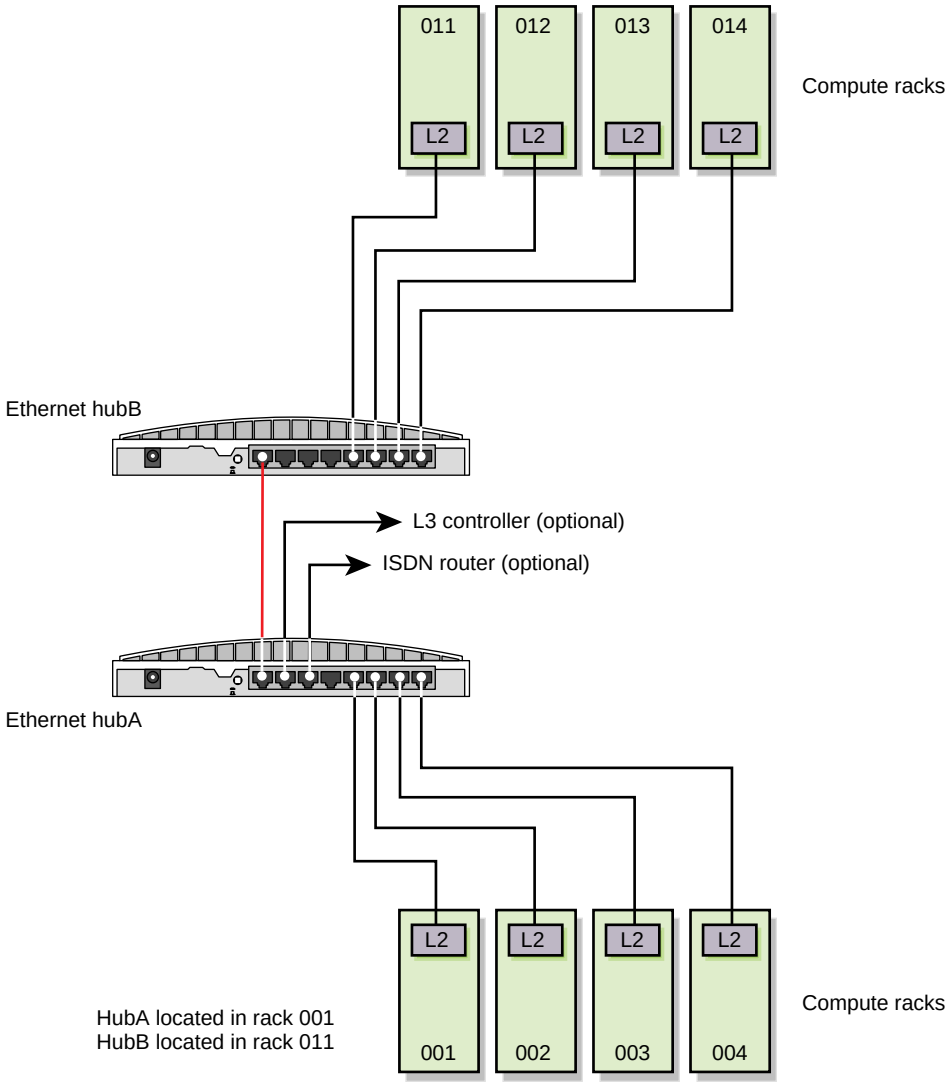

Figure 3-30 Multiple-hub Network in 256-processor System

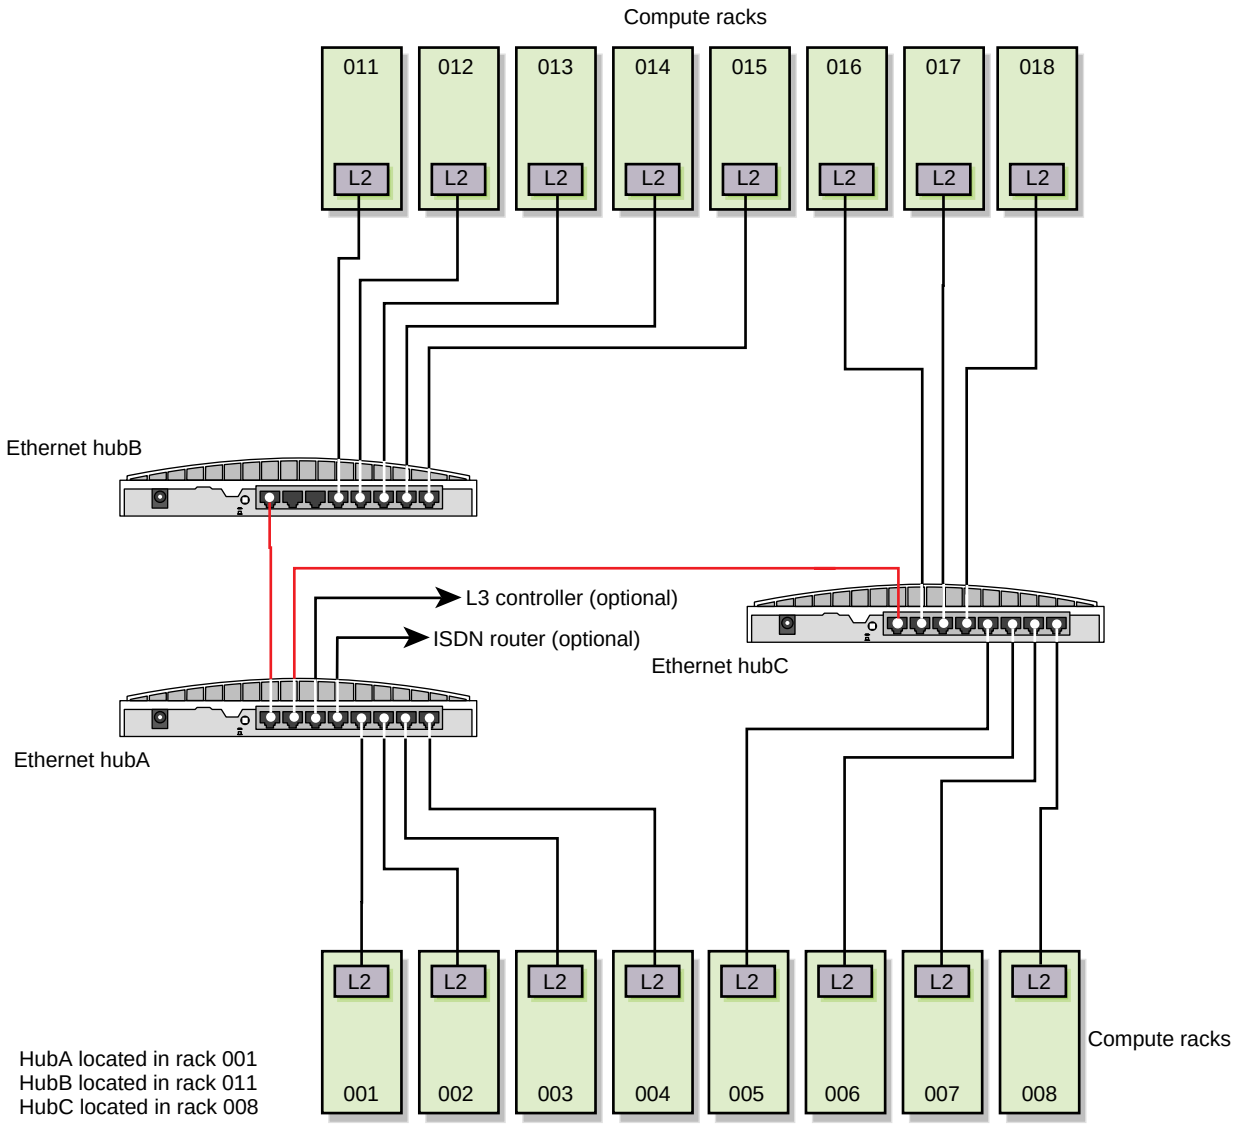

Figure 3-31 Multiple-hub Network in 512-processor System

3.19 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 3-32 ISDN Router (Rear View)

3.20 System Cooling

All SGI Origin 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. 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-, X-bricks) to determine the sustainable channel bandwidth.

4.1 C-brick (Compute Node)

Each C-brick consists of two or four MIPS R12000A processors running at 400 MHz, each with an 8-MB secondary cache. Each processor is capable of executing two floating-point instructions per cycle, which supports a peak speed of 800 MFLOP/s. The peak and sustained bandwidths for each of the C-bricks channels are listed in Table 4-1. The memory is a distributed shared memory (DSM) scheme, in which the memory is physically partitioned among the nodes but is accessible by all nodes. 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 ful-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 MBs/ full-duplex ~560 MB/s each direction
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 4-8 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 Bandwidth
18-GB FC Disk Drive	(TBD)
36-GB FC Disk Drive (10K RPM)	~ 30 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 1200-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	2 modes, software selectable 800 MB/s or 1.2 GB/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 36. 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 1.2 GB/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 X-brick (XIO Expansion)

The X-brick provides four expansion slots for legacy XIO interface cards, such as HIPPI, digital video, GSN. For more information about the X-brick, refer to the “XIO Expansion (X-brick)” on page 38. Refer to Table 4-5 for the bandwidth values.

Table 4-5 Bandwidth Characteristics of the X-brick

Description	Peak Bandwidth
Xtown2 Ports A and B	2 modes, software selectable 800 MB/s or 1.2 GB/s
XIO Ports	800 MB/s full-duplex

4.6 Supported XIO Cards

The X-brick supports the XIO cards that are listed in Table 4-6 for First Customer Ship (FCS). Additional XIO cards will be added to this table as they qualify.

Table 4-6 Supported XIO Cards

Marketing Code	Description	Bandwidth (Peak)
XT-DIVO	Digital video XIO card	TBD
XT-DIVO-DVC	Digital video I/O DVC Onyx2	TBD
XT-FDDI-D	XIO 1-port FDDI Dual Attach Station (DAS) (The FDDI card is a PCI card installed in a shoehorn adapter.)	TBD
XT-GSN-C-1XIO	GSN adapter single port XIO card copper	800 MB/s (Full Duplex)
XT-GSN-C-2XIO	GSN adapter two port XIO card copper (Note: This is a two-card set that uses two XIO slots in the X-brick)	800 MB/s per port (1600 MB/s total) (Full Duplex)
XT-HD	XIO high definition video I/O	TBD
XT-HIPPI-800-SER	Single-port serial HIPPI XIO card	100 MB/s
XT-VME-6U	XIO to VME adapter 6 U	TBD
XT-VME-9U	XIO to VME adapter 9 U	TBD
XT-ATM-OC3C-4P	Four port ATM OC3	19 MB/s per port

Table 4-7 XIO Card - Brick and System Maximum Quantities

Marketing Code	Minimum Required IRIX Level	Maximum Cards per Brick	Maximum Cards per System (Fully Tested)	Maximum Cards per System (Goal - not tested)
Networking Controllers				
XT-ATM-OC3C-4P	6.5.9	4	0	16
	6.5.10	4	0	16
	6.5.11	4	N/A	16
XT-FDDI-D	6.5.9	4	4	6
	6.5.10	4	4	6
	6.5.11	4	N/A	6
XT-GSN-C-1XIO-EA	6.5.9	2	2	2
	6.5.10	2	2	2
	6.5.11	2	N/A	2
XT-GSN-C-2XIO-EA	6.5.9	2	2	2
	6.5.10	2	2	2
	6.5.11	2	N/A	2
Miscellaneous Controllers				
XT-DIVO	6.5.9	TBD	TBD	TBD
XT-DIVO-DVC	6.5.9	TBD	TBD	TBD
XT-HD	6.5.9	TBD	TBD	TBD
XT-HIPPI-800-SER	6.5.9	4	4	16
XT-VME-6U	6.5.9	4	4	4
XT-VME-9U	6.5.9	4	4	4

Note: Shaded boxes indicate card is not fully qualified.

4.7 Supported PCI Cards

The PCI cards listed in Table 4-8 are supported in the I, P, and X-bricks for First Customer Ship (FCS). Additional PCI cards will be added to this table as they qualify.

Table 4-8 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-LVD-2P	QLogic QLA1280	33-MHz / 64-bit	Two-port ultra2 SCSI low voltage differential	80 MB/s (wide) 40 MB/s (narrow)
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
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

Table 4-8 Supported PCI Cards

Marketing Code	Vendor's P/N	PCI Bus Clock / Data	Description	Bandwidth (Peak) (per port)
Miscellaneous				
PCI-AUD-D1000	SGI	33-MHz / 32-bit	PCI 8 Channel Digital Audio Card	2.88 MB/s
PCI-SER-10002	SGI	33-MHz / 32-bit	Two-port RS-232/RS422 Serial Port Card	.0143 MB/s

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

Table 4-9 PCI Card - Brick and System Maximum Quantities

Marketing Code	Minimum Required IRIX Level	Maximum Cards per Brick	Maximum Cards per System (Fully Tested)	Maximum Cards per System (Goal - not tested)
Storage Controllers **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.				
PCI-SCSI-DF-2P	6.5.9	I-brick - 2 per bus (04 total) P-brick - 2 per bus (12 total)	24 see note	84
	6.5.10	I-brick - 2 per bus (04 total) P-brick - 2 per bus (12 total)	24 see note	84
	6.5.11	I-brick - 2 per bus (04 total) P-brick - 2 per bus (12 total)	N/A	84
PCI-SCSI-LVD-2P	6.5.9	I-brick - 2 per bus (04 total) P-brick - 2 per bus (12 total)	12 see note	42
	6.5.10	I-brick - 2 per bus (04 total) P-brick - 2 per bus (12 total)	12 see note	42
	6.5.11	I-brick - 2 per bus (04 total) P-brick - 2 per bus (12 total)	N/A	42
PCI-FC-1POPT-A	6.5.9	I-brick - 2 per bus (04 total) P-brick - 2 per bus (12 total)	24 see note	84
	6.5.10	I-brick - 2 per bus (04 total) P-brick - 2 per bus (12 total)	24 see note	84
	6.5.11	I-brick - 2 per bus (04 total) P-brick - 2 per bus (12 total)	N/A	84
PCI-FC-1PCOP-A	6.5.9	I-brick - 2 per bus (04 total) P-brick - 2 per bus (12 total)	24 see note	84
	6.5.10	I-brick - 2 per bus (04 total) P-brick - 2 per bus (12 total)	24 see note	84
	6.5.11	I-brick - 2 per bus (04 total) P-brick - 2 per bus (12 total)	N/A	84

Table 4-9 PCI Card - Brick and System Maximum Quantities

Marketing Code	Minimum Required IRIX Level	Maximum Cards per Brick	Maximum Cards per System (Fully Tested)	Maximum Cards per System (Goal - not tested)
Networking Controllers **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. **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.				
PCI-GIGENET-OR	6.5.9	I-brick - 1 on bus 2 (01 total) P-brick - 1 per bus (06 total)	10 see note	30
	6.5.11	I-brick - 1 on bus 2 (01 total) P-brick - 1 per bus (06 total)	20 see note	30
PCI-GIGENET-C	6.5.9	I-brick - 1 on bus 2 (01 total) P-brick - 1 per bus (06 total)	10 see note	30
	6.5.11	I-brick - 1 on bus 2 (01 total) P-brick - 1 per bus (06 total)	20 see note	30
PCI-ATMOC3-1P (see ATM note above)	6.5.9	I-brick - 2 on bus 2 (02 total) P-brick -1 per bus (06 total)	8 see note	16
	6.5.10	I-brick - 2 on bus 2 (02 total) P-brick -1 per bus (06 total)	8 see note	16
PCI-ATMOC12-1P (see ATM note above)	6.5.9	I-brick - 1 on bus 2 (01 total) P-brick -1 per bus (06 total)	8 see note	16
	6.5.10	I-brick - 1 on bus 2 (01 total) P-brick -1 per bus (06 total)	8 see note	16
Miscellaneous Controllers				
PCI-AUD-D1000	6.5.9	I-brick - 2 on bus 1 (02 total) I-brick - 2 on bus 2 (02 total) P-brick -2 per bus (06 total)	6	6
PCI-SER-10002	6.5.9	TBD	TBD	TBD

4.8 Supported Tape Drives

The tape drives listed in Table 4-10 are available at first customer ship (FCS).

Table 4-10 Supported Tape Drive Subsystems

Peripheral Product Designator	Interface Type	Storage Capacity (Mbytes)	Transfer Rate (MBytes)
DLT 7000	SCSI only	35,000	5.0
FUJITSU M2483/M2485/M2488	SCSI only	24,000	3.0
IBM 3480/3490/3490E	SCSI only	800	4.5
IBM 3590 Magstar	SCSI only	20,000	9.0
STK 9840	SCSI and FC	20,000	9.0
STK 4480 (18 track)	SCSI only	200	3.0
STK 4490 Silverton (36 track)	SCSI only	800	3.0
STK 9490 Timberline (36 track)	SCSI only	800	6.0
SONY AIT-2	SCSI only	60,000	6.0

4.9 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 4-11 for the bisection bandwidth for common system sizes.

Table 4-11 Bisection Bandwidths

Number of processors	System Bisection Bandwidth	Bisection Bandwidth per CPU
8	3.2 Gbytes/s	0.4 Gbytes/s
16	6.4 Gbytes/s	0.4 Gbytes/s
32 Note 1	6.4 Gbytes/s	0.2 Gbytes/s
32 Note 2	12.8 Gbytes/s	0.4 Gbytes/s
64	12.8 Gbytes/s	0.2 Gbytes/s
128	12.8 Gbytes/s	0.1 Gbytes/s
256	25.6 Gbytes/s	0.1 Gbytes/s
512	51.2 Gbytes/s	0.1 Gbytes/s

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

Note 2: This configuration connects the two routers together with four NUMAlink cables (high bandwidth model).

5. System Interconnection Fabric Drawings

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

5.1 2-processor through 16-processor Systems

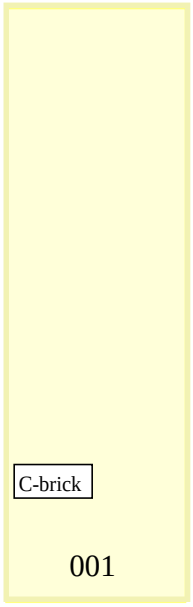

2 or 4 Processors

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

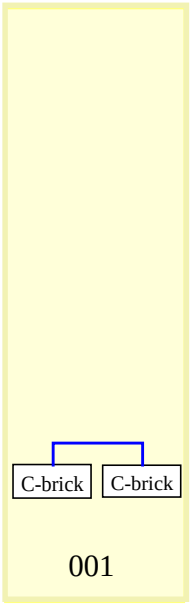

4 to 8 Processors

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

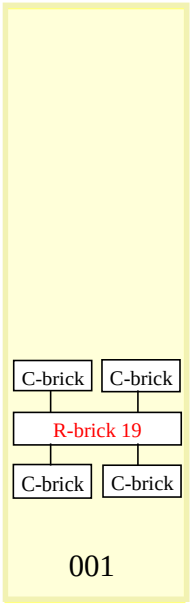

8 to 16 Processors

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

5.2 32-processor through 64-processor Systems

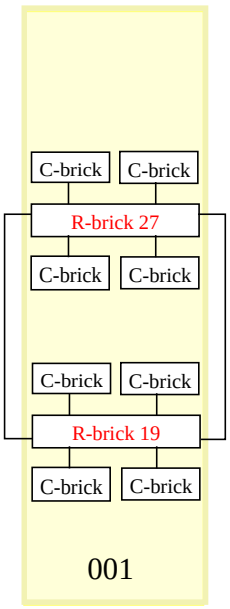

32 Processors

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

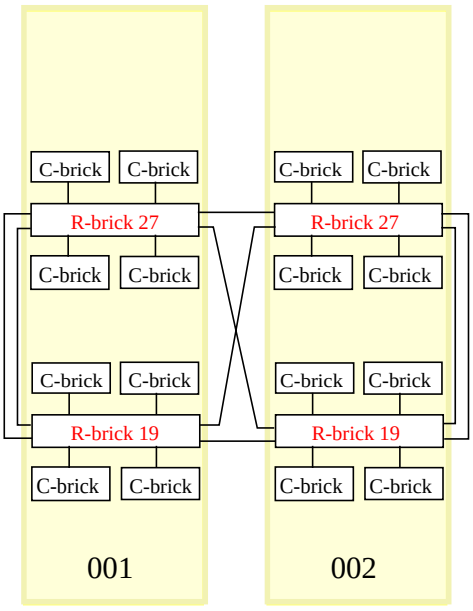

64 Processors

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

5.3 96-processor System

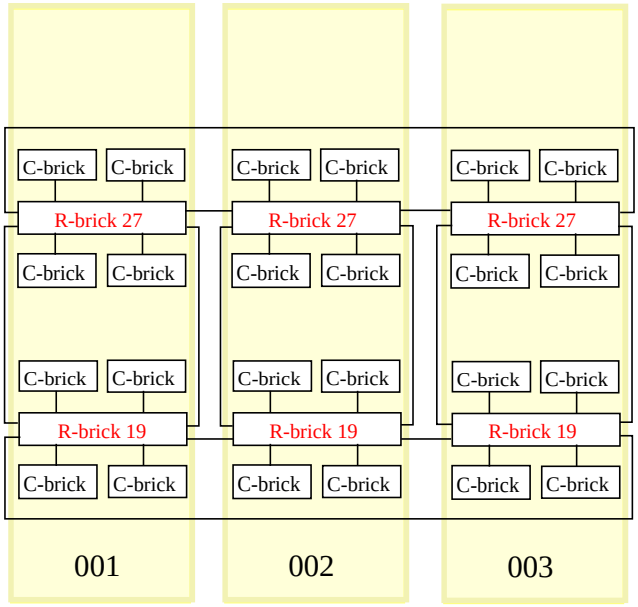

Quantities			
3	Racks	6	R-bricks
36	Cables	24	C-bricks

5.4 128-processor System

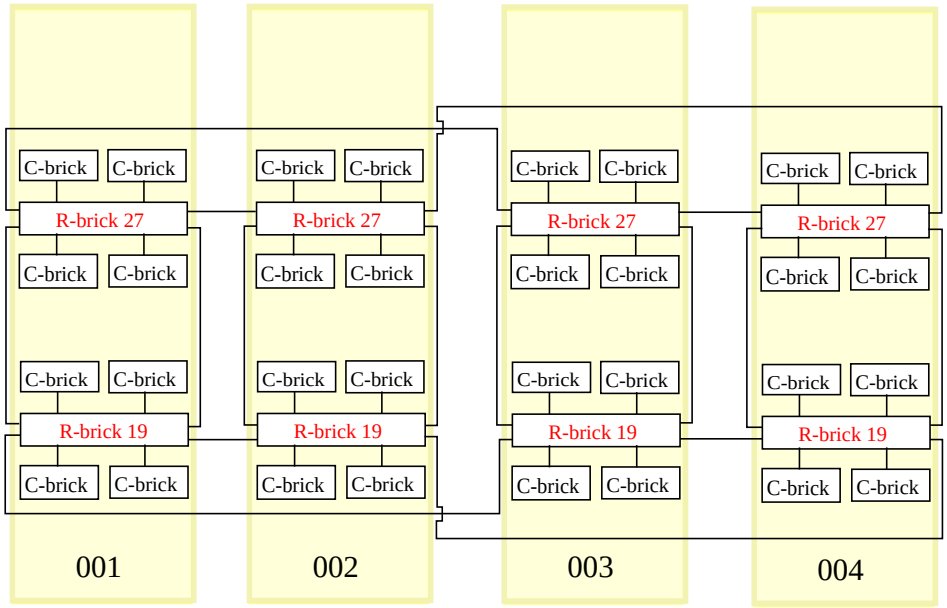

Quantities			
4	Racks	8	R-bricks
48	Cables	32	C-bricks

5.5 160-processor System

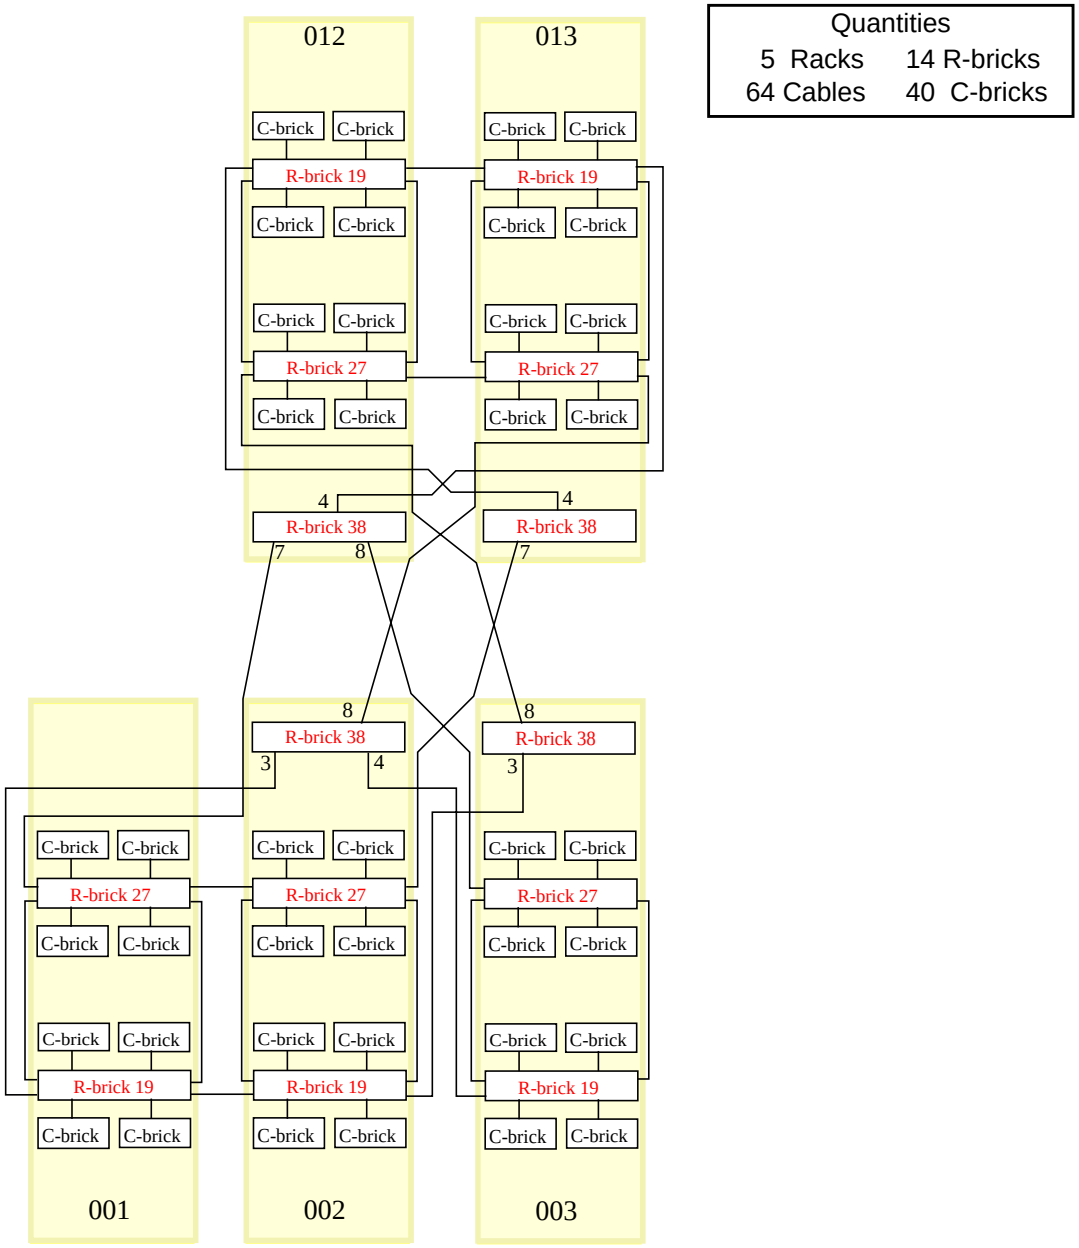

5.6 192-processor System

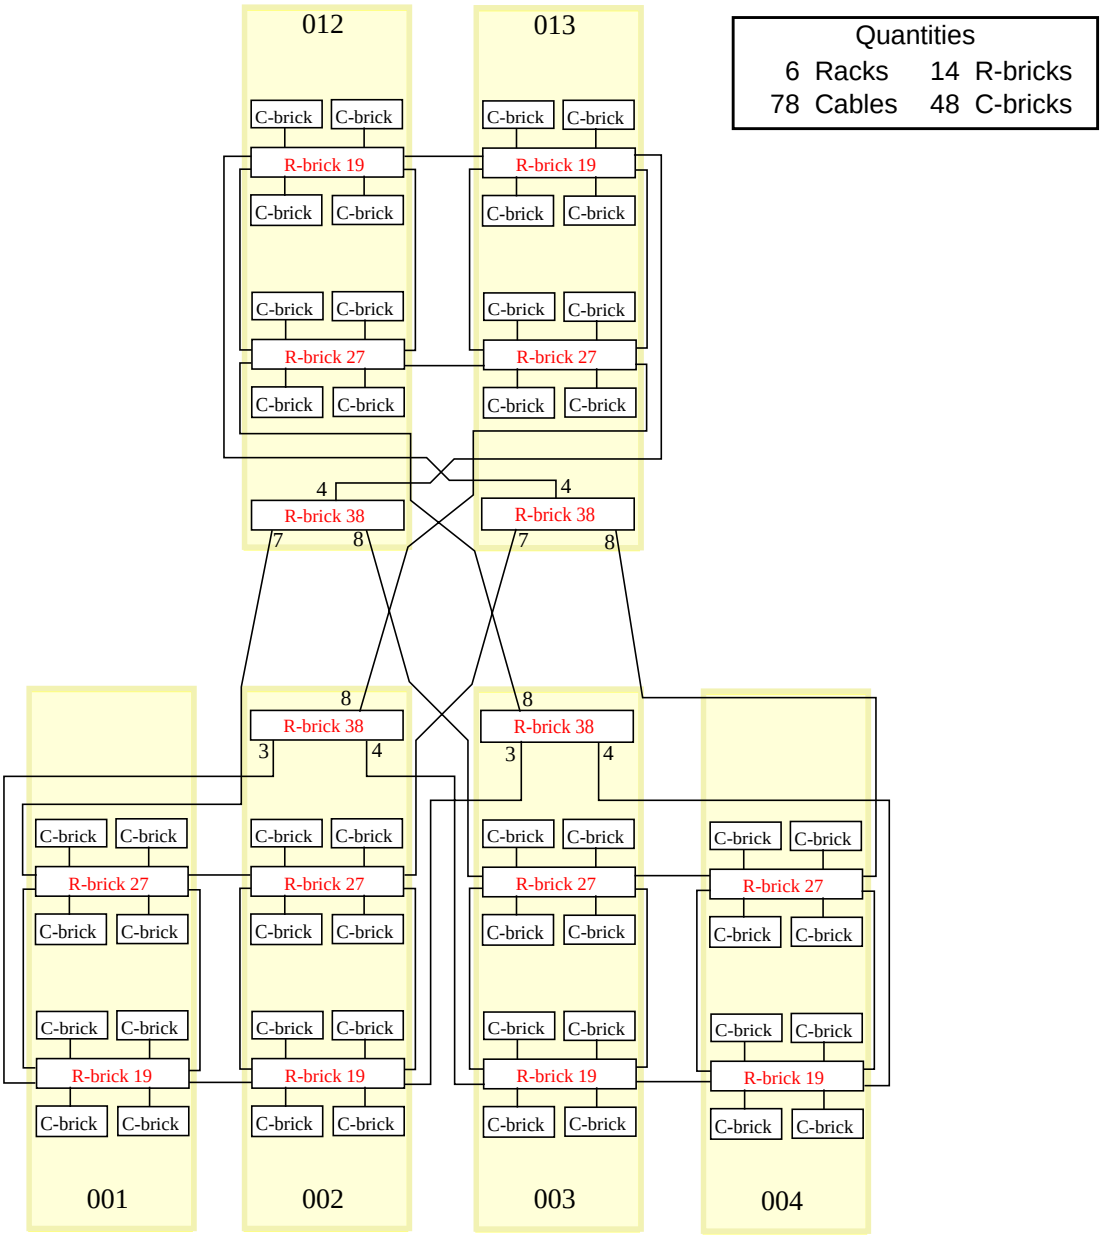

5.7 224-processor System

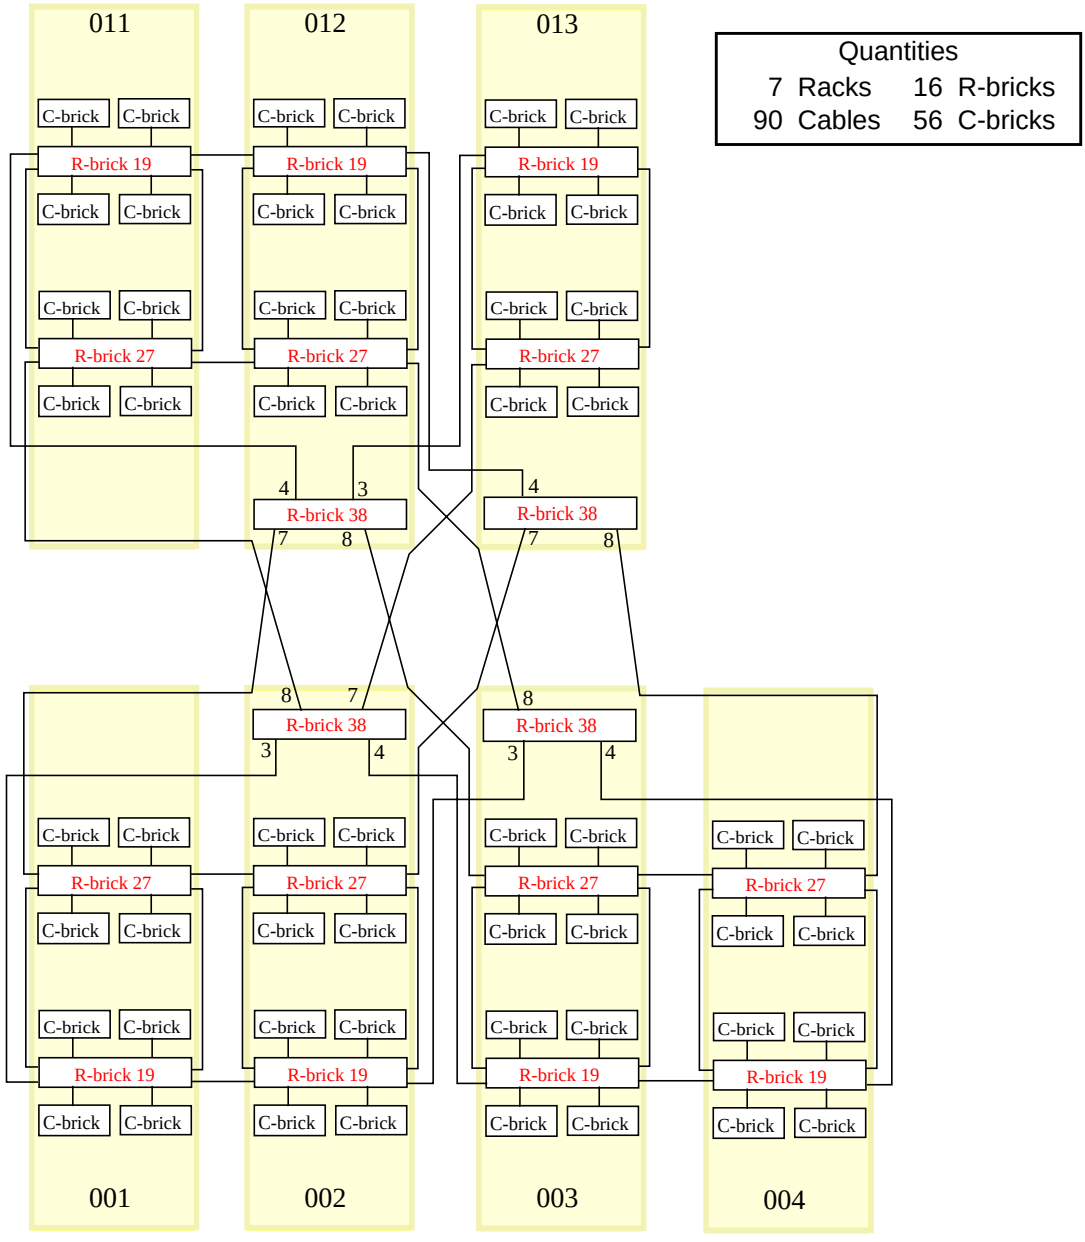

5.8 256-processor System

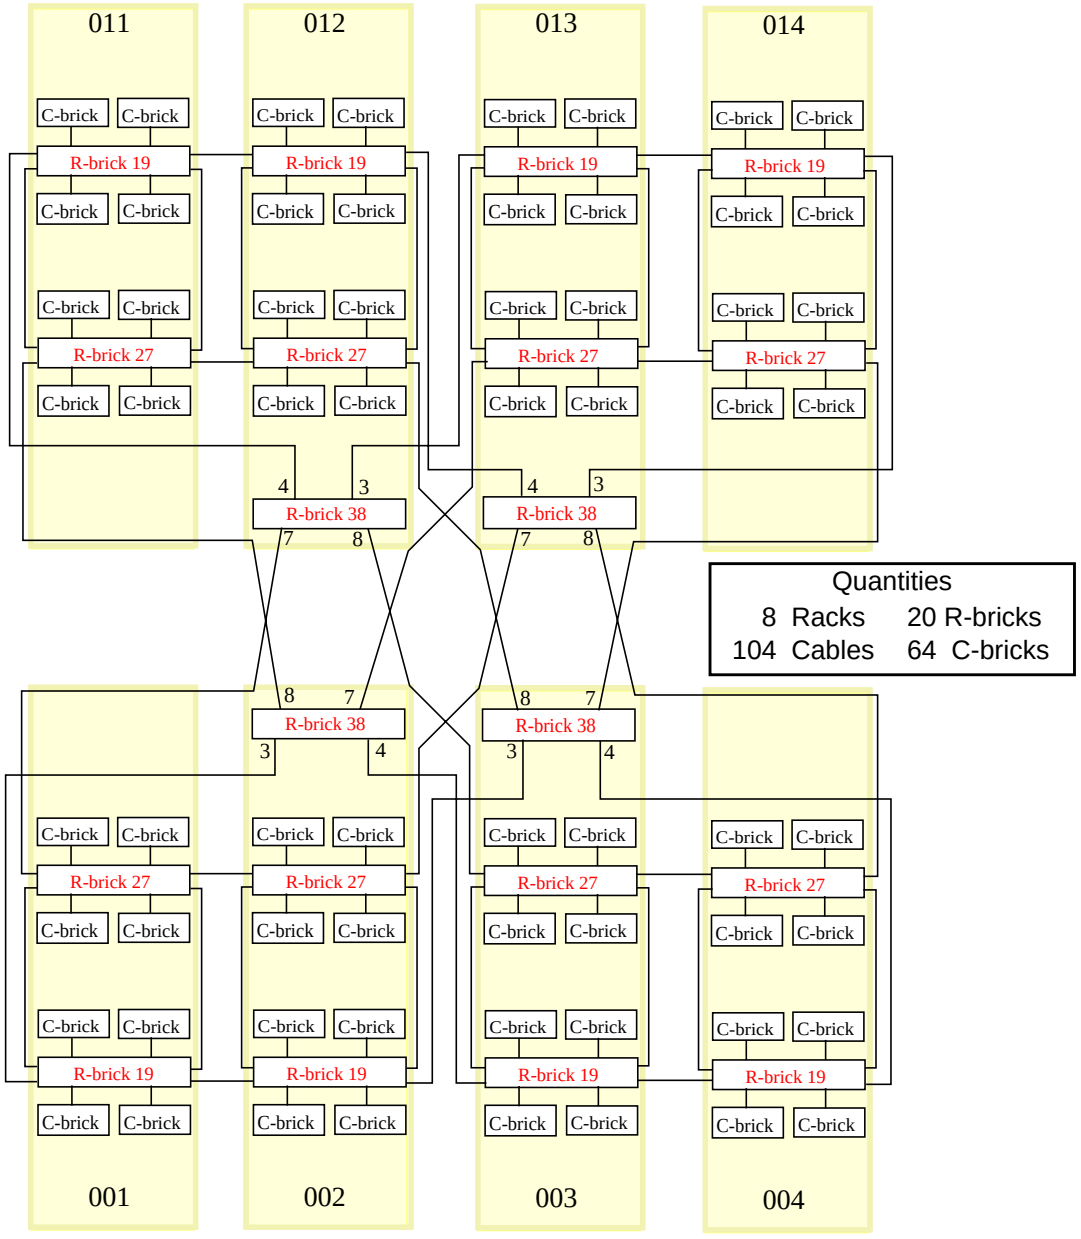

5.9 288-processor System

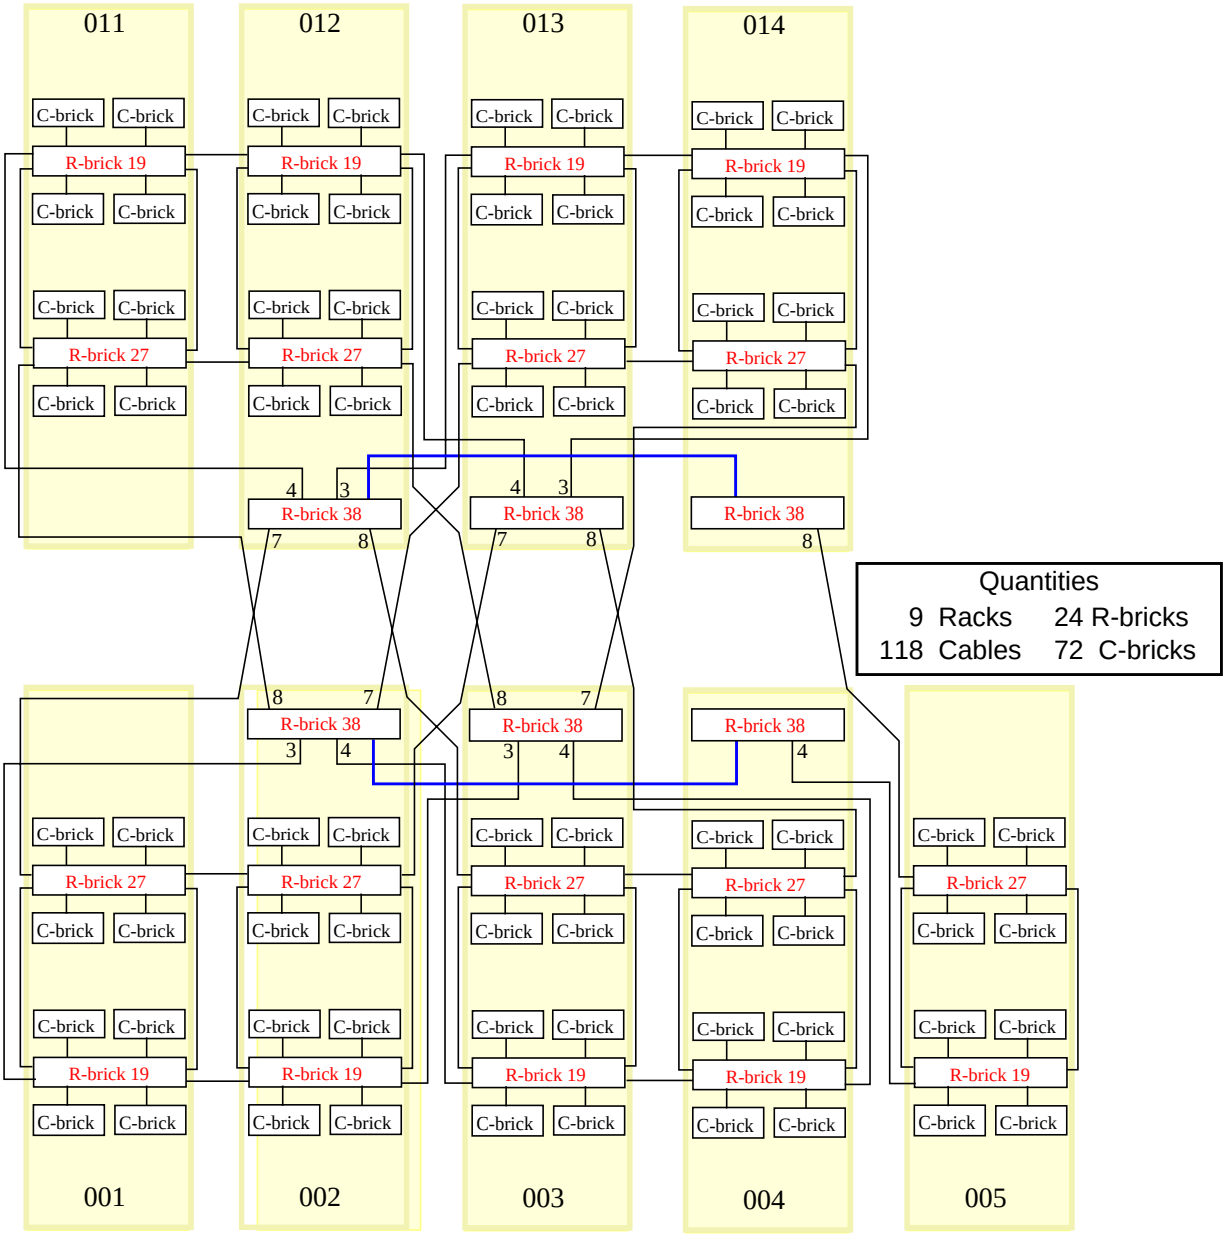

5.10 320-processor System

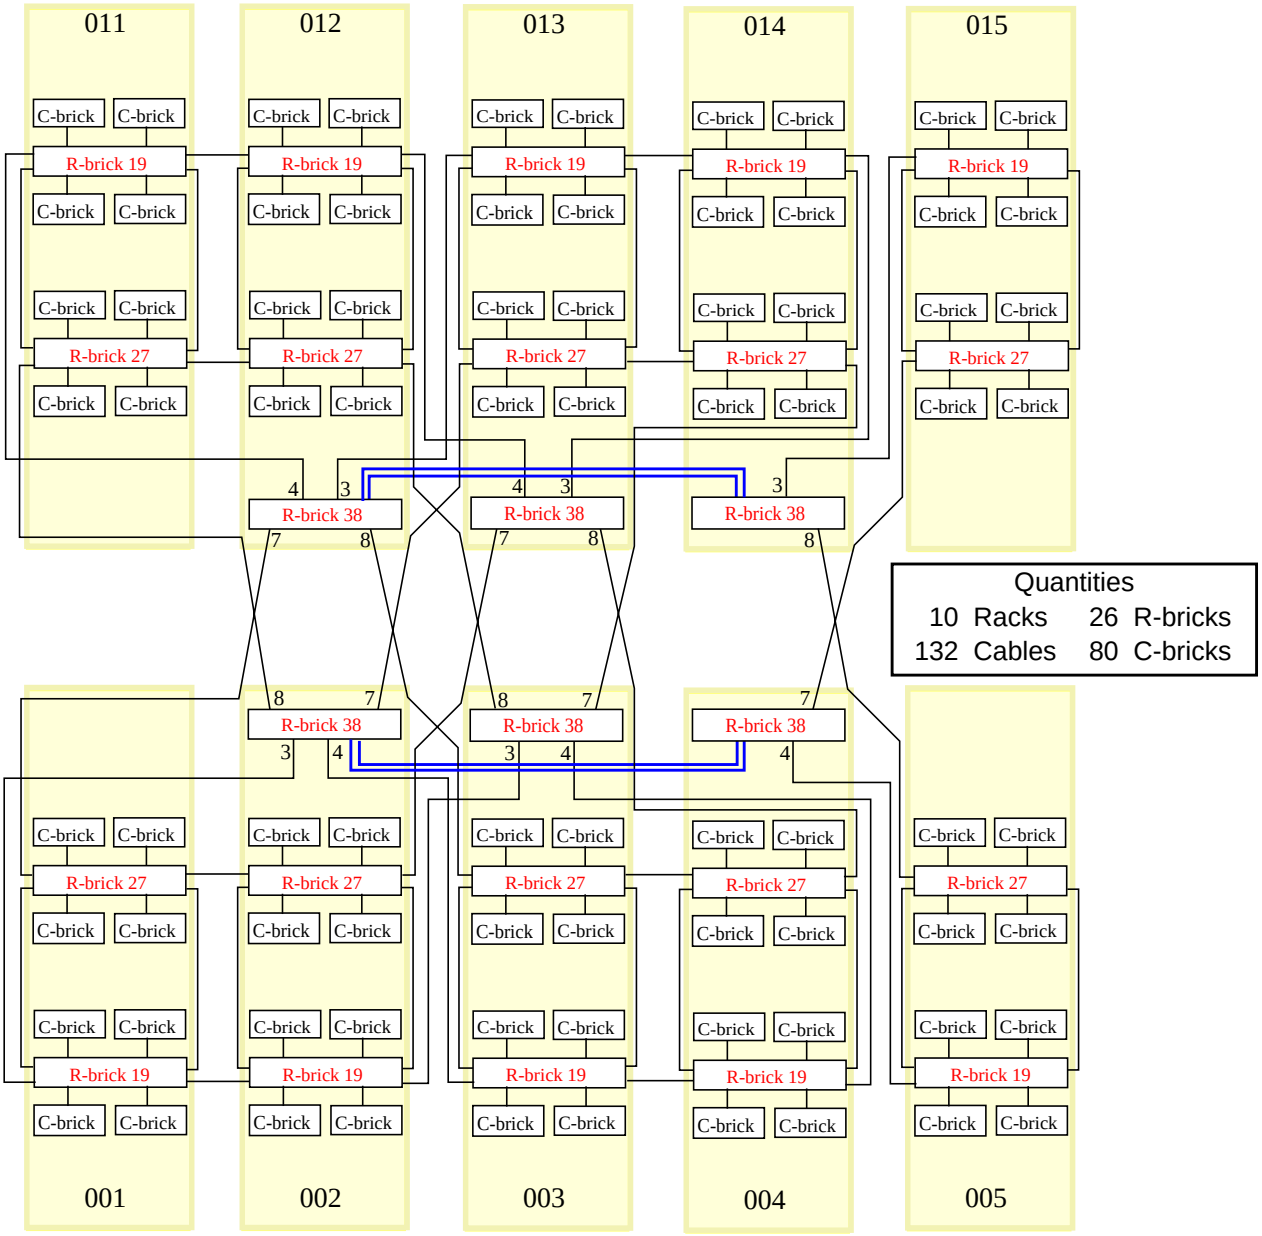

5.11 352-processor System

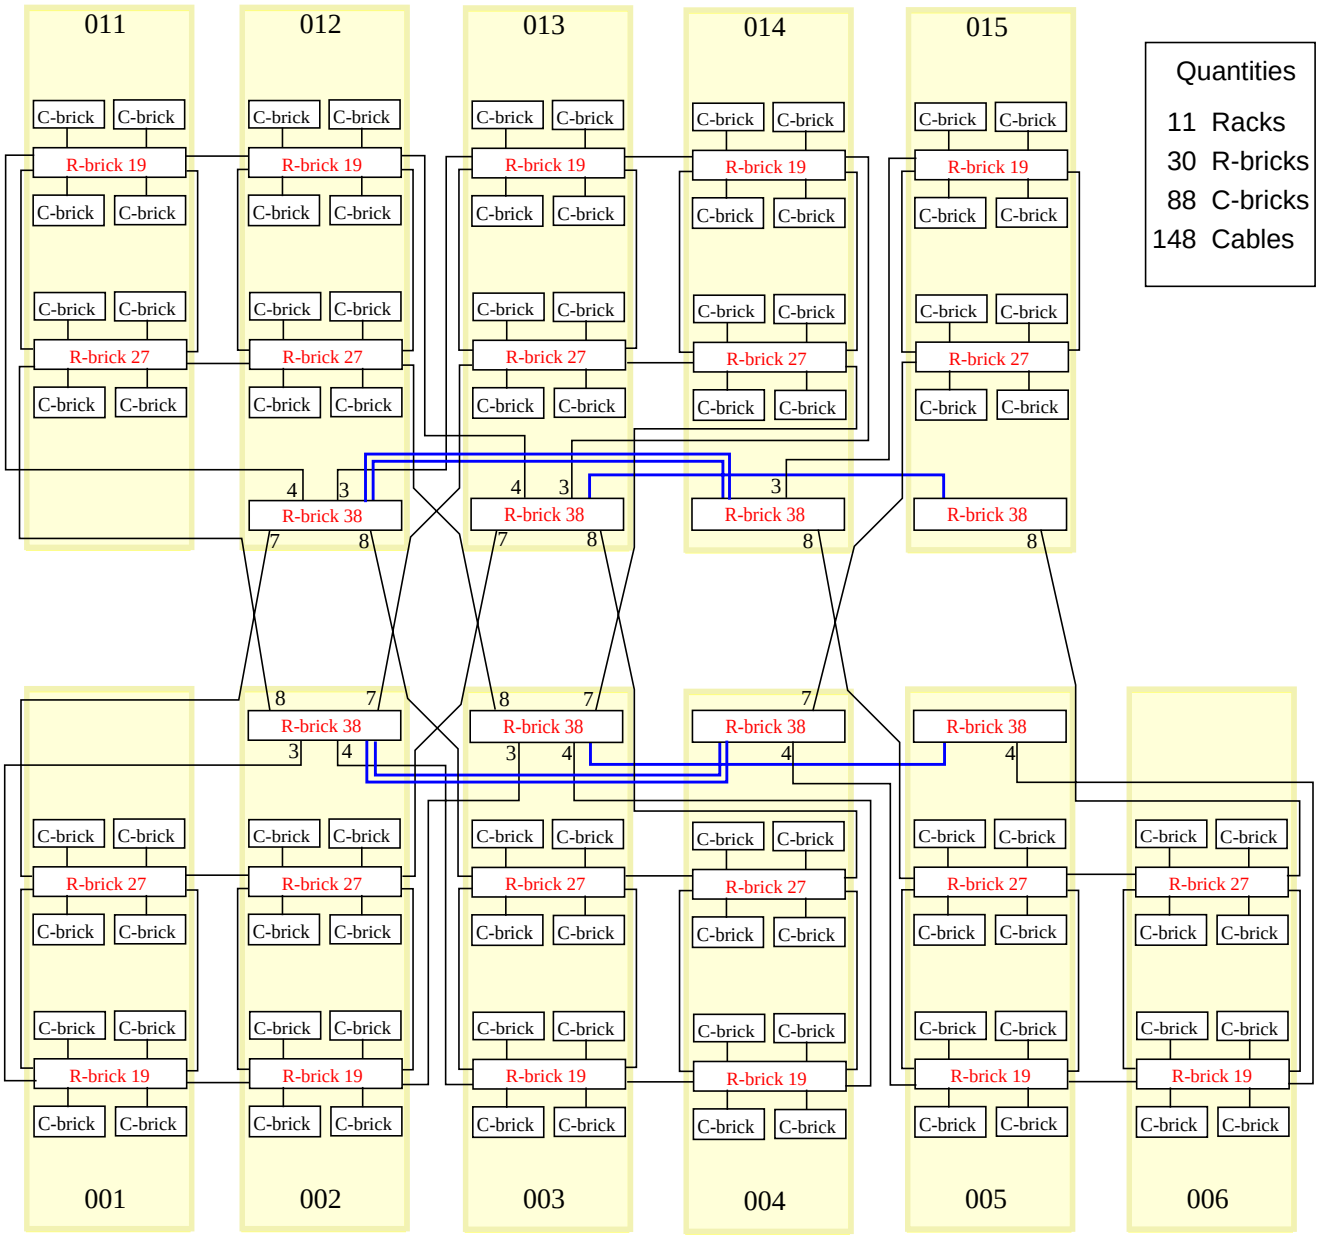

5.12 384-processor System

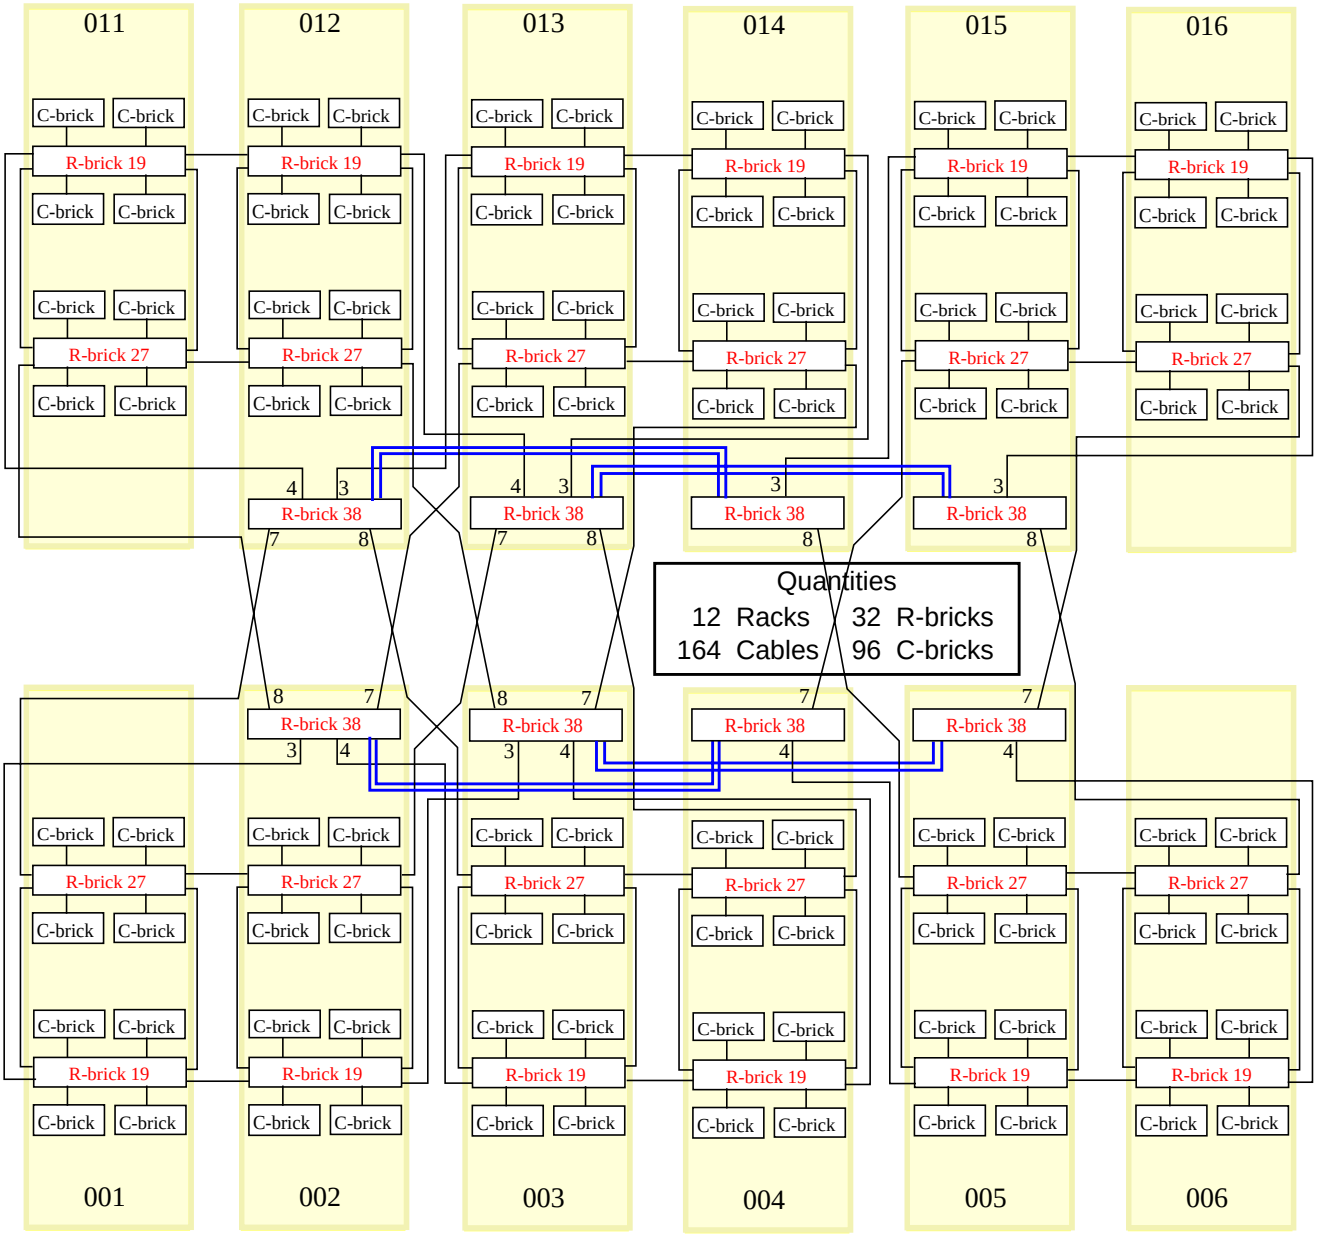

5.13 416-processor System

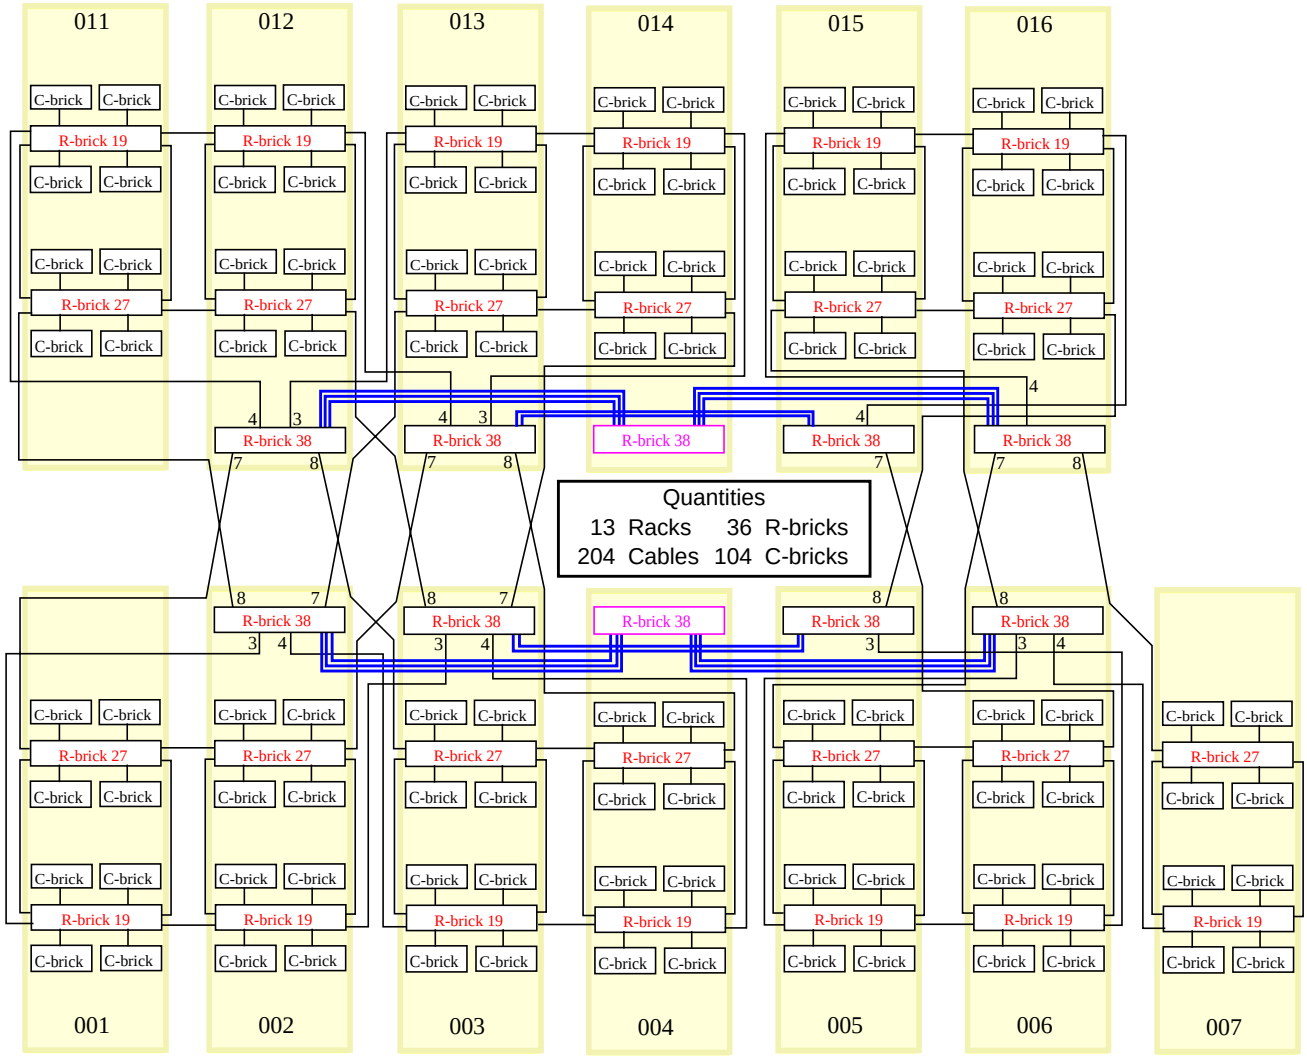

5.14 448-processor System

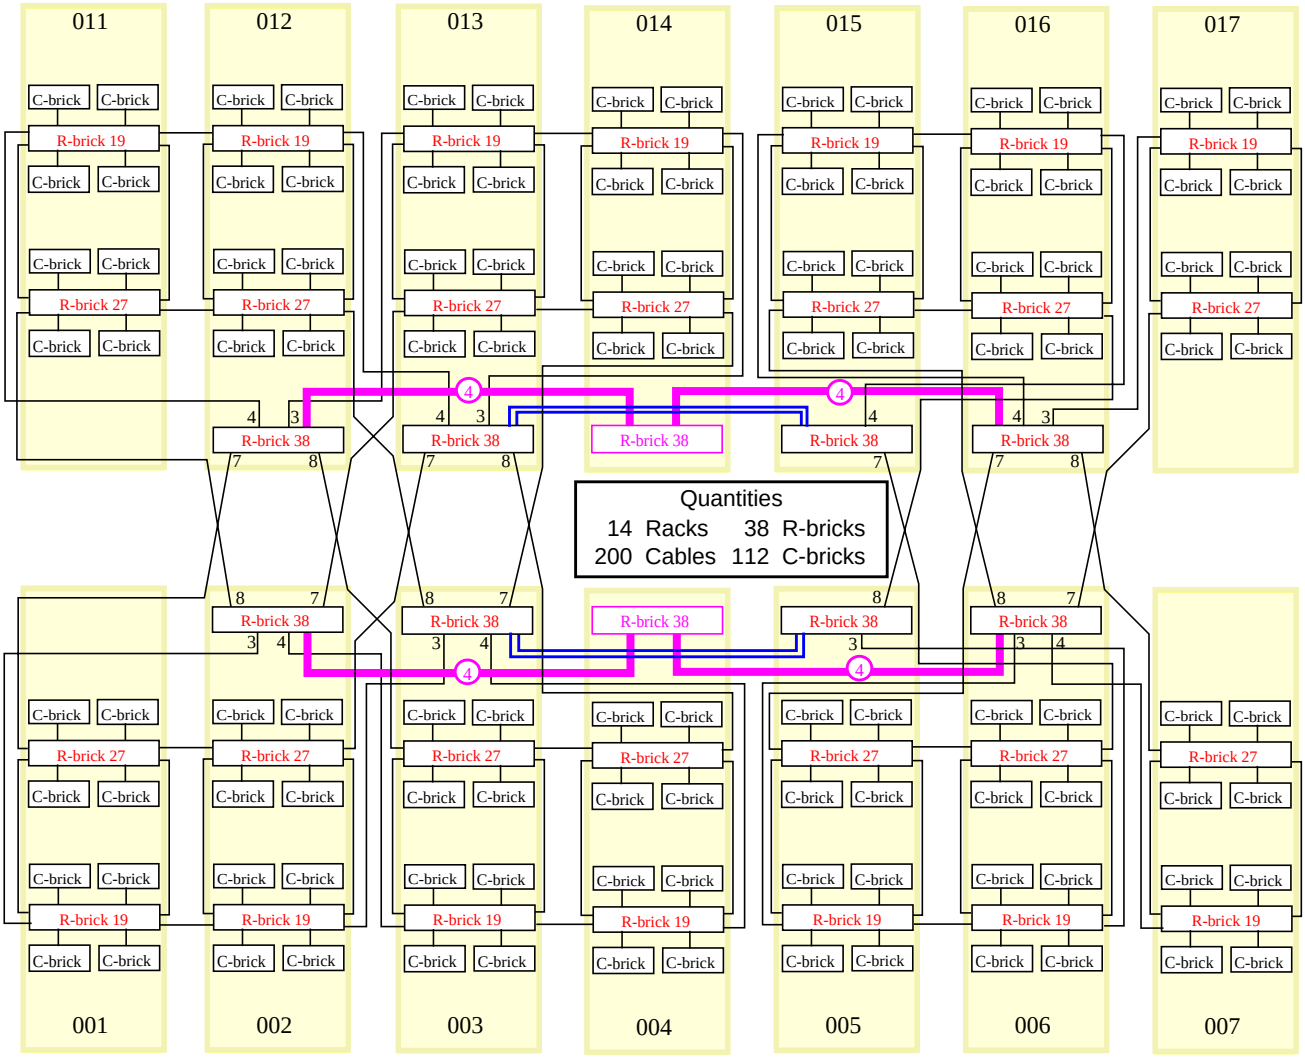

5.15 480-processor System

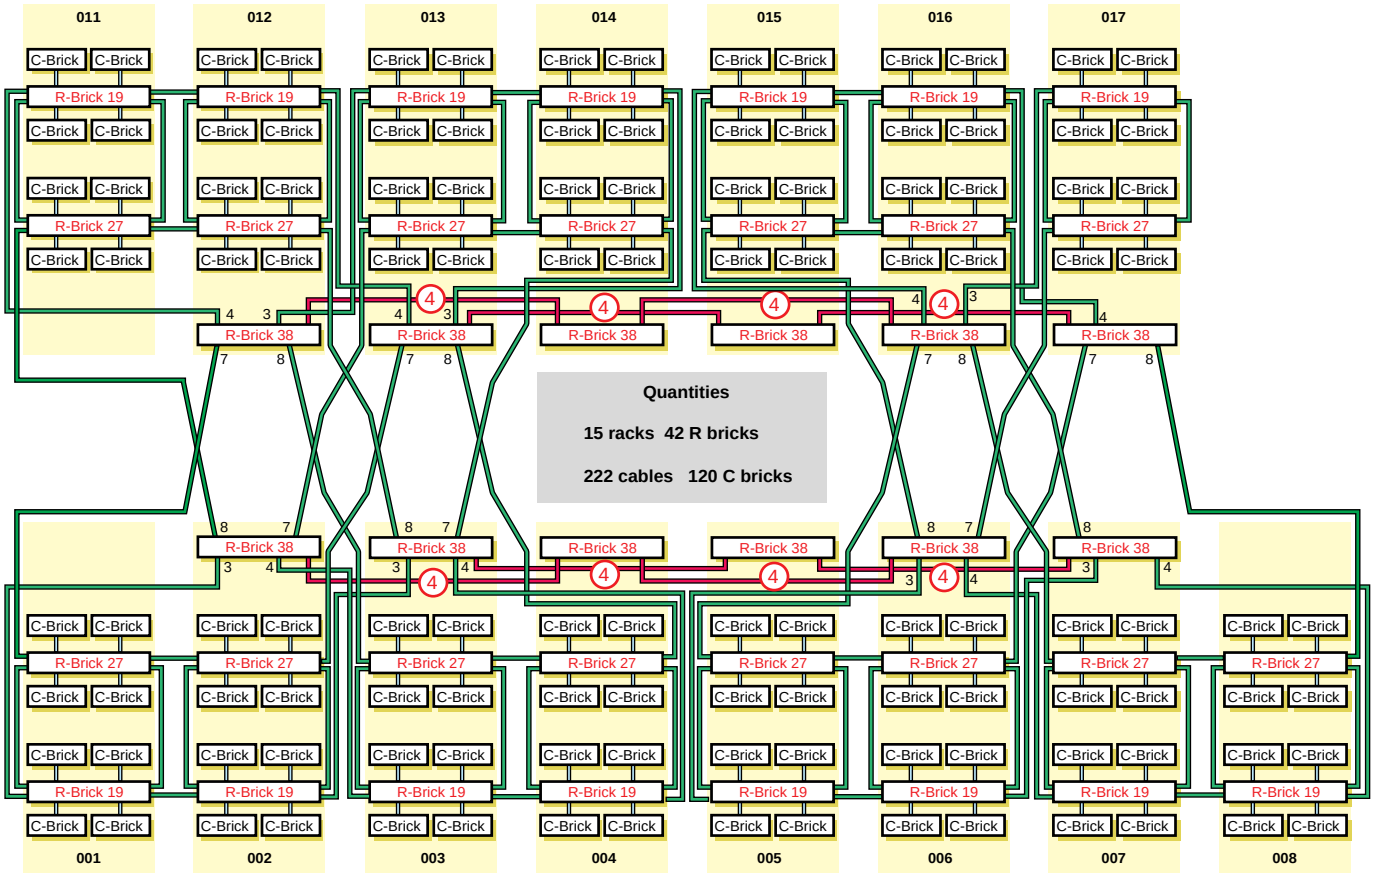

5.16 512-processor System

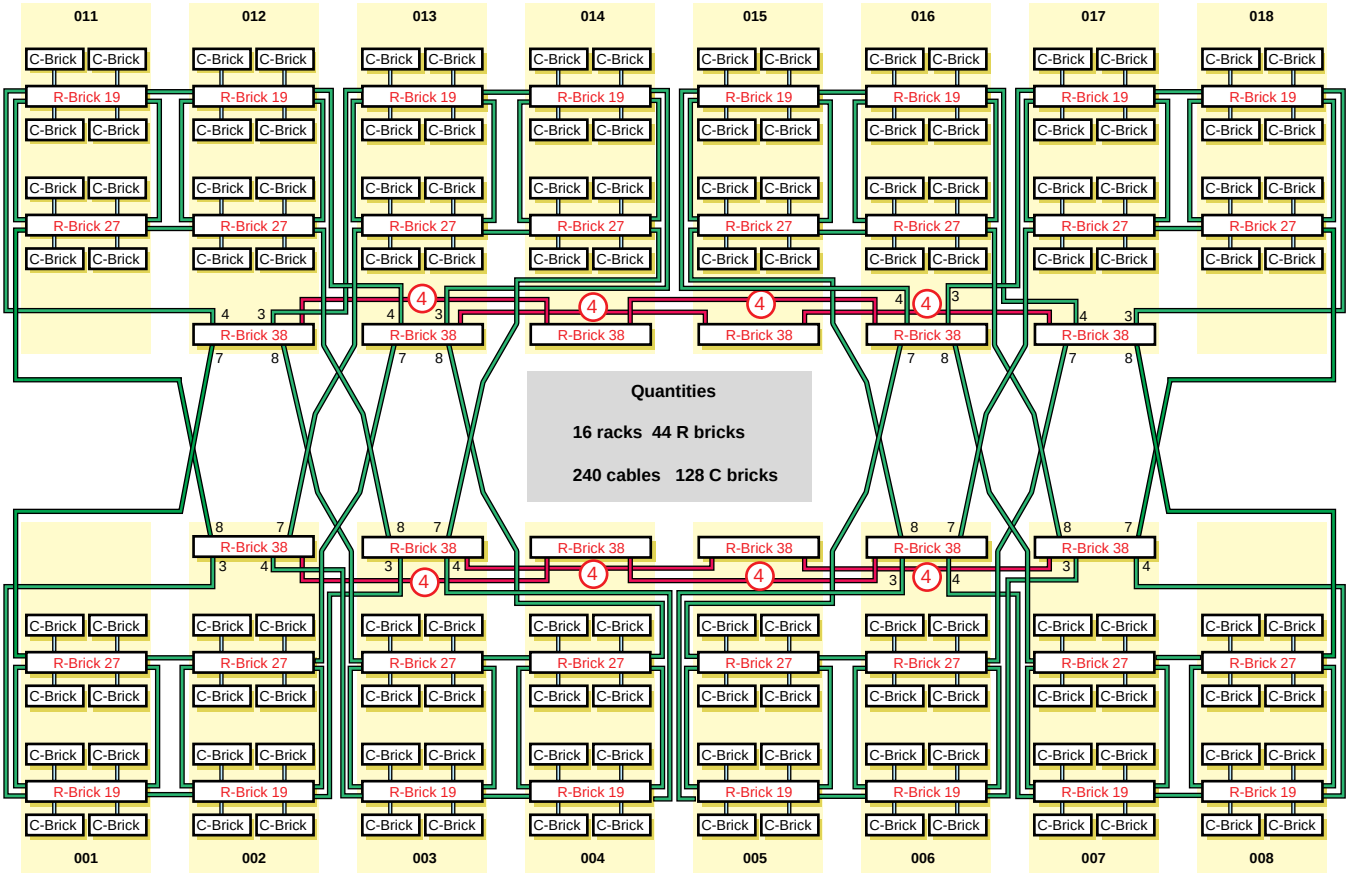

6. System Partitioning

6.1 Partitioning Overview

Partitioning is defined as the ability to take a multi-node distributed shared memory (DSM) system, such as one of the SGI Origin 3000 series, 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.

6.2 Definition of Terms

There has been much discussion on the terminology used to described 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 IRIX 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 Origin 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 total 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 Origin 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 2 partitions	with with	2 C-bricks 1 C-brick	8 processors 4 processors
	1 partition 1 partition	with with	1 C-brick 3 C-bricks	4 processors 12 processors
	2 partitions	with	2 C-bricks	8 processors
8	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 1 partition	with with	1 C-brick 15 C-bricks	4 processors 60 processors
	1 partition 2 partitions	with with	8 C-bricks 4 C-bricks	32 processors 16 processors
	1 partition 1 partition 2 partitions	with with with	8 C-bricks 4 C-bricks 2 C-bricks	32 processors 16 processors 8 processors
	1 partition 4 partitions	with with	8 C-bricks 2 C-bricks	32 processors 8 processors
	8 partitions	with	2 C-bricks	8 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
32	1 partition	with	32 C-bricks	128 processors
	2 partitions	with	16 C-bricks	64 processors
	4 partitions	with	8 C-bricks	32 processors
	1 partition 1 partition	with with	1 C-brick 31 C-bricks	4 processors 124 processors
	1 partition 2 partitions	with with	16 C-bricks 8 C-bricks	64 processors 32 processors
	1 partition 1 partition 2 partitions	with with with	16 C-bricks 8 C-bricks 4 C-bricks	64 processors 32 processors 16 processors
	1 partition 4 partitions	with with	16 C-bricks 4 C-bricks	64 processors 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	1 partition	with	64 C-bricks	256 processors
	2 partitions	with	32 C-bricks	128 processors
	4 partitions	with	16 C-bricks	64 processors
	8 partitions	with	8 C-bricks	32 processors
	16 partitions	with	4 C-bricks	16 processors
	1 partition 1 partition	with with	1 C-brick 63 C-bricks	4 processors 252 processors
	1 partition 2 partitions	with with	32 C-bricks 16 C-bricks	128 processors 64 processors
	1 partition 1 partition 2 partitions	with with with	32 C-bricks 16 C-bricks 8 C-bricks	128 processors 64 processors 32 processors
	1 partition 4 partitions	with with	32 C-bricks 8 C-bricks	128 processors 32 processors
	32 partitions	with	2 C-bricks	8 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
128	1 partition	with	128 C-bricks	512 processors
	2 partitions	with	64 C-bricks	256 processors
	4 partitions	with	32 C-bricks	128 processors
	8 partitions	with	16 C-bricks	64 processors
	16 partitions	with	8 C-bricks	32 processors
	32 partitions	with	4 C-bricks	16 processors
	1 partition 2 partitions	with with	64 C-bricks 32 C-bricks	256 processors 128 processors
	1 partition 1 partition 2 partitions	with with with	64 C-bricks 32 C-bricks 16 C-bricks	256 processors 128 processors 64 processors
	1 partition 1 partition 1 partition 2 partitions	with with with with	64 C-bricks 32 C-bricks 16 C-bricks 8 C-bricks	256 processors 128 processors 64 processors 32 processors
	1 partition 1 partition 1 partition 1 partition 2 partitions	with with with with with	64 C-bricks 32 C-bricks 16 C-bricks 8 C-bricks 4 C-bricks	256 processors 128 processors 64 processors 32 processors 16 processors

Note: 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.

Shaded partitions are either not supported or special order only.

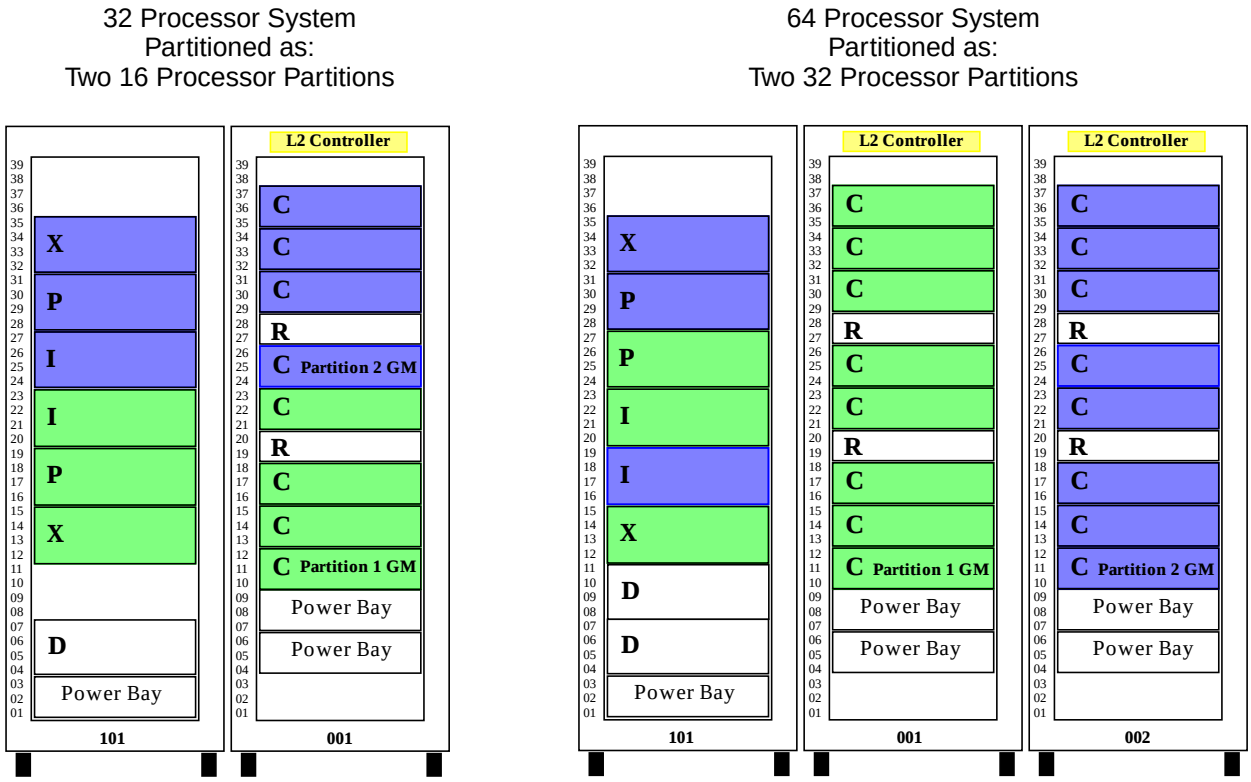

Figure 6-1 SGI Origin 3800 Systems with Even Partitions

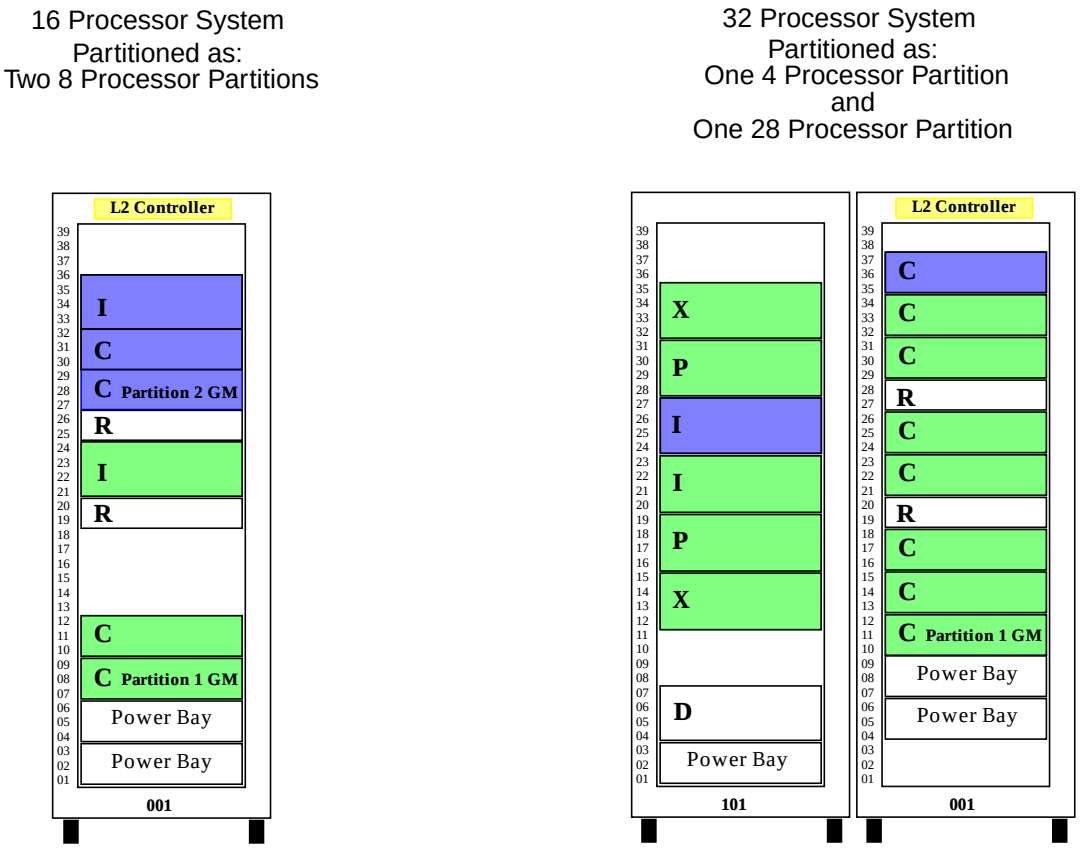

Figure 6-2 SGI Origin 3800 Series Systems with Uneven Partitions

It is important when configuring systems with two uneven partitions that the location of the single C-brick in the single C-brick partition be located in the highest location of the highest rack number. Therefore, a 64 processor system partitioned as two partitions with one partition containing only one C-brick, the location of the single C-brick must be in slot 35 of rack 002.

7. Brick Locations and Xtown2 Cable Configurations

This section defines the standard brick locations within the SGI Origin 3000 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 Origin 3200 Brick Placement Guidelines

1. The SGI Origin 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-, X-, or D-brick that is located at U10
- One power bay that is located at U01

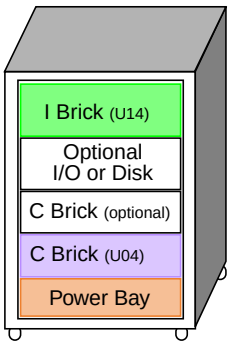

17-U Rack

7.1.2 SGI Origin 3200C System Brick Placement Guidelines

The SGI Origin 3200C system contains:

- From 1 to 3 nodes per tall rack
- Each node consists of:
  - Two (4P) C-bricks
  - One I-brick
- One switch (16-port Myrinet)
- One LAN card per node (Myrinet LAN PCI card)
- From 1 to 2 Power Bays per rack

The first node is located:

- C-brick in U07
- C-brick in U10
- I-brick in U13

The second node is located:

- C-brick in U17
- C-brick in U20
- I-brick in U23

The third node is located:

- C-brick in U27
- C-brick in U30
- I-brick in U33

The switch is located at U37.

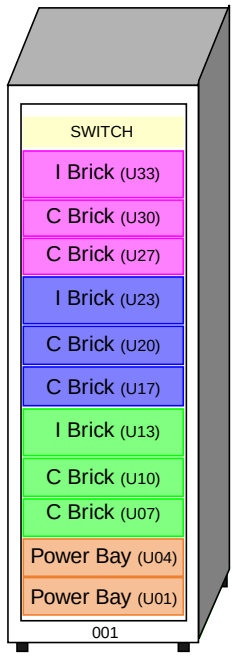

3 Node System (24 Processors)

7.1.3 Single-rack SGI Origin 3400 Brick Placement Guidelines

A single-rack SGI Origin 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	U07, U10, U13, U16, U27
Six	U07, U10, U13, U16, U27, U30
Seven	U07, U10, U13, U16, U27, U30, U33
Eight	U07, U10, U13, U16, U27, U30, U33, U36

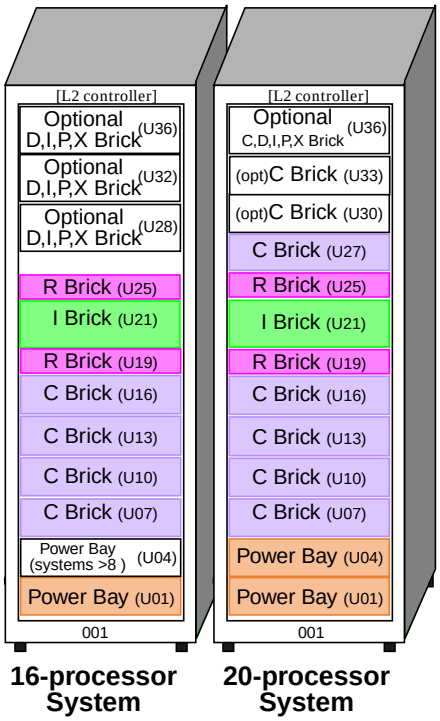

- The first R-brick is located at U19 and the second R-brick is located at U25. Note: The R-brick in location U25 is not used until the number of C-bricks exceeds four.
- The first I-brick is located at U21. Up to three additional I, P, or X-bricks can be placed starting at U36 and working downward. If three additional I/O bricks are required, then there cannot be a utility shelf in the rack.
- D-bricks, if configured, are located in the highest available I/O brick location. Note: D-bricks are slotted after the I/O bricks are slotted.
- 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 Origin 3400 Brick Placement Guidelines

A multiple-rack SGI Origin 3400 system contains multiple combinations of bricks and power bays. The C and R-bricks reside in a compute rack and the I-, P-, X-, and D-bricks reside 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	U10
Two	U10, U13
Three	U10, U13, U16
Four	U10, U13, U16, U21
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

- Two R-bricks; the first R-brick is located at U19 and the second R-brick is located at U27.
- One or two power bays; the first power bay is located at U04. The second power bay is located at U07. Note: Compute racks that have more than eight 48 Vdc power connections require two power bays.

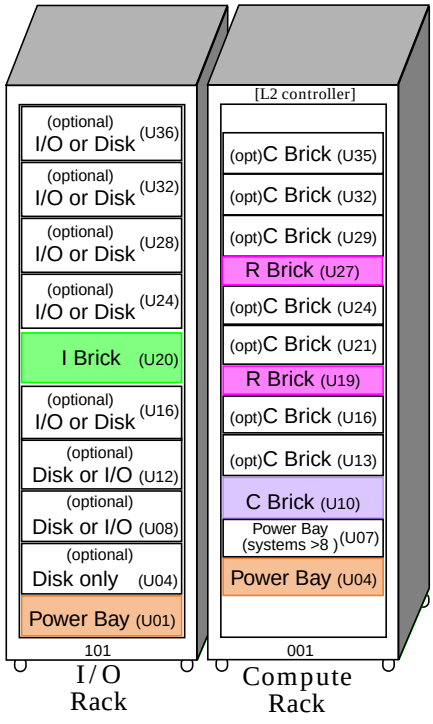

Front View

The I/O rack contains:

- One I-brick that is located at U20 of rack 101.
- Additional I-, P-, or X-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-, P-, or X-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-, P-, and X-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 Origin 3800 Brick Placement Guidelines

The SGI Origin 3800 system contains multiple combinations of bricks and power bays. The C and R-bricks are located in compute racks and the I-, P-, X-, and D-bricks are located in separate I/O racks. In systems that require many D-bricks, the D-bricks can be placed in a 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. You must fill one compute rack before you add additional compute racks. The C-bricks reside the following locations:

Qty C-brick	Location in Compute Rack
One	U10
Two	U10, U13
Three	U10, U13, U16
Four	U10, U13, U16, U21
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

- One or two R-bricks; the first R-brick is located at U19 and the second R-brick is located at U27.
- One or two power bays; the first power bay is located at U04. The second power bay is located at U07.
  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 in the first I/O rack.
- Additional I, P, or X-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, place the additional I-, P-, or X-bricks on 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-, P-, and X-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.

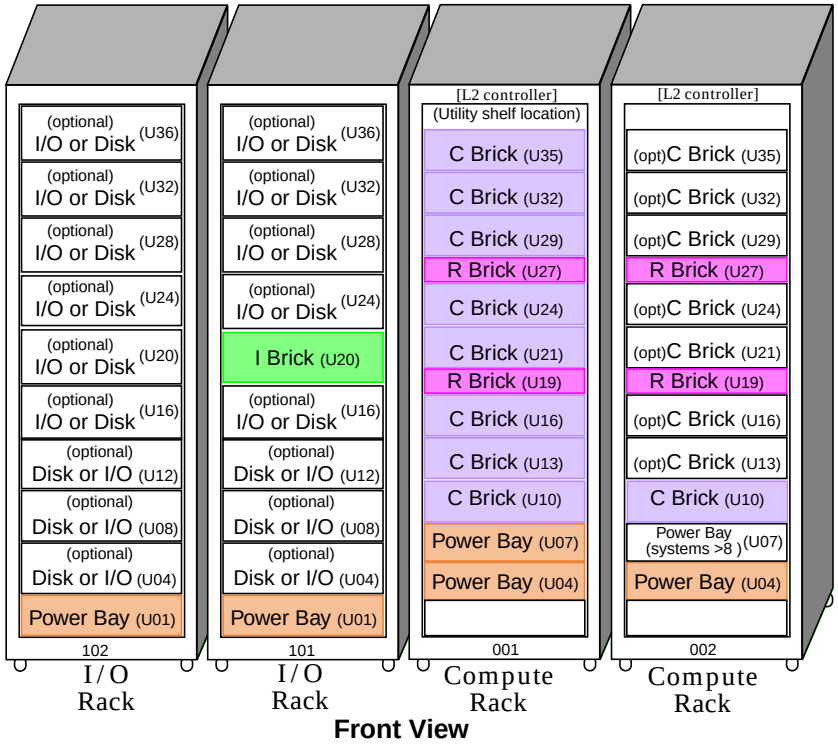

Front View

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 C-bricks-to-I/O bricks within a rack and C bricks-to-I/O bricks between adjacent racks. A 2-meter cable also connects a C brick in any location to an I/O brick within the same rack. When you connect 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.

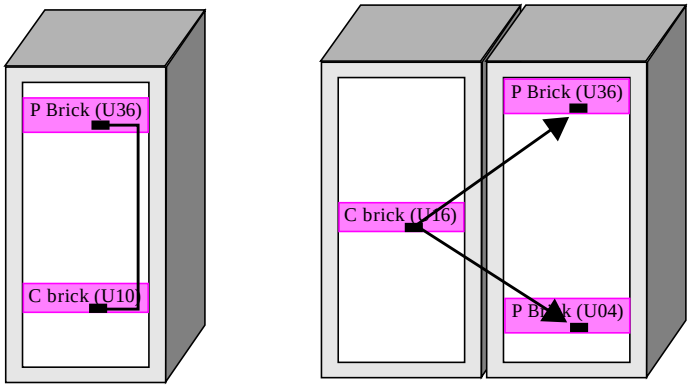

Figure 7-1 NUMAlink Cabling from a Rack and to an Adjacent Rack

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 2U of the C-brick.

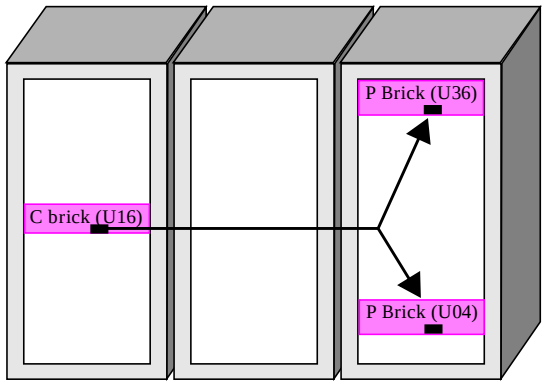

Figure 7-2 NUMAlink Cable Spanning One Full Rack

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.

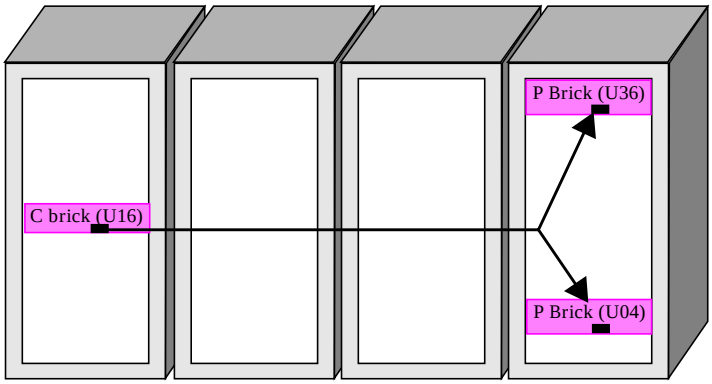

Figure 7-3 NUMAlink Cable Spanning Two Full Racks

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.

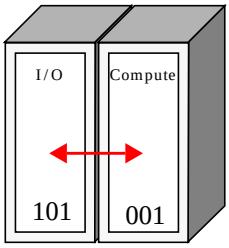

- In systems with two compute racks and two I/O racks:
  - Rack 001 connects to Rack 101; use 2-meter cables
  - Rack 002 connects to Rack 101; use 3-meter cables
  - Rack 002 connects to Rack 102; use 4-meter cables.

  Note: If the I/O brick is within 2 U of the same location as the C-brick, then a 3-meter cable can be used. For example, if the C-brick is at location U13 and the I/O brick is at location U12, a 3-meter cable can be used.

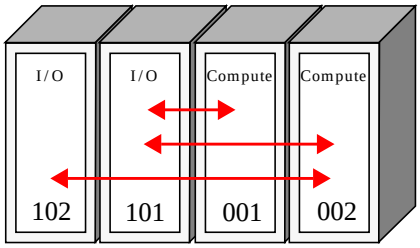

- In systems with three compute racks and three I/O racks:
  - Rack 001 connects only to Rack 101; use 2-meter cables
  - Rack 002 connects to Rack 101; use 3-meter cables
  - Rack 002 connects to Rack 102; use 4-meter cables
  - Rack 003 connects to Rack 111; use 2-meter cable.

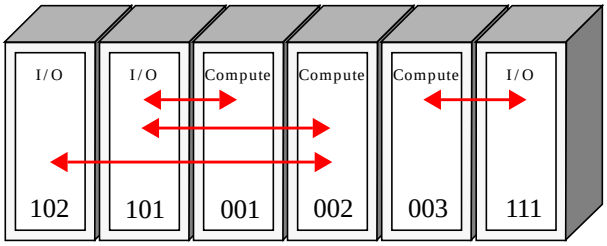

- In systems with four compute racks and four I/O racks:
  - Rack 001 connects only to Rack 101; use 2-meter cables
  - Rack 002 connects to Rack 101; use 3-meter cables
  - Rack 002 connects to Rack 102; use 4-meter cables
  - Rack 003 connects to Rack 111; use 3-meter cables
  - Rack 003 connects to Rack 112; use 4-meter cables
  - Rack 004 connects to Rack 111; use 2-meter cables

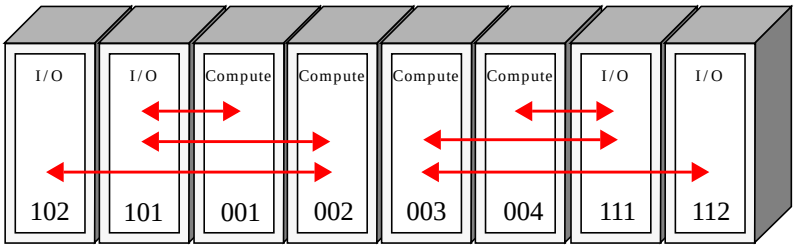

8. Configuration Guidelines

This section provides guidelines for configuring SGI Origin 3000 series systems. These systems have a PCI-based I/O subsystem for the primary I/O and a Crosstalk I/O subsystem that is compatible with XIO boards from Origin 2000 and Octane systems.

8.1 System Configuration Process

To configure a system within the SGI Origin 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 Origin 3200 systems are configured with 2, 4, 6, or 8 processors.
  System sizes are increased in 2-or 4-processor increments.
  Note: the 4-processor configuration is offered as:
  One 4 processor C-brick or two C-bricks, each with two processors (2×2).
- SGI Origin 3200C systems consists of 8-processor nodes. The systems can scaled to hundreds and thousands of processors:
  upgrades are in one node (8 processor) increments.
- SGI Origin 3400 systems can be configured with 4, 8, 16, or 32 processors;
  upgrades are in 4-processor increments.
- SGI Origin 3800 systems are configured with 16 to 512 processors;
  upgrades are as follows.
  - If the existing system has less than 128 processors;
    the system sizes increase in 4-processor increments up to 128 processors.
  - If the existing system is equal or greater than 128 processors;
    the system sizes increases in 32-processor increments.

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 (standard or 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.

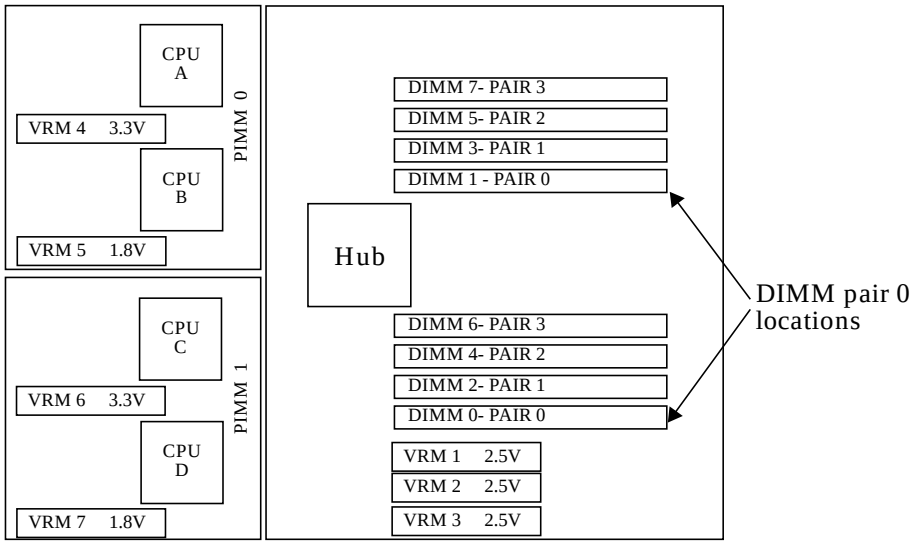

Figure 8-1 Memory DIMM and Processor PIMM locations

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-D (plus) One MEM-1G	One MEM-2G-D (plus) Two MEM-512	
4.0 Gbytes	1.0 GBs	Two MEM-2G-D	Four MEM-1G	
6.0 Gbytes	1.5 GBs	Three MEM-2G-D	Two MEM-2G-D (plus) Two MEM-1G	
8.0 Gbytes	2.0 GBs	Four MEM-2G-D		
CUSTOM MEMORY SIZES				
2.5 Gbytes	625 MBs	One MEM-2G-D (plus) One MEM-512	Two MEM-1G (plus) One MEM-512	
3.5 Gbytes	875 MBs	One MEM-2G-D (plus) Three MEM-512	One MEM-2G-D (plus) One MEM-1G (plus) One MEM-512	
4.5 Gbytes	1.125 GBs	Two MEM-2G-D (plus) One MEM-512		
5.0 Gbytes	1.25 GBs	Two MEM-2G-D (plus) One MEM-1G	Two MEM-2G-D (plus) Two MEM-512	
7.0 Gbytes	1.75 GBs	Three MEM-2G-D-D (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 Origin 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.

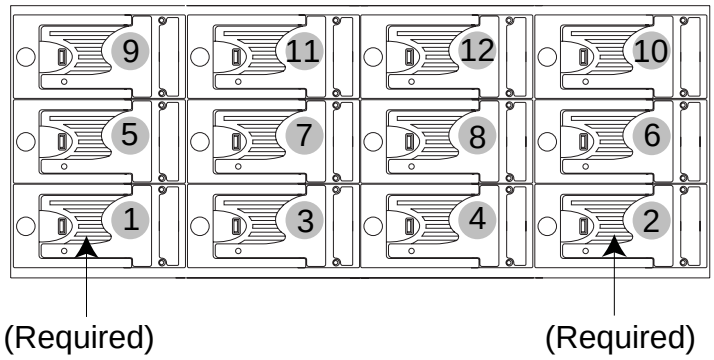

Figure 8-2 Disk Drive Installation Order

8.4 Guidelines for Determining I/O Brick Locations

Due to the flexibility that has been designed into the SGI Origin 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 and XIO 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 Origin 3400 System.

1. Install the I-brick in slot U21.
2. Install remaining I/O bricks starting in slot 36 and working downward in the following order: X-bricks first, then P-bricks, then I-bricks, then D-bricks

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:

3. Install the P-bricks. The P-bricks are alternately installed above and below the I-bricks.
4. Install the X-bricks. The X-bricks are installed alternately above and below the P-bricks.

Note: The X-bricks are installed last. They are either the highest or the lowest position brick of the I-, P-, or X-bricks. The X-brick has the smallest quantity of cards; therefore, it is positioned at the lowest or highest positions of the rack.

---

4. Install D-bricks in the lower section of the rack starting at U04. Any remaining D-bricks that cannot be slotted in the lower section are placed in the upper section of the rack starting at the first available slot above the I/O bricks.

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 or XIO 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 four 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 87.

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 87

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.

8.5.4 Configuring the X-brick

The following are guidelines for configuring XIO cards in the X-brick:

- Distribute XIO cards evenly among the X-bricks.
- Distribute XIO cards evenly among the slots of each X-brick.

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

- Fully compatible with XIO slots in SGI Origin 2000 and Octane systems.
- Supports XIO card power levels up to 50 watts per card.
- The formula to determine the number of required X bricks is:

  #XB = TotalXIO / 4

  Where:

  #XB is the number of X-bricks required.

  TotalXIO is the total number of XIO cards required.

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

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:

TPB = 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:

#DPS = 1 + (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 J for brick power requirements.

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

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

The power bay connects to the bricks with either a 1-meter or a 2-meter power cord. The power bay located in U04 of a rack uses a 1-meter power cord to connect to bricks in locations U04 through U14. Bricks in locations above U14 require a 2-meter power cord.

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

The power bay located in U07 of a rack uses 2-meter power cords for connecting to bricks in location U24 and above.

8.6.1 Installing Distributed Power Supplies (DPS)

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

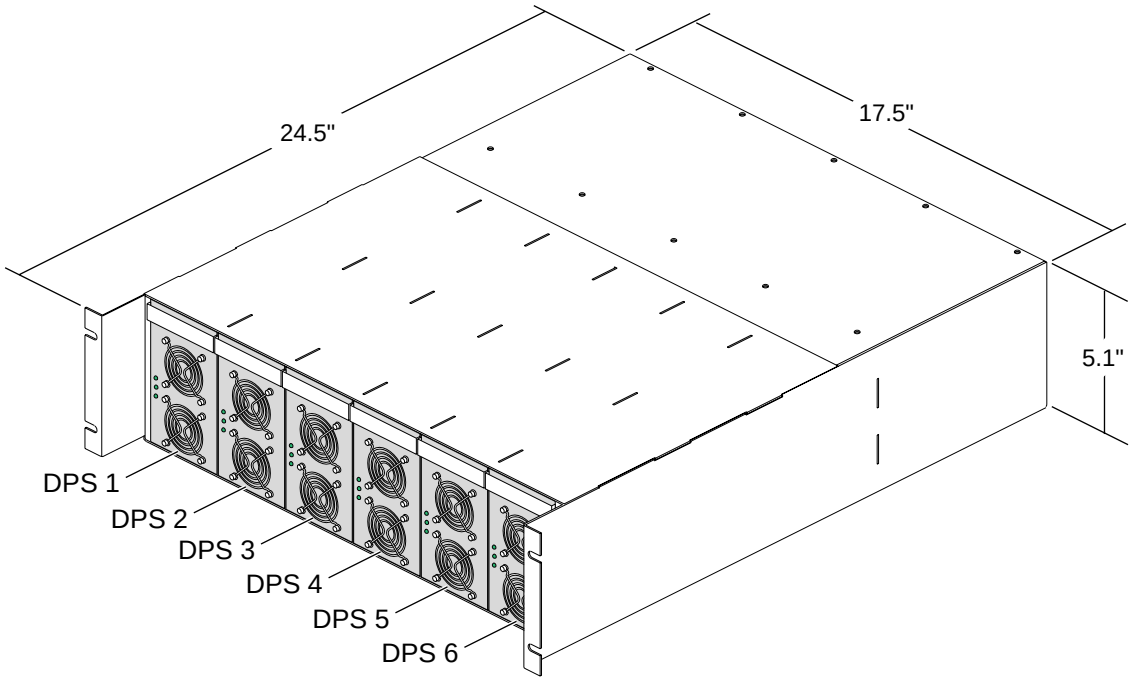

Figure 8-3 DPS Power Bay Locations

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 or through the R-brick's NUMAlink connection to 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 the C-brick's 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 Origin 3200 System

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

2 Processor System

- One short rack with trim
- One power bay with three power supplies
- One single-phase power distribution strip
- One 2-processor C-brick without memory
- One I-brick with Fibre Channel PCI card and cable
- One 18-GB system disk
- One 1-meter NUMAlink cable
- Two 1-meter power cables
- System software (IRIX) - SC4-ASE-6.5

4 Processor System

- One short rack with trim
- One power bay with three power supplies
- One single-phase power distribution strip
- One 4-processor C-brick without memory
- One I-brick with Fibre Channel PCI card and cable
- One 18-GB system disk
- One 1-meter NUMAlink cable
- Two 1-meter power cables
- System software (IRIX) - SC4-ASE-6.5

4 Processor System

- One short rack with trim
- One power bay with three power supplies
- One single-phase power distribution strip
- Two 2-processor C-bricks without memory
- One I-brick with Fibre Channel PCI card and cable
- One 18-GB system disk
- Two 1-meter NUMAlink cables
- Three 1-meter power cables
- System software (IRIX) - SC4-ASE-6.5

8 Processor System

- One short rack with trim
- One power bay with three power supplies
- One single-phase power distribution strip
- Two 4-processor C-bricks without memory
- One I-brick with Fibre Channel PCI card and cable
- One 18-GB system disk
- Two 1-meter NUMAlink cables
- Three 1-meter power cables
- System software (IRIX) - SC4-ASE-6.5

Each of the four SGI Origin 3200 system bundles require that you select from the 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-, P-, or X-brick
- Customer configurable XIO cards, if an X-brick is configured
- 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 Origin 3200C System

The SGI Origin 3200C system is scalable from 16 processors up.

Single Rack 24 Processor System

- One tall rack with trim
- Two power bays: one with five power supplies, one with three power supplies
- One single-phase or three-phase PDU
- Three 8-processor nodes without memory - each node contains:
  - Two 4-processor C-bricks
  - One I-bricks with Fibre Channel PCI card and cable
  - One 18-GB system disk
- One 16-port Myrinet LAN switch
- Three Myrinet LAN PCI cards
- Three 10ft Myrinet LAN cables
- One Utility Shelf
- Two 1-meter NUMAlink cables
- Three 1-meter power cords
- Six 2-meter power cords
- Three System software (IRIX) - SC4-ASE-6.5
- One SC4-ACEIRIX-1.0-4N

8.8.3 SGI Origin 3400 System

Four basic marketing codes define the SGI Origin 3400 system.

4 Processor System

- One tall rack with trim
- One power bay with four power supplies
- One single-phase or three-phase PDU
- One 4-processor C-brick without memory
- Two R-bricks
- One I-brick with Fibre Channel PCI card and cable
- One 18-GB system disk
- One L2 controller
- Three 1-meter NUMAlink cables
- One 2-meter NUMAlink cable
- Three 1-meter power cords
- One 2-meter power cord
- System software (IRIX) - SC4-ASE-6.5

8 Processor System

- One tall rack with trim
- One power bay with four power supplies
- One single-phase or three-phase PDU
- Two 4-processor C-bricks without memory
- Two R-bricks
- One I-brick with Fibre Channel PCI card and cable
- One 18-GB system disk
- One L2 controller
- Four 1-meter NUMAlink cables
- One 2-meter NUMAlink cable
- Three 1-meter power cords
- Two 2-meter power cord
- System software (IRIX) - SC4-ASE-6.5

16 Processor System

- One tall rack with trim
- One power bay with four power supplies
- One single-phase or three-phase PDU
- Four 4-processor C-bricks without memory
- Two R-bricks
- One I-brick with Fibre Channel PCI card and cable
- One 18-GB system disk
- One L2 controller
- Six 1-meter NUMAlink cables
- One 2-meter NUMAlink cable
- Three 1-meter power cords
- Four 2-meter power cord
- System software (IRIX) - SC4-ASE-6.5

32 Processor System

- One tall rack with trim
- Two power bays with eight power supplies
- Two single-phase PDUs or one three-phase PDU
- Eight 4-processor C-bricks without memory
- Two R-bricks
- One I-brick with Fibre Channel PCI card and cable
- One 18-GB system disk
- One L2 controller
- Ten 1-meter NUMAlink cables
- One 2-meter NUMAlink cable
- Three 1-meter power cords
- Eight 2-meter power cords
- System software (IRIX) - SC4-ASE-6.5

Each of the four SGI Origin 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-, P-, or X-bricks
- Customer configurable XIO cards, if an X-brick is configured
- 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

8.8.4 SGI Origin 3800 Systems

The SGI Origin 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, X, and D-bricks (I/O rack). The SGI Origin 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:

16 Processor System

- One tall rack with trim
- One power bay with four power supplies
- One single-phase or three-phase PDU
- Four 4-processor C-bricks without memory
- Two R-bricks
- One L2 controller
- One Ethernet cable (L2)
- Six 1-meter NUMAlink cables
- Three 1-meter power cables
- FIve 2-meter power cables
- System software (IRIX) - SC4-ASE-6.5

32 Processor System

- One tall rack with trim
- Two power bays with eight power supplies
- Two single-phase PDUs or one three-phase PDU
- Eight 4-processor C-bricks without memory
- Two R-bricks
- One L2 controller
- Ten 1-meter NUMAlink cables
- Three 1-meter power cables
- Seven 2-meter power cables
- System software (IRIX) - SC4-ASE-6.5

The base I/O rack has the following components:

- One tall rack with trim
- One power bay with five power supplies
- One single-phase or three-phase PDU
- One I-brick with Fibre Channel PCI card and cable
- One 18-GB system disk
- One P-brick
- Two 2-meter NUMAlink cable
- Two 2-meter power cables

To complete the SGI Origin 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-, P-, or X-brick
- Customer configurable XIO cards, if an X-brick is configured
- 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 Origin 3000 Series

There are three areas in which the SGI Origin 3000 series 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 Origin 3200 systems can have C-bricks that contain two or four processors. A C-bricks that contains two processors can be upgraded by adding an additional PIMM (a PIMM contains two processors). A system with a single four-processor C-brick can be upgraded by adding an additional four-processor C-brick.

SGI Origin 3400 systems are configured with four-processor C-bricks. SGI Origin 3400 systems can be expanded to a maximum of 8 four-processor C-bricks.

SGI Origin 3800 systems are configured with four-processor C-bricks. SGI Origin 3800 systems can be expanded to a maximum of 128 four-processor C-bricks. Each compute rack can hold a maximum of eight C-bricks; therefore, additional racks must be configured when the number of C-bricks exceeds 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 Gbyte, 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 four memory options.:

Table 9-1 Memory Marketing Codes

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

If a customer requires more memory, but the memory in the existing C-bricks is filled; additional memory can be added by increasing the number of C-bricks up to the system maximum (refer to “Processor Capacity” on page 137).

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.

9.4 PCI Card Expansion

When adding PCI cards to an existing system refer to the SGI Origin 3000 Series Configuration Rules document located at:

http://wwwcf.americas.sgi.com/PUBLIC/tech_pub/

Section 5 of the configuration rules document describes the steps used to configure PCI cards in I-bricks and P-bricks.

Once you have determined the slots in the I-brick or P-brick where you are going to install the new PCI cards, the next step is to determine the process for installing the new cards. The IRIX level, the type of PCI card, and the cards device driver determines wether or not you will be able to hot plug the new PCI cards.

Use the following guidelines and Table 9-2 to determine if you will be able to hot plug PCI cards in your specific SGI Origin 3000 system. The following guidelines use the term “Hot Insertion”. The definition of “Hot Insertion” is the ability to install PCI cards in a PCI slot while the system is operating, but requires taking administrative actions to power the targeted PCI slot down prior to inserting the PCI card. The “Hot Removal” definition is the same as “Hot Insertion” only removing the PCI card from the system.

SUPPORTED ACTIONS

1. Hot insertion/removal of a 33 MHz card in the 33 MHz slots (bus 1) of the I-brick is dependent on the ability of the device driver and the IRIX level.

2. Hot removal of a 33 MHz card from a 66 MHz capable bus running at 33 MHz when the other slot on the bus has a PCI card installed is supported.

3. Hot insertion/removal of 66 MHz cards is dependent on the ability of the device driver and the IRIX level.

UNSUPPORTED ACTIONS

1. Hot insertion of a 33 MHz card in a 66 MHz bus is not supported. This includes both inserting a 33 MHz card in an empty bus that is defaulted to 66 MHz and installing a 33 MHz card in a 66 MHz bus where the bus already has a 66 MHz card attached.

2. Hot removal of a 33 MHz card from a 66 MHz capable bus running at 33 MHz when the other slot on the bus is empty is not supported.

Table 9-2 PCI Card Hot Swap Support Schedule

Marketing Code	Hot Insertion Supported	Hot Removal Supported
PCI-FC-1PCOP-A	IRIX 6.5.11	IRIX 6.5.13
PCI-FC-1POPT-A	IRIX 6.5.11	IRIX 6.5.13
PCI-ATMOC3-1P	TBD	TBD
PCI-ATMOC12-1P	TBD	TBD
PCI-GIGENET-C	TBD	TBD
PCI-GIGENET-OR	TBD	TBD
PCI-SCSI-DF-2P	IRIX 6.5.11	IRIX 6.5.13
PCI=SCSI-LVD-2P	IRIX 6.5.11	IRIX 6.5.13
PCI-AUD-D1000	IRIX 6.5.12	TBD
PCI-SER-10002	IRIX 6.5.12	TBD

10. SGI Origin 3000 Series Software Requirements

10.1 Required Software

SGI Origin 3000 series systems use the IRIX 6.5 Advanced Server Environment (ASE) software . The marketing code for this package is SC4-ASE-6.5. The following software is included in ASE:

- IRIX 6.5 operating system
- License for NFS (from Sun)
  NFS is a file-mounting / sharing tool.
- License for Database Accelerator (DBA from SGI)
  DBA provides optimized database scalability and enhanced manageability of decision support systems and data warehouse systems.
- License for host access to HP (referred to as HP MIB)
- License for CA Unicenter TNG Framework
  Basis for Computer Associates large-enterprise system management tool for heterogeneous environments.
- Temporary license for EnlightenDSM (from Enlighten Software Solutions, Inc.)
  Distributed system management for small enterprises and work groups
- Temporary license for Roboinst (from SGI)
  Roboinst provides remote software distribution for heterogeneous networked environments.
- CD pack for Hot Mix (SGI)
- CD pack for SGI freeware

10.2 Optional Software

To access the online price book for a list of optional software go to the following URL:

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 Origin 3200 System Configuration

System Definition:

- Single 17-U high 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 (1) 2-P PIMM Add (1) 4-P C-brick Add (1) 2-P PIMM 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 (1) 2-P PIMM Add (2) 2-P PIMMs
(qty 1) 4-P C-brick (qty 1) 2-P C-brick	Add (1) 2-P PIMM
(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)
- One L3 controller (optional)

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

SGI Origin 3200 Systems

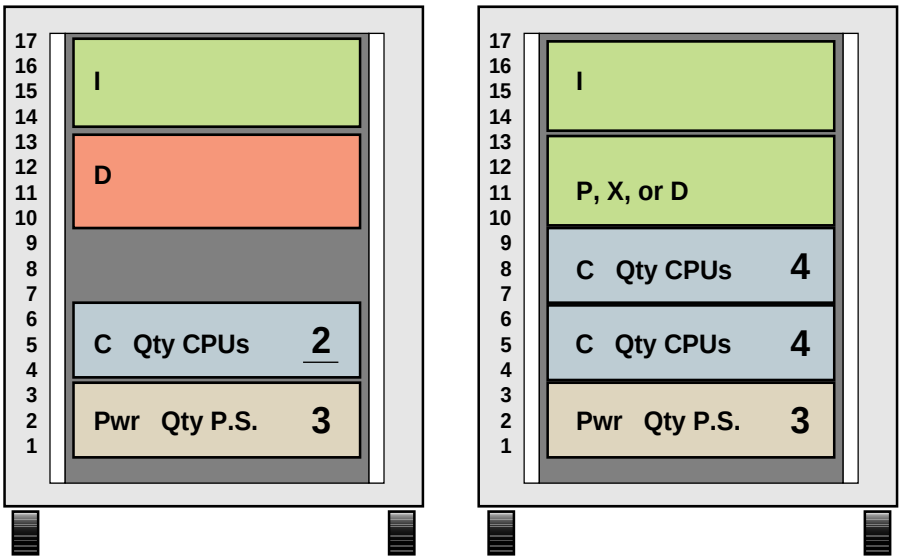

Note: Disks can be placed in a separate SGI or third-party disk rack.

SGI Origin 3200C System Configuration

System Definition:

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

Main Memory:

- Minimum node memory size 1 GB: (two C-bricks, each with one 512-MB bank)
- Maximum node memory size 16 GB: (two C-bricks, each with four 2-GB banks)

Processors:

- Minimum of two C-bricks (8 processors per node)
- Maximum of two C-bricks (8 processors per node)

I/O:

- One additional I/O brick can be added per node

Nodes:

- Minimum of one node per rack, Maximum of three nodes per rack
- Minimum of two nodes per switch, Maximum of 16 nodes per switch
- Minimum of two nodes per system

Disk:

- Additional D-brick rack(s) for disks as required

Miscellaneous:

- Minimum of 1 power bay, maximum of 2 power bays per rack
- L3 controller (optional)

SGI Origin 3200C System

3 Node
24-processor System

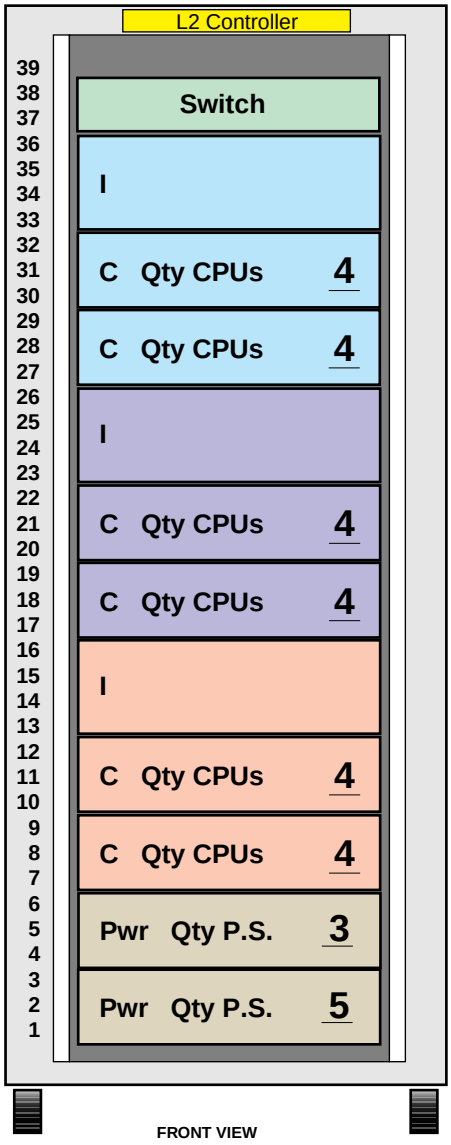

Note:
Additional I/O or D-brick racks not shown.

SGI Origin 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 two R-bricks	Add one through seven 4-processor C-bricks
Two 4-processor C-bricks and two R-bricks	Add one through six 4-processor C-bricks
Four 4-processor C-bricks and two R-bricks	Add one through four 4-processor C-bricks
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
- L3 controller (optional)

SGI Origin 3400 Systems

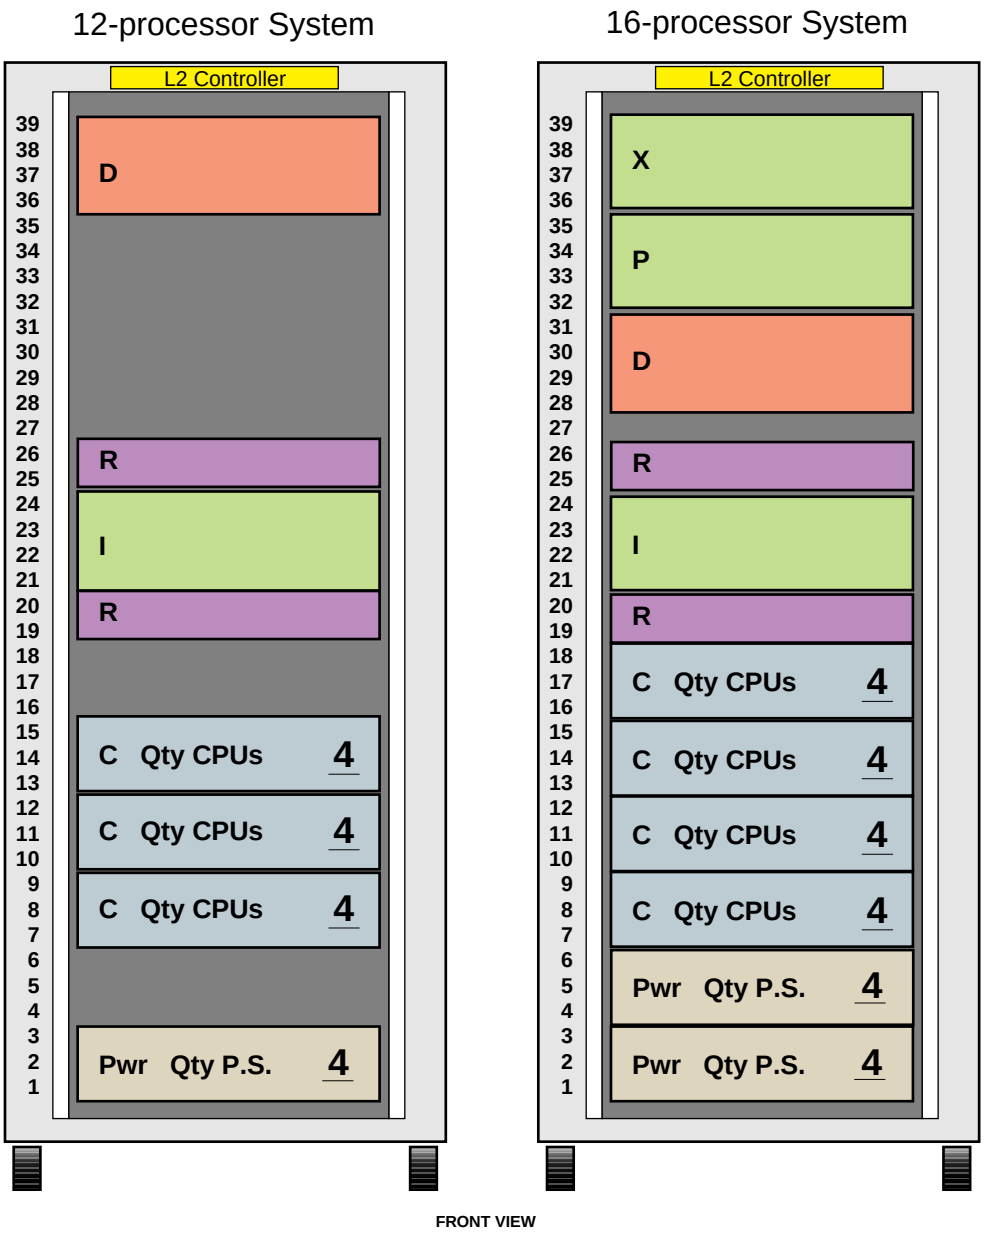

Note:
Additional D-brick racks not shown

32-processor System

SGI Origin 3400 System

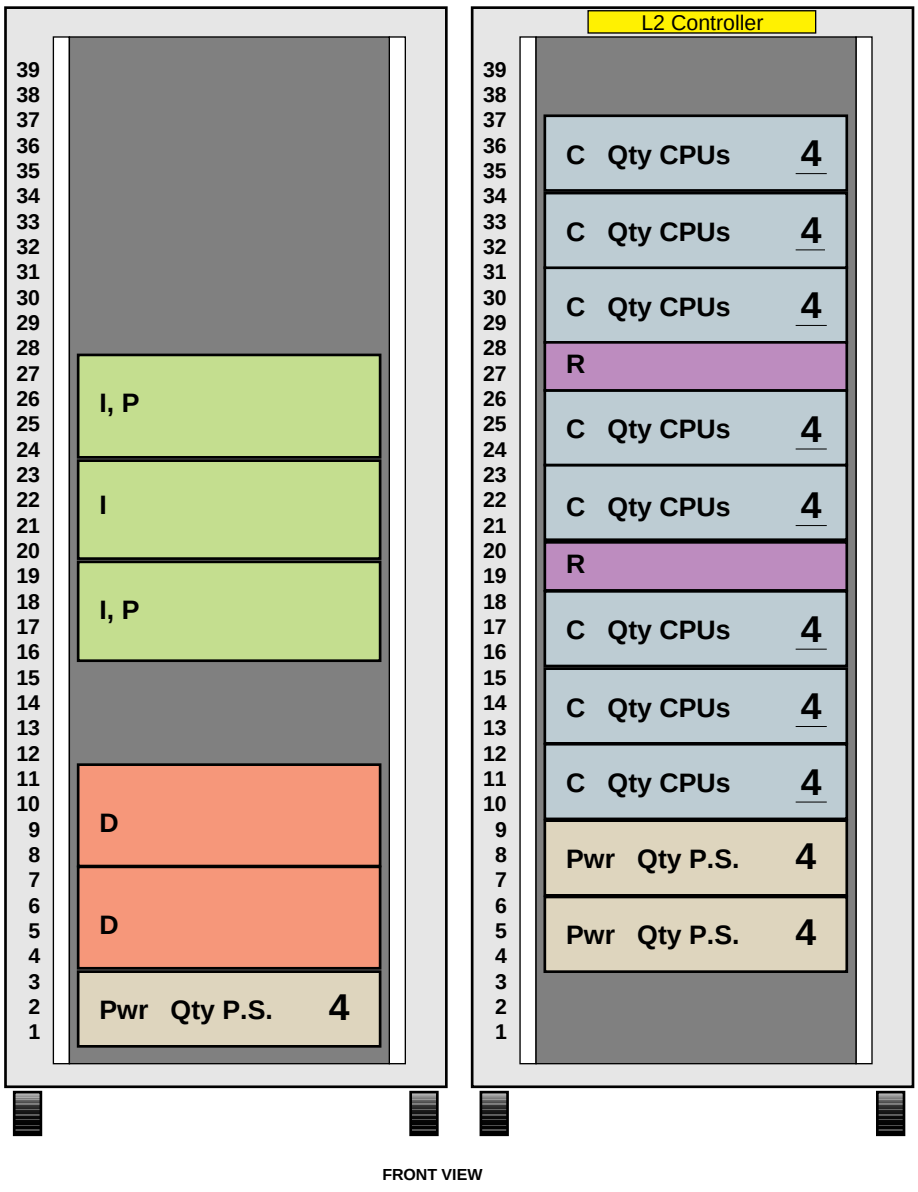

Note:
Additional racks with D-bricks not shown

SGI Origin 3800 System Configuration

System definition:

Main Memory:

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

Processors:

- 16 to 512 processors (increased in increments of 4, 128, and 256 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 three 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 mandatory for all compute racks
- L3 controller is optional

16-processor (Minimum System)

SGI Origin 3800 System

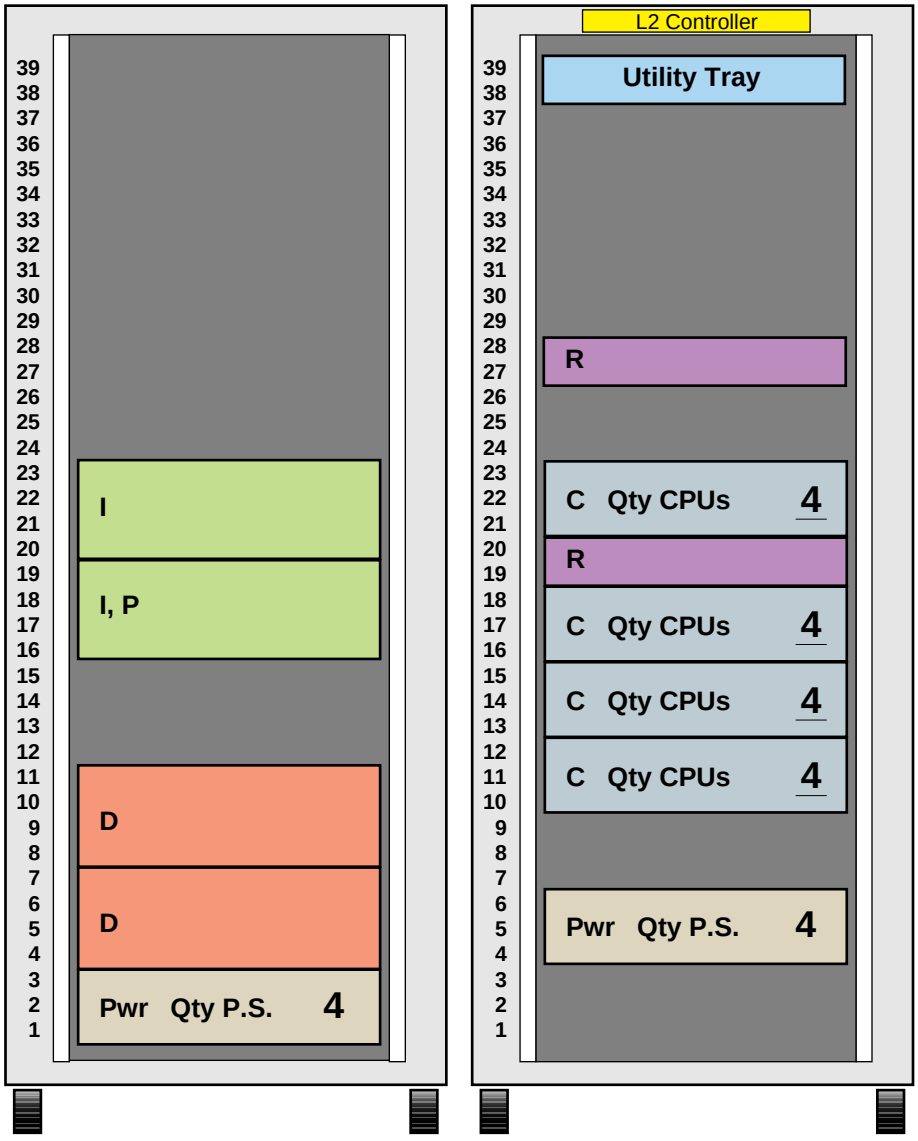

SGI 3800 System

32-processor System

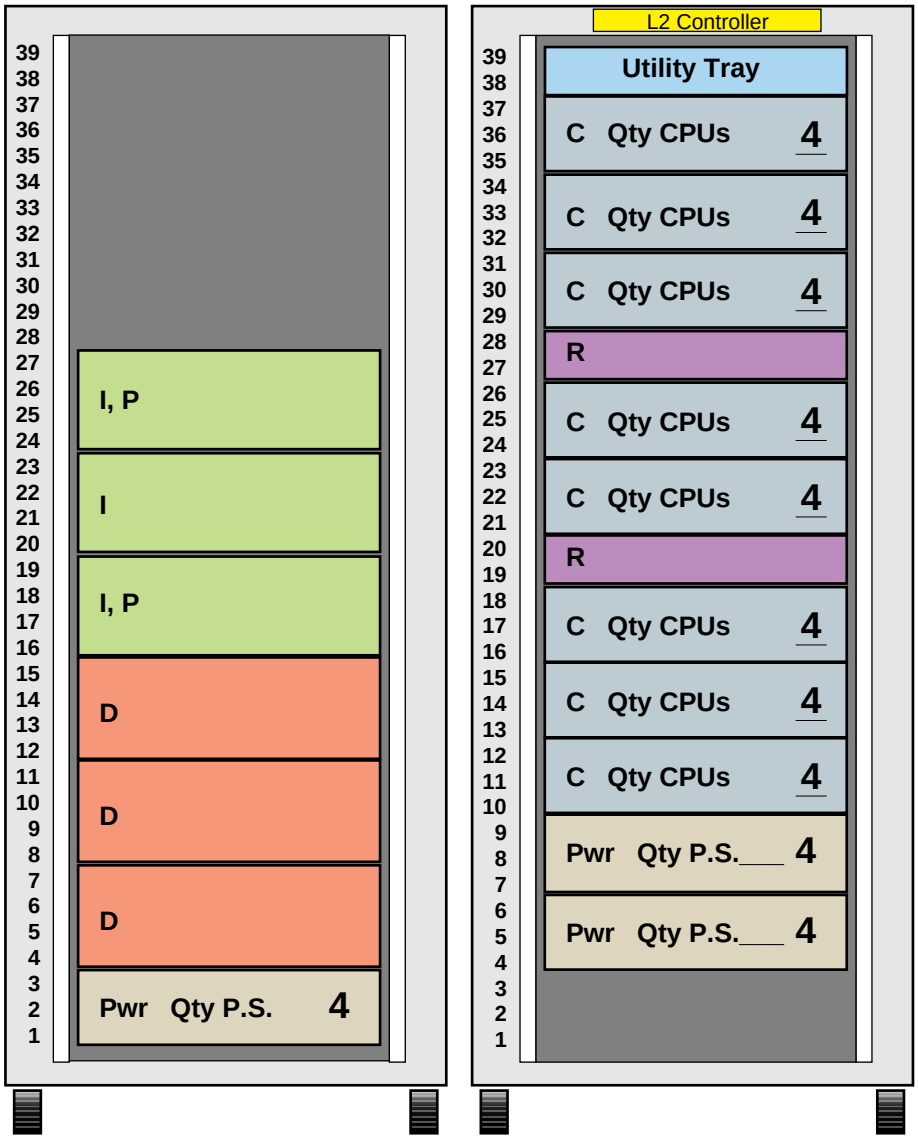

SGI Origin 3800 System
128-processor system

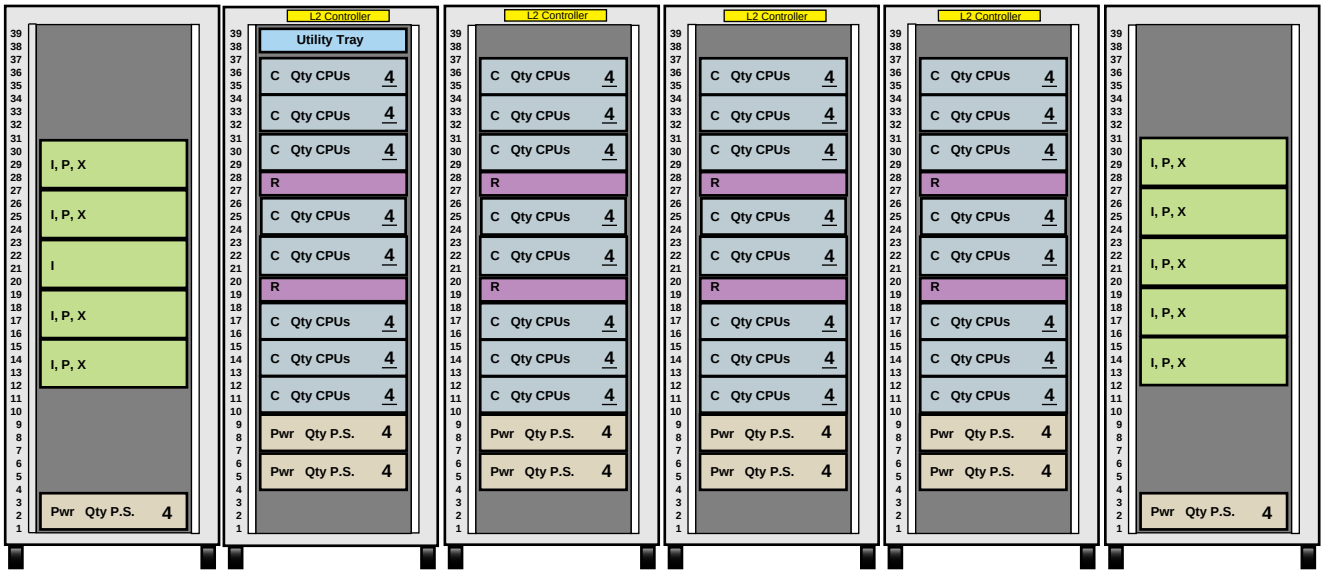

Note: Racks that contain the D bricks are not shown.

SGI Origin 3800 System
One quadrant of 512-processor System

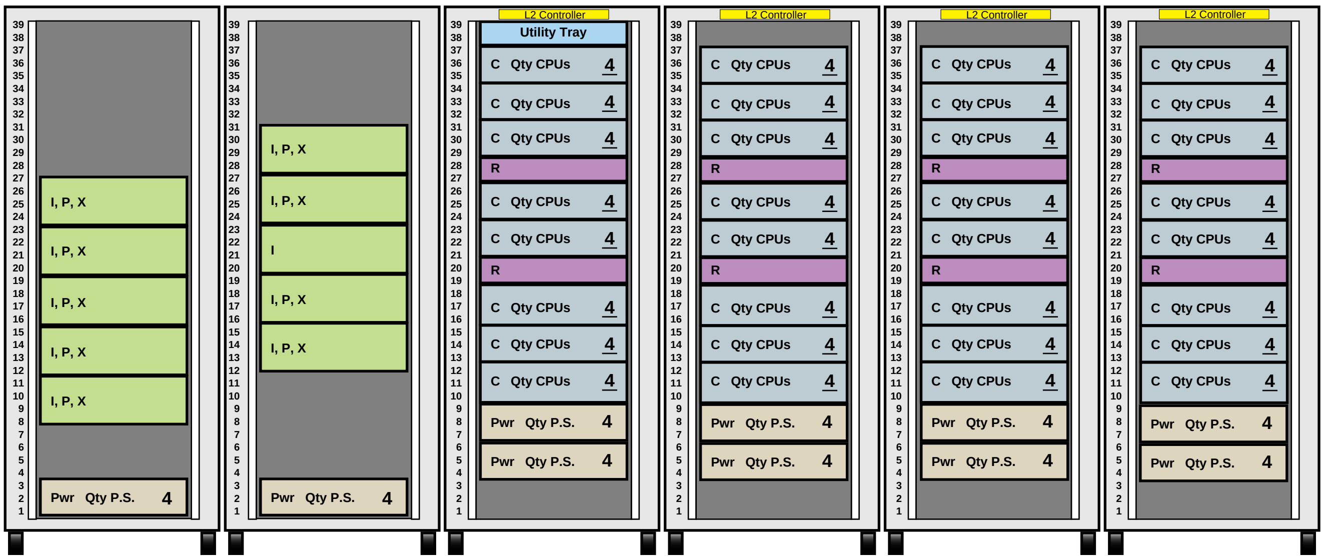

Note: Racks that contain the D bricks are not shown.

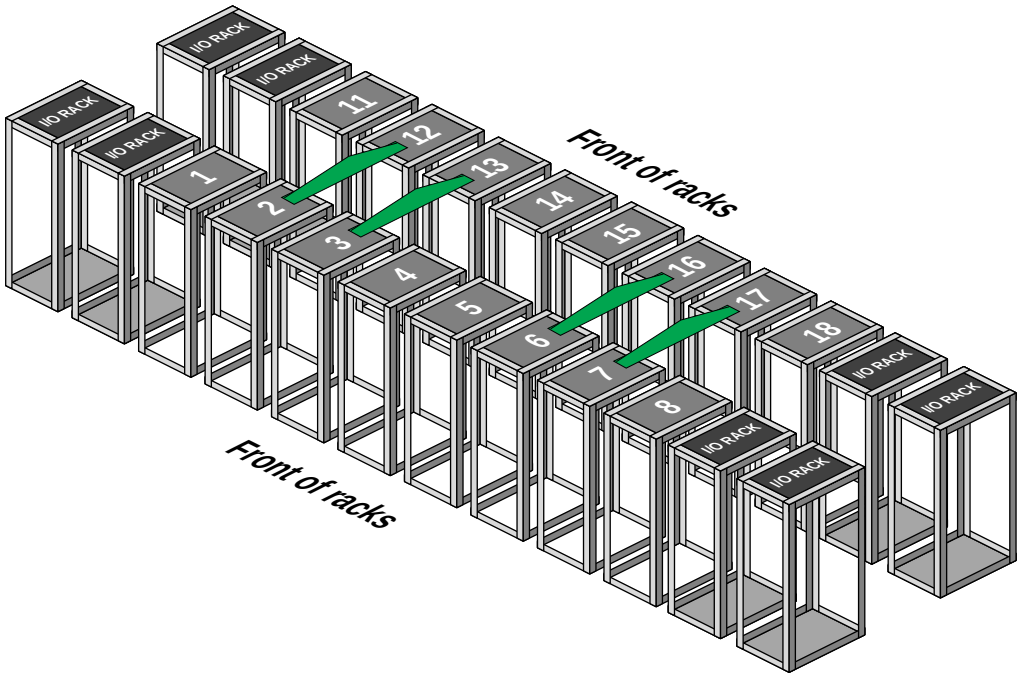

512-processor SGI Origin 3800 System

D-brick Rack Configuration

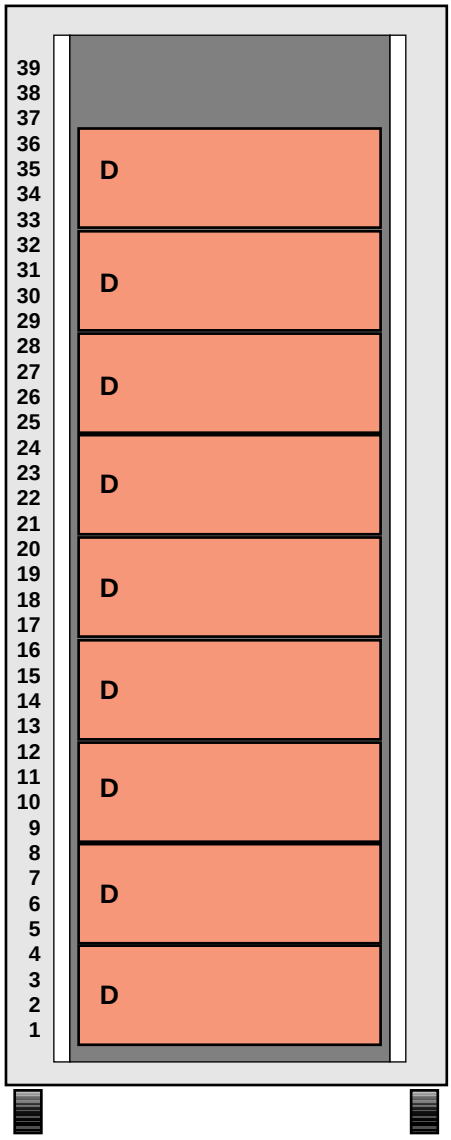

Mechanical and Electrical Specifications

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. (914 mm) (48.0 in. recommended) 36.0 in. (914 mm) None
Short Rack Electrical Requirements	
Voltage Nominal Tolerance	North America / Japan / International 200-240 Vac / 230 Vac 180-254 Vac / 180-254 Vac
Frequency Nominal Tolerance	North America / Japan / International 50/60 Hz / 50 Hz 47-63 Hz / 47-63 Hz
Phases	Single-Phase
Power Requirements (maximum) Short Rack	 1.18 kVA (1.16 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)	4.78 Kbtu/hr (Based on 1.40 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 Nominal Tolerance	North America / Japan / International 200-240 Vac / 230 Vac 180-254 Vac / 180-254 Vac
Frequency Nominal Tolerance	North America / Japan / International 50/60 Hz / 50 Hz 47-63 Hz / 47-63 Hz
Phases	Single-Phase or Optional Three-Phase
Power Requirements (maximum) Compute Rack I/O Rack Disk Rack	 2.63 kVA (2.58 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)	10.24 Kbtu/hr (Based on 3.00 kW)	
Cooling Requirement	Ambient air	
Airflow: (Intake, Front; Exhaust, Rear)	Less than 3200 CFM	

System / Channel Bandwidths

Description	Clock Frequency	Sustained Bandwidth	Peak Bandwidth
Main Memory Bandwidth	200 MHz	3200 MB/s	3200 MB/s
NUMAlink 3 Channel Bandwidth	800 MHz	1420 MB/s each direction	3.2 GB/s (full-duplex) 1.6 GB/s each direction
Xtown2 Channel Bandwidth	600 MHz	~ 1066 MB/s half duplex ~ 1744 MB/s full-duplex ~ 872 MB/s each direction	2.4 GB/s (full-duplex) 1.2 GB/s each direction
Xtown2 Channel Bandwidth	400 MHz	~700 MB/s half duplex 1120 MBs/ full-duplex ~560 MB/s each direction	1.6 GB/s (full-duplex) 800 MB/s each direction
PCI Channel Bandwidth	33 MHz	N/A	128 MB/s in 32-bit mode 256 MB/s in 64-bit mode
	66 MHz	N/A	256 MB/s in 32-bit mode 512 MB/s in 64-bit mode

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
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)
PIMM	Processor integrated multi-module
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
XIO	SGI proprietary I/O channel, rated at 800 MB/s full-duplex

Rear View of Bricks

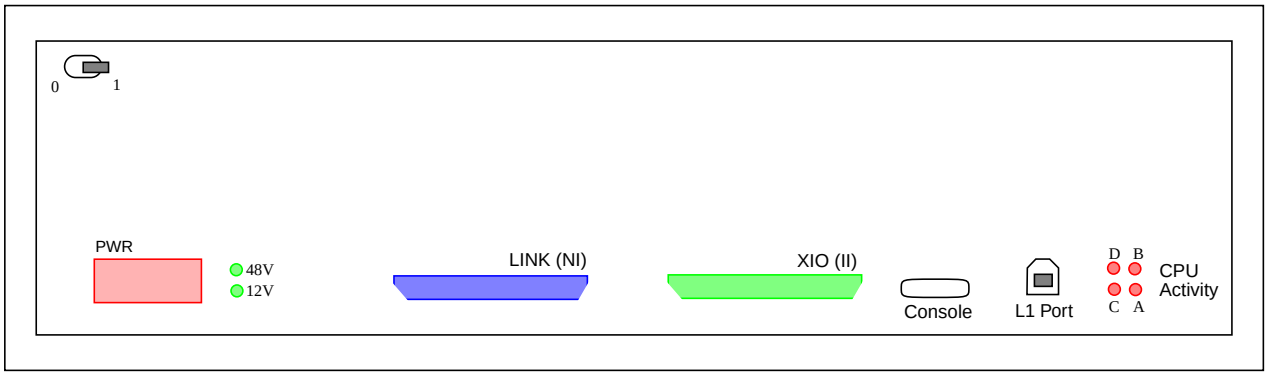

C Brick
(Rear View)

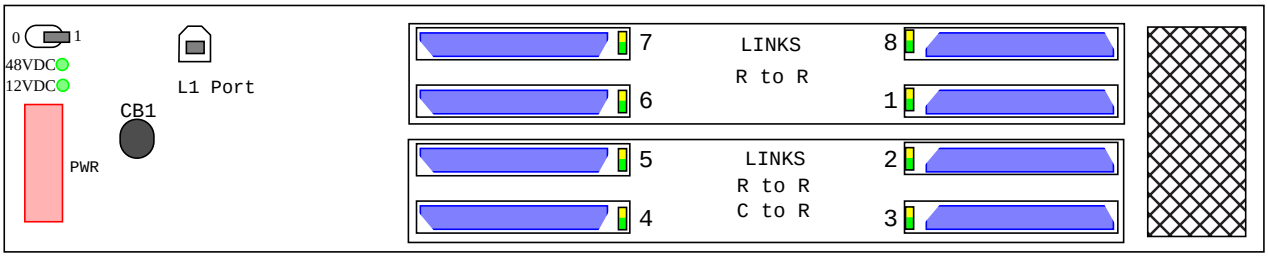

R Brick
(Rear View)

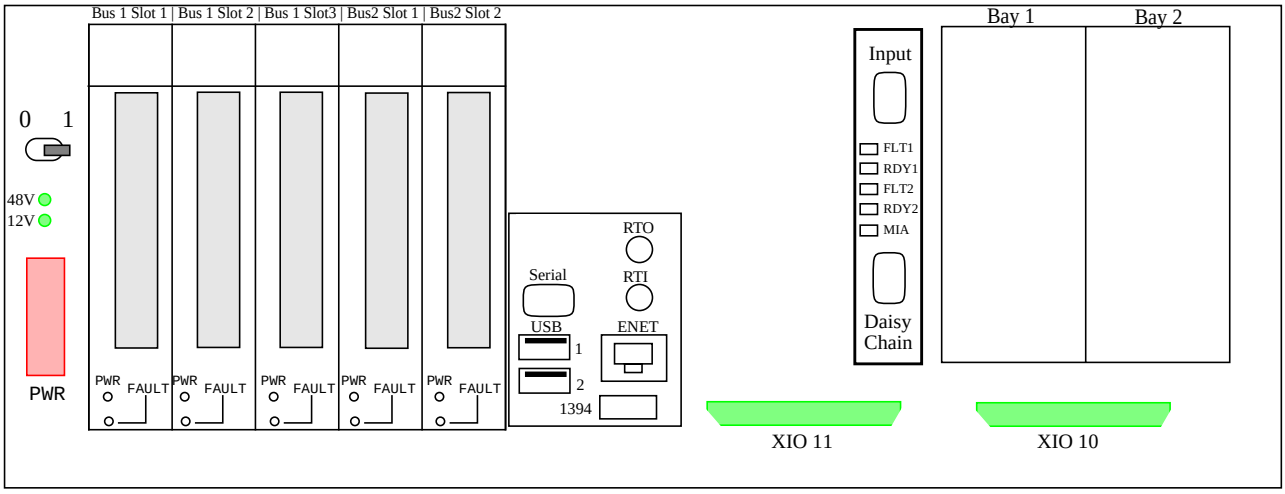

I Brick
(Rear View)

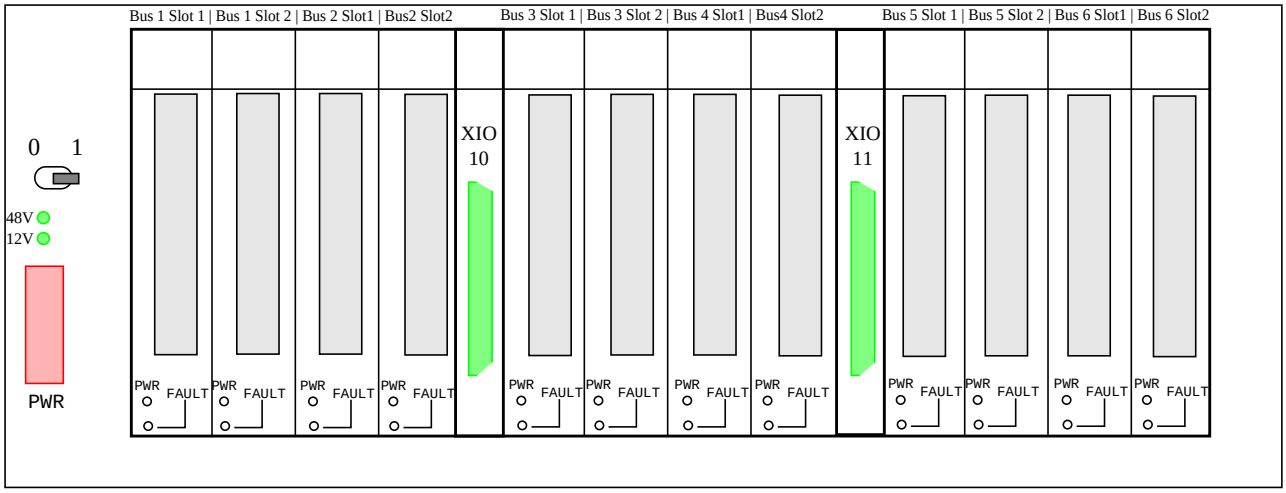

P Brick
(Rear View)

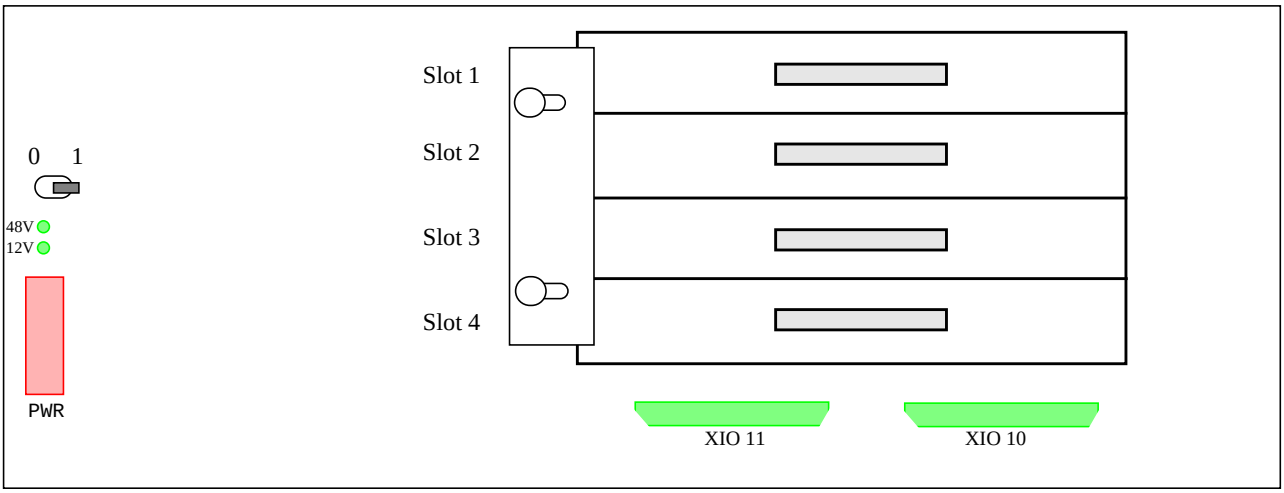

X-brick
(Rear View)

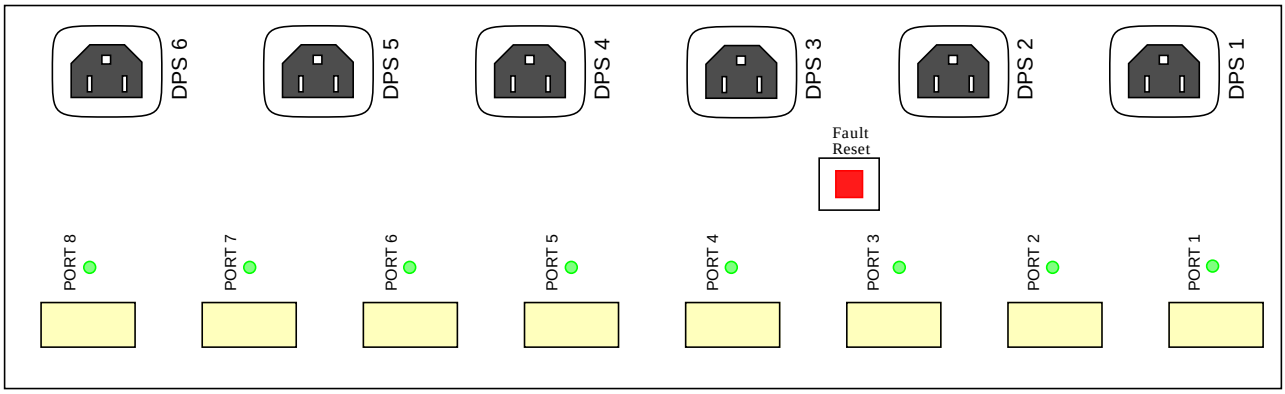

Power Bay
(Rear View)

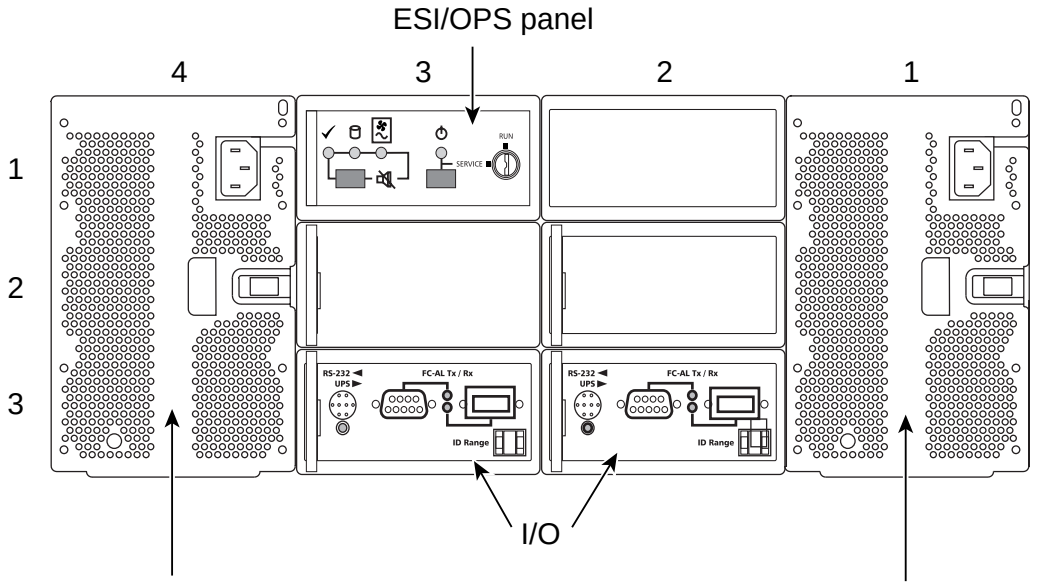

D-brick
(Rear View)

Component Dimensions, Weights, and Power

Product Name	Dimensions	Weight	Power	Number 0f U
Tall Rack	74 in. H × 32 in. W × 48 in. D	400 lbs (empty)	N/A	N/A
Short Rack	34 in. H × 24 in. W × 42 in. D	200 lbs (empty)	N/A	N/A
C-brick	5.06 in. H × 17.19 in. W × 27.80 in. D	65 lbs (fully loaded)	~308 watts	3
D-brick	6.64 in. H × 17.50 in. W × 27.74 in. D	94 lbs (fully loaded)	~400 VA	4
G-brick	31.5 in. H × 19.0 in. × 20.0 in. D	215 lbs (fully loaded)	2000 watts	18
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
X-brick	6.64 in. H × 17.50 in. W × 27.74 in. D	69 lbs (fully loaded)	~225 watts	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

System Site Requirements

System	Number of Processors	Number of Compute Racks	Number of I/O Racks	System Footprint (sq ft)	Power Required (kW)	Cooling Requirement (kBtu)
Origin 3200	8	Compute and IO in one rack		7.16	1.16	3.96
Origin 3400	16	Compute and IO in one rack		10.73	2.22	7.57
Origin 3400	32	Compute and IO in one rack		10.73	2.77	9.45
Origin 3400	32	1	1	21.28	4.72	16.10
Origin 3800	16	1	1	21.28	2.22	7.57
Origin 3800	32	1	1	21.28	4.72	16.10
Origin 3800	64	2	1	31.83	7.30	24.91
Origin 3800	128	4	1	52.93	12.46	42.51
Origin 3800	256	8	2	149.80	24.92	85.03
Origin 3800	512	16	4	299.09	49.84	170.05

Index

N

O

P

R

S

T

U

V

W

X