Product Overview
SGI™ Origin™ 3000 and
SGI™ SNIA 3000 Server Series

SGI Confidential & Proprietary Information - For Internal Recipients Only

Document Number 108-xxx-xxx

CONTRIBUTORS

Written by Nancy Heller
Illustrated by Kelly Begley and Nancy Heller
Edited by Allison Gosbin
Production by Rhonda Kunsman

Record of Revision

Version	Description
001	July 2000 Original printing.
002	December 2000 This revision includes the following technical changes: • Noted that the P brick no longer supports the G brick • Updated Figures 1-3 through 1-7 • Reduced the number of power supplies in the I/O rack power bay from 5 to 4
00x	?? 2001 This revision adds the SGI SNIA 3000 series server information.

Contents

Figures

Tables

Chapter 1

Product Overview

This document provides overview information for the SGI Origin 3000 series servers and the SGI SNIA 3000 series servers. For information that applies to both the SGI Origin 3000 series servers and the SGI SNIA 3000 series servers, the name 3000 series servers is used throughout this document.

The 3000 series servers are distributed shared memory (DSM) systems where each compute node (2 or 4 processors) has local memory that it shares with the other compute nodes in the system.

The 3000 series servers are based on SGI NUMA (formerly ccNUMA), a cache-coherent non-uniform memory access architecture, which ensures that the caches of the processors contain valid data. For example, if a processor alters the data in a cache location and another processor has a copy of that data in its cache, the processor that holds the copy will be notified that the memory location no longer contains valid data. SGI NUMA architecture also supports varied access times for local and remote memory references.

All 3000 series servers are air cooled by fans and receive power from power bays. The power supplies and cooling fans have the following RAS (reliability, availability, and serviceability) features:

- The power supplies and cooling fans are N+1 redundant, which means that the system contains additional fans and power supplies so that if a power supply or fan fails, the system runs without interruption.
- The power supplies and cooling fans can be hot swapped, which means that a fan or power supply can be physically removed from the system without causing an interruption or taking any administrative action.

Some other RAS features of the 3000 series servers include:

- Single-bit error correction and double-bit error detection (SECDED) for memory and secondary cache
- Parity protection for a processor's primary cache
- Power-on self tests

For a complete list of the RAS features, refer to the *System Architecture* document, publication number 108-0240-00x.

1.1 SGI Origin 3000 Series Server

The SGI Origin 3000 series servers:

- Scale from 2 to 512 processors
- Support the MIPS R12000, MIPS R12000A, MIPS R14000 processors and their successors
- Use a peripheral component interface (PCI) based I/O system as the primary I/O system; a secondary I/O system supports the legacy Crosstalk I/O (XIO) system of the SGI Origin 2000 and Silicon Graphics Octane systems
- Use the IRIX operating system that includes the baseline features of the IRIX 6.5 release

The SGI Origin 3000 series servers consist of peripheral devices, I/O interfaces, routers, and compute nodes that are linked by a NUMAlink 3 interconnect network. These system components are packaged in standard 19-inch rack-mounted subassemblies that are referred to as bricks.

An SGI Origin 3000 series server can consist of the following bricks:

- C brick - Houses the compute node that provides the compute functionality and memory for the system.

  The key components of the compute node are (refer to Figure 1-1):

  - 2 or 4 MIPS processors (R12000, R12000A, R14000, or their successors)

    Note: Each processor has access to two 32-Kbyte primary caches and a 4- or 8-Mbyte secondary cache. The size of the secondary cache depends on the processor type.

  - Memory - includes a directory for cache coherence
  - Hub - allows communication among the node components, I/O interfaces, and the NUMAlink 3 interconnect network

    Note: The hub is also referred to as Bedrock.

- R brick - Houses a router that routes information between C bricks, either directly or through other R bricks.

  Note: Routers make up the NUMAlink 3 interconnect network.

- I brick - An I/O interface that provides the base I/O functions for the SGI Origin 3000 series servers, which includes 1 or 2 system disks for boot functions. The I brick also houses 5 PCI cards.
- P brick - An I/O interface that houses 12 PCI cards.
- X brick - An I/O interface that houses four half-height XIO slots. These slots are compatible with the XIO slots of the Origin 2000 and Octane systems.
- D brick - A peripheral device that supports two Fibre Channel loops and houses 12 disk drives.
- G brick - Houses the board set of the SGI Onyx 3000 graphics system.

  Note: The RAS features of the power supplies and fans that are mentioned on page 1-1 do not apply to the G brick.

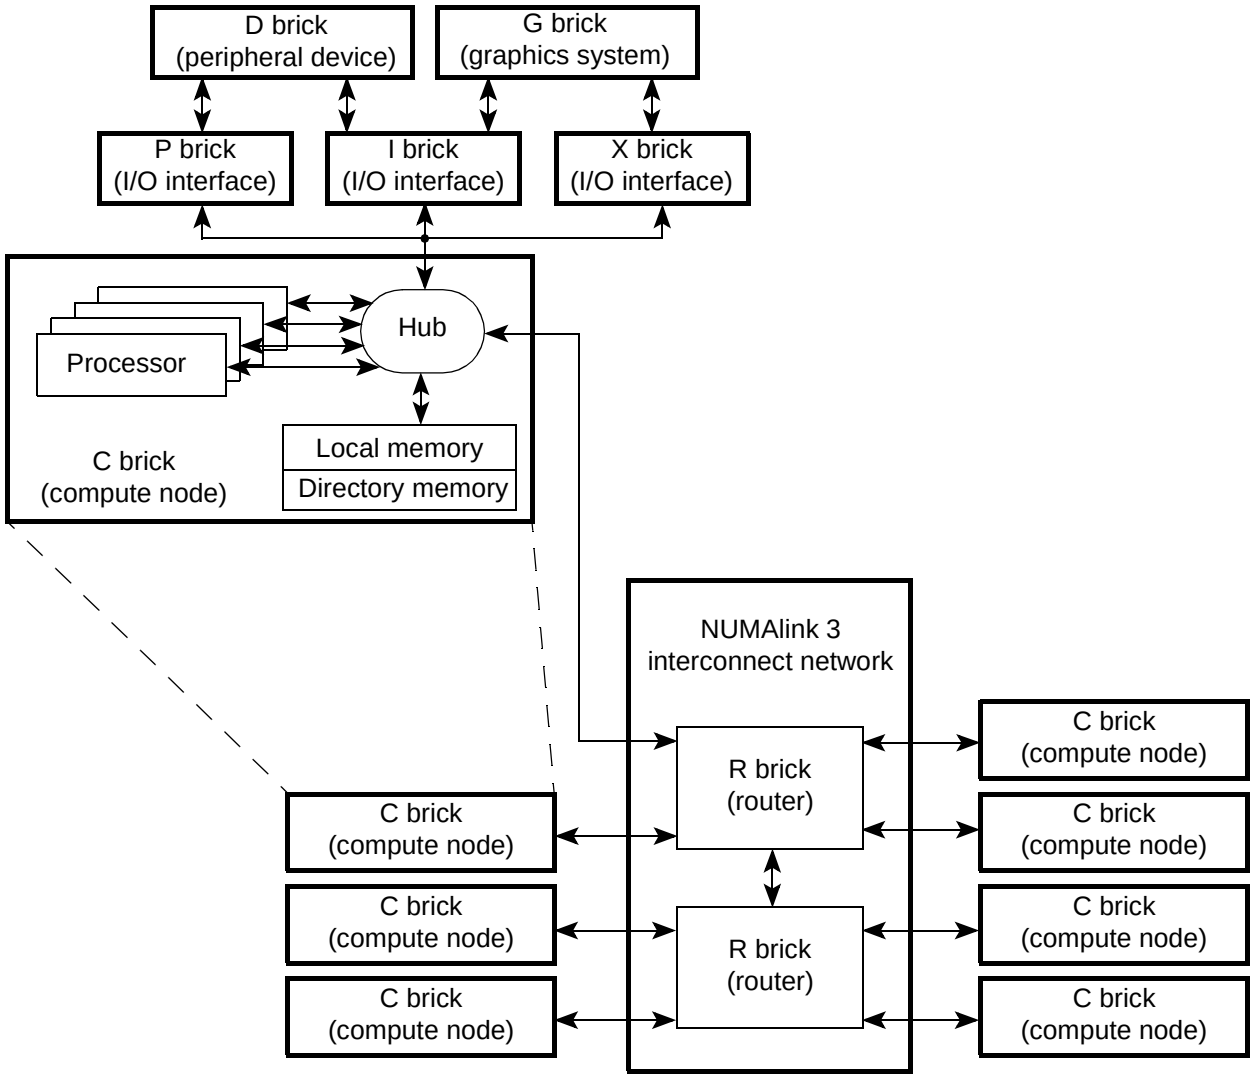

Figure 1-1 System Block Diagram

1.1.1 SGI Origin 3000 Series Server Configurations

The SGI Origin 3000 series servers can be expanded by increasing the number of bricks or memory size. Table 1-1 lists the minimum and maximum system configurations.

Table 1-1 Minimum and Maximum System Configurations

System Component	Minimum	Maximum
Processors	2	512
Memory	512 Mbytes (1 C brick)	1024 Gbytes (128 C bricks)
I/O channels	1	128 (1 I/O channel per C brick)

The SGI Origin 3000 series product line includes three system classifications:

- The SGI Origin 3200 server consists of 2 to 8 processors.
- The SGI Origin 3400 server consists of 4 to 32 processors.
- The SGI Origin 3800 server consists of 16 to 512 processors.

1.1.1.1 SGI Origin 3200 Server

An SGI Origin 3200 server consists of 2 to 8 processors and uses one 17-unit (U) rack. (One U is 1.75 in.) The 17-U rack contains (refer to Figure 1-2):

- 1 or 2 C bricks
  - Minimum amount of memory - 512 Mbytes (1 C brick that has 2 banks of memory)
  - Maximum amount of memory - 16 Gbytes (2 C bricks; each C brick has 8 banks of memory)
- 1 or 2 I/O bricks (At least one of the I/O bricks must be an I brick; the second I/O brick can be an I, P, or X brick)
- 1 D brick (optional)
- 1 power bay that contains 3 power supplies

Note: The SGI Origin 3200 server does not use an R brick.

The SGI Origin 3200 server might also include:

- Additional D bricks in separate racks (maximum of 9 D bricks per rack)
- A rack (L2) controller or system (L3) controller (both are optional)

An SGI Origin 3200 server is upgraded by individual bricks or by increments of 2 processors. For example, if the SGI Origin 3200 server has 1 C brick that contains 2 processors, the system can be upgraded to a 4-processor system by adding 2 processors to the C brick. An SGI Origin 3200 system that has one 4-processor C brick can be upgraded to an 8-processor system by adding another 4-processor C brick.

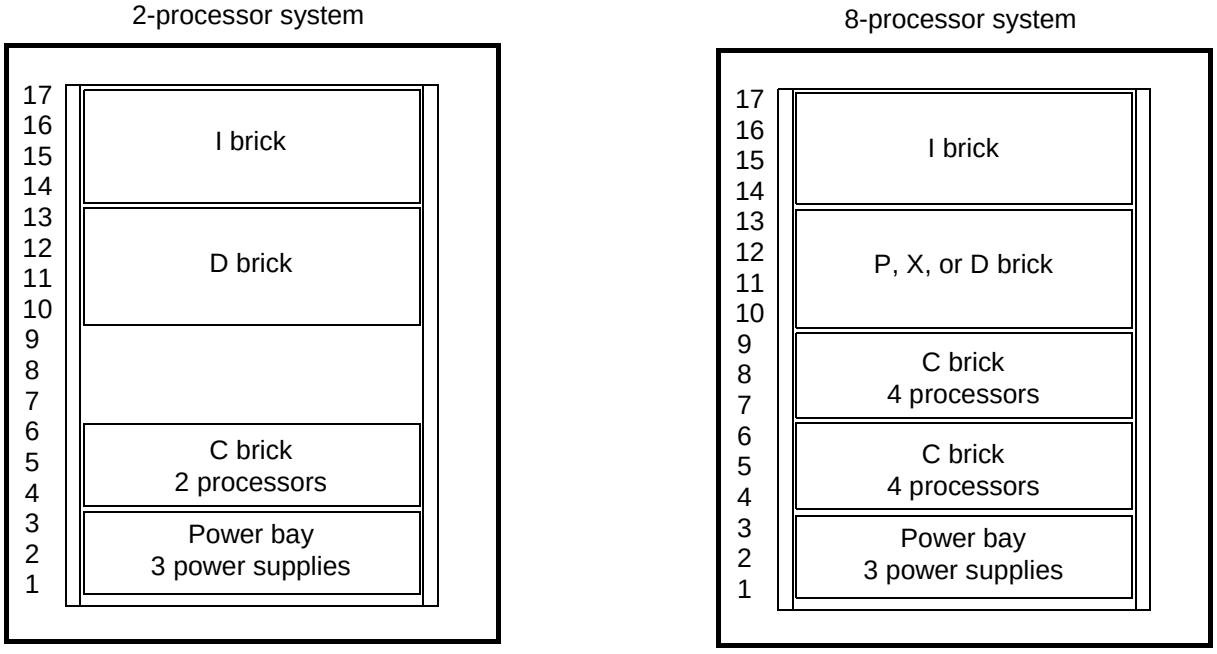

Figure 1-2 SGI Origin 3200 Servers

1.1.1.2 SGI Origin 3400 Server

An SGI Origin 3400 server consists of 4 to 32 processors (refer to Figure 1-3) and uses one 39-U compute rack that contains:

- 1 to 8 C bricks
  - Minimum amount of memory - 512 Mbytes (1 C brick that has 2 banks of memory)
  - Maximum amount of memory - 64 Gbytes (8 C bricks; each C brick has 8 banks of memory)
- 1 I brick
- 2 R bricks
- 1 or 2 power bays; each power bay contains 4 power supplies
- 1 L2 controller

The SGI Origin 3400 server might also include:

- 1 to 3 D bricks in the compute rack
- Additional D bricks in separate racks (maximum of 9 D bricks per rack)
- Additional I/O bricks in the compute rack
- Additional I/O bricks in separate (I/O) racks
- An L3 controller (optional)

An SGI Origin 3400 server is upgraded by individual bricks and by increments of 4 processors. For example, to upgrade a 16-processor system to a 20-processor system, a 4-processor C brick is added to the system.

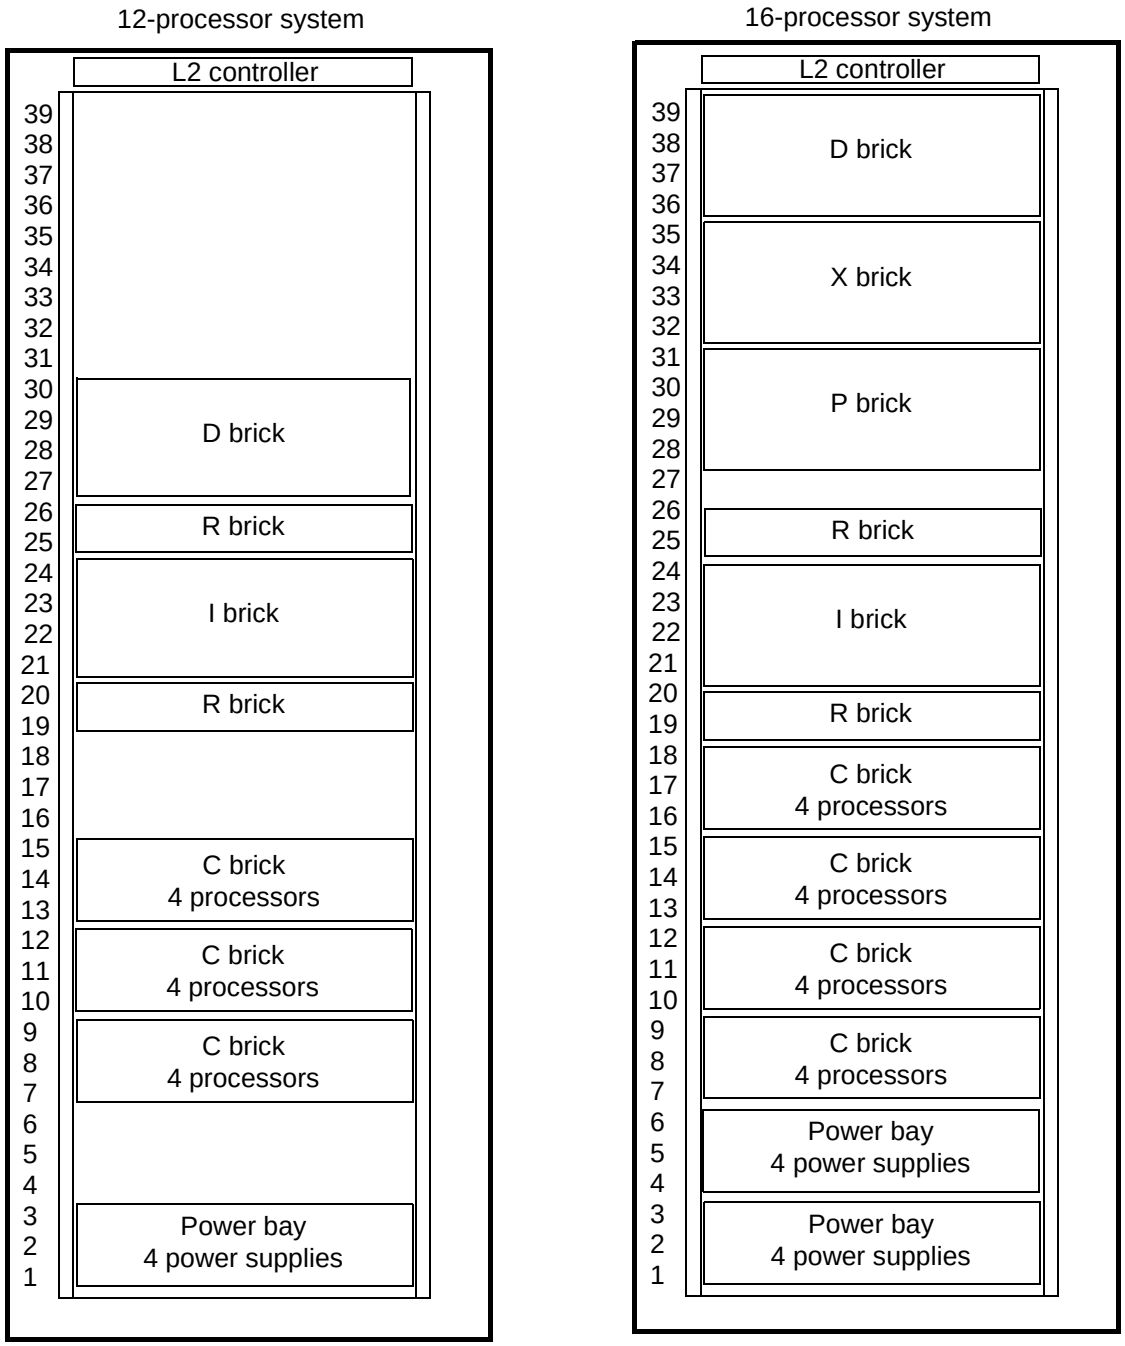

Figure 1-3 SGI Origin 3400 Servers (D-brick and I/O racks not shown)

1.1.1.3 SGI Origin 3800 Server

An SGI Origin 3800 server consists of 16 to 512 processors (refer to Figure 1-4, Figure 1-5, Figure 1-6, and Figure 1-7) and uses 1 to 16 compute racks that contain:

- 4 to 8 C bricks per compute rack
  - Minimum amount of memory - 2 Gbytes (1 rack that contains 4 C bricks)
  - Maximum amount of memory - 1024 Gbytes (16 racks; each rack contains 8 C bricks)
- 2 or 3 R bricks per compute rack

  Note: The third R brick is a MetaRouter or repeat router. A MetaRouter routes information between R bricks (for systems that have more than 128 processors). A repeat router routes information between MetaRouters and routers (for systems that have more than 256 processors).
- 1 or 2 power bays per compute rack; each power bay contains 4 power supplies
- 1 L2 controller per compute rack

The SGI Origin 3800 server also consists of 1 to 8 I/O racks that contain:

- 1 to 8 I, P, or X bricks per I/O rack

  Note: One of the I/O racks must contain an I brick.
- D bricks (optional)
- 1 power bay that contains 4 power supplies per I/O rack

The SGI Origin 3800 server also includes:

- D-brick racks
- G bricks (optional)

  Note: When an Origin 3000 series server contains a G brick(s), it is referred to as an SGI Onyx 3000 series server.
- An L3 controller (optional)

An SGI Origin 3800 server is upgraded by:

- Increments of 4 processors when the existing system has between 16 and 128 processors
- Increments of 32 processors when the existing system has 128 processors or more
- Individual bricks or racks for I/O

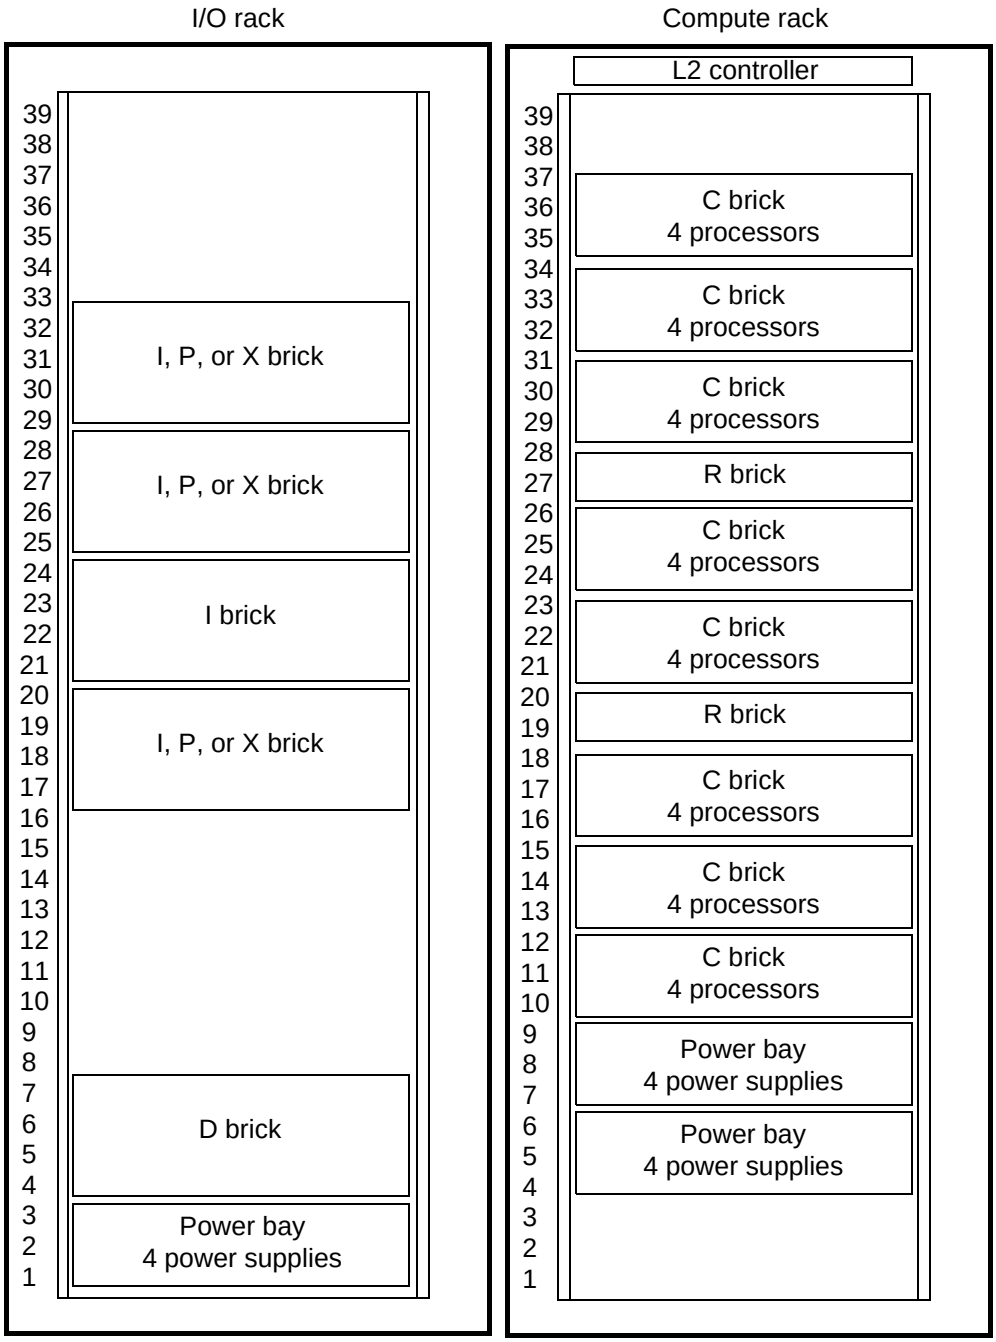

Figure 1-4 32-processor SGI Origin 3800 Server (D-brick racks not shown)

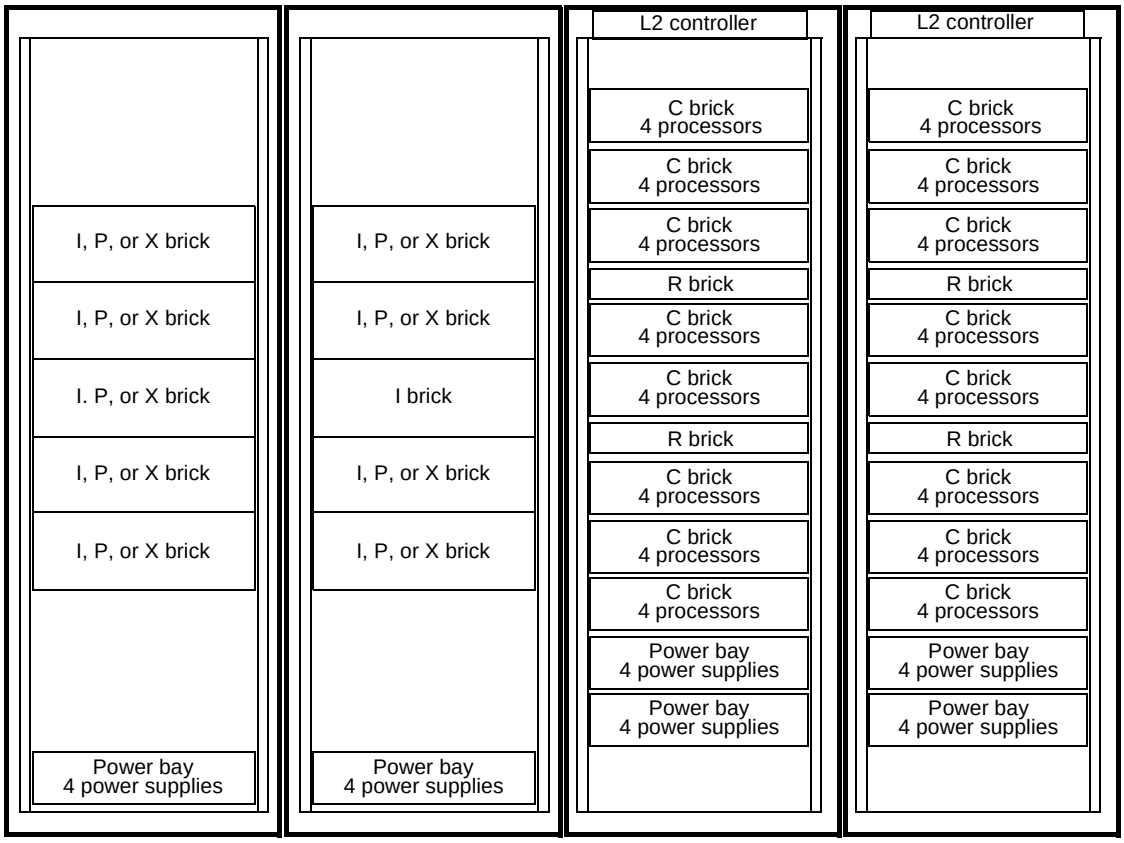

Figure 1-5 64-processor SGI Origin 3800 Server (D-brick racks not shown)

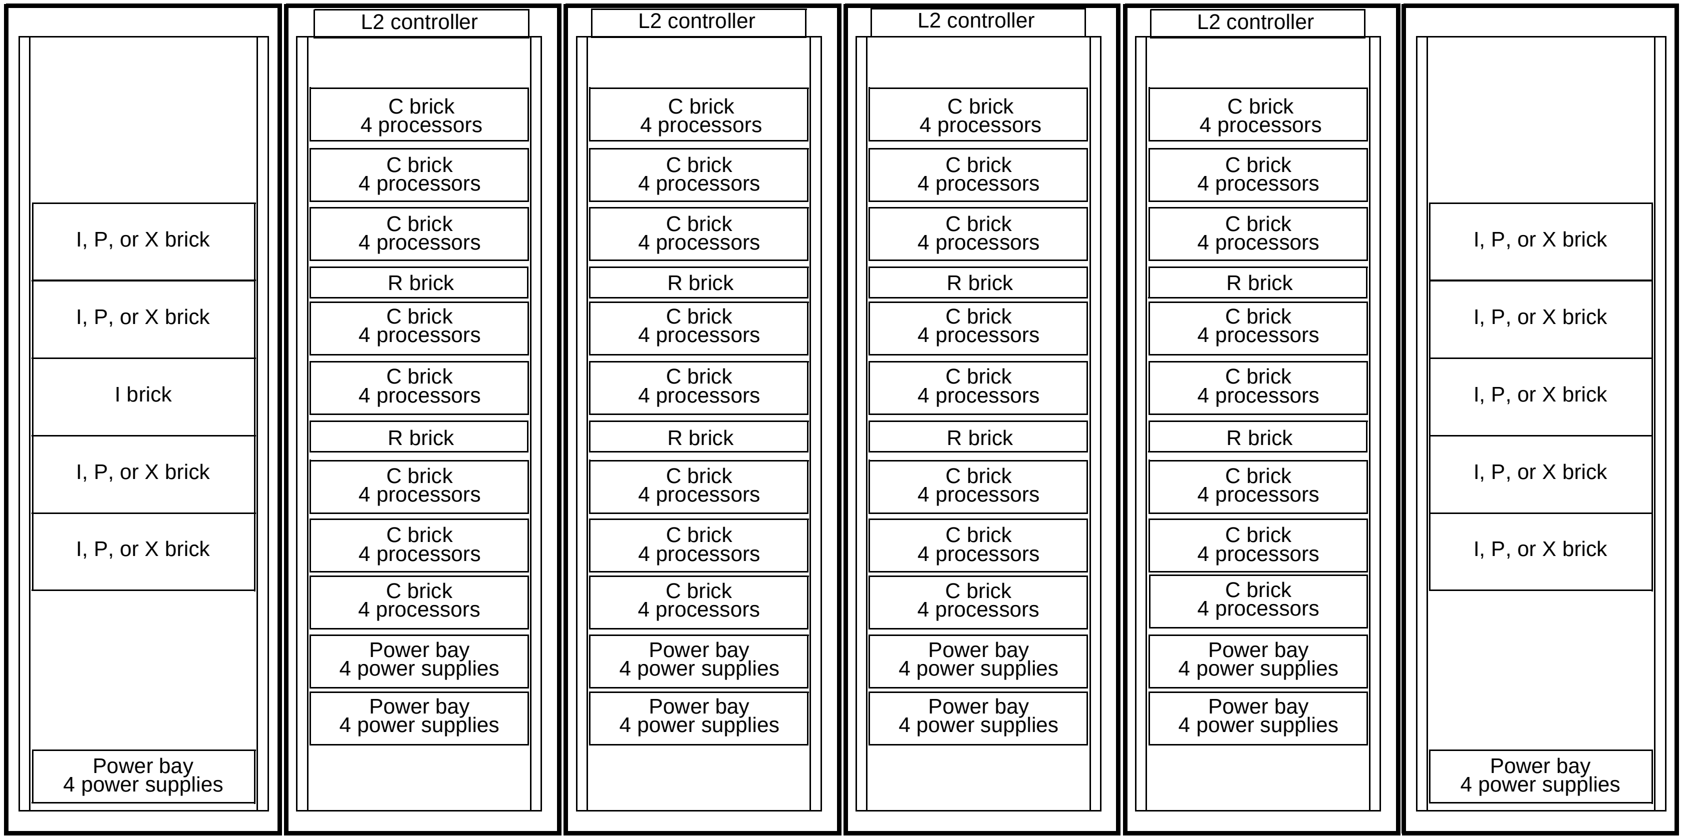

Figure 1-6 128-processor SGI Origin 3800 Server (D-brick racks not shown)

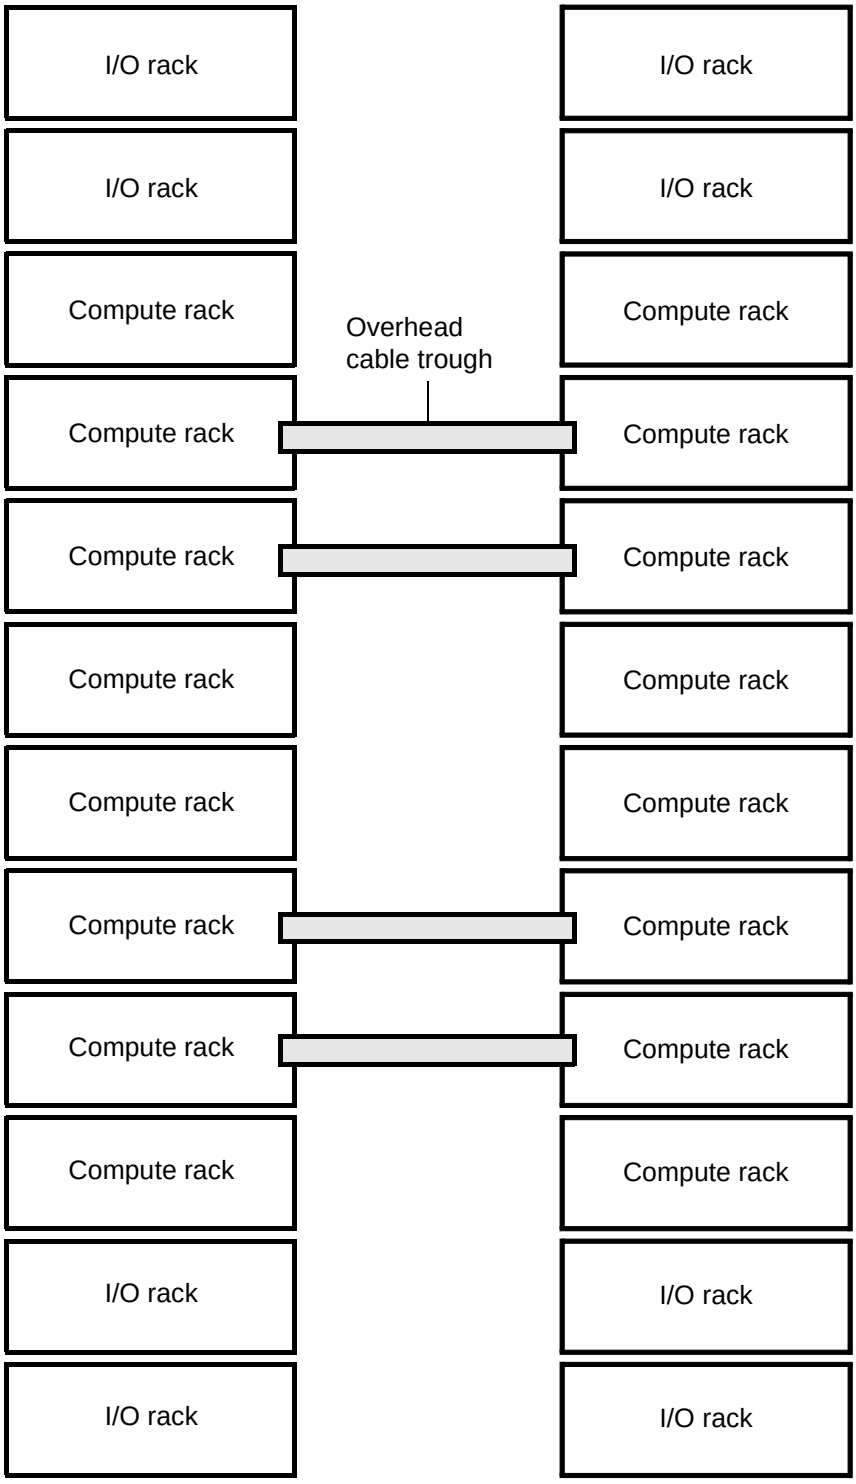

Figure 1-7 512-processor SGI Origin 3800 Server (D-brick racks not shown)

1.2 SGI SNIA 3000 Series Server

The SGI SNIA 3000 series servers:

- Scale from 2 to 8 processors in a single-system image

  Note: At a later date, the SGI SNIA 3000 series servers will scale from 2 to 64 processors in a single-system image, and they will support partitioning in 64-processor to 512-processor systems and clustering.

- Support the Intel Itanium processors
- Use a peripheral component interface (PCI) based I/O system
- Use the Linux operating system

The SGI SNIA 3000 series servers consist of peripheral devices, I/O interfaces, routers, and compute nodes that are linked by a NUMAlink 3 interconnect network. These system components are packaged in standard 19-inch rack-mounted subassemblies that are referred to as bricks.

An SGI SNIA 3000 series server can consist of the following bricks:

- C brick - Houses the compute node that provides the compute functionality and memory for the system. The key components of the compute node are (refer to Figure 1-8):
  - 2 or 4 Intel Itanium processors
  - Memory - includes a directory for cache coherence
  - Hub (Bedrock) ASIC - allows communication among the node components, I/O interfaces, and the NUMAlink 3 interconnect network
  - Synergy ASIC - allows communication between the hub ASIC and the Intel Itanium processors
- R brick - Houses a router that routes information between C bricks, either directly or through other R bricks.

  Note: Routers make up the NUMAlink 3 interconnect network.
- I brick - An I/O interface that provides the base I/O functions for the SGI SNIA 3000 series servers, which includes 1 or 2 system disks for boot functions. The I brick also houses up to 5 PCI cards.
- P brick - An I/O interface that houses 12 PCI cards.
- D brick - A peripheral device that supports two Fibre Channel loops and houses 12 disk drives.

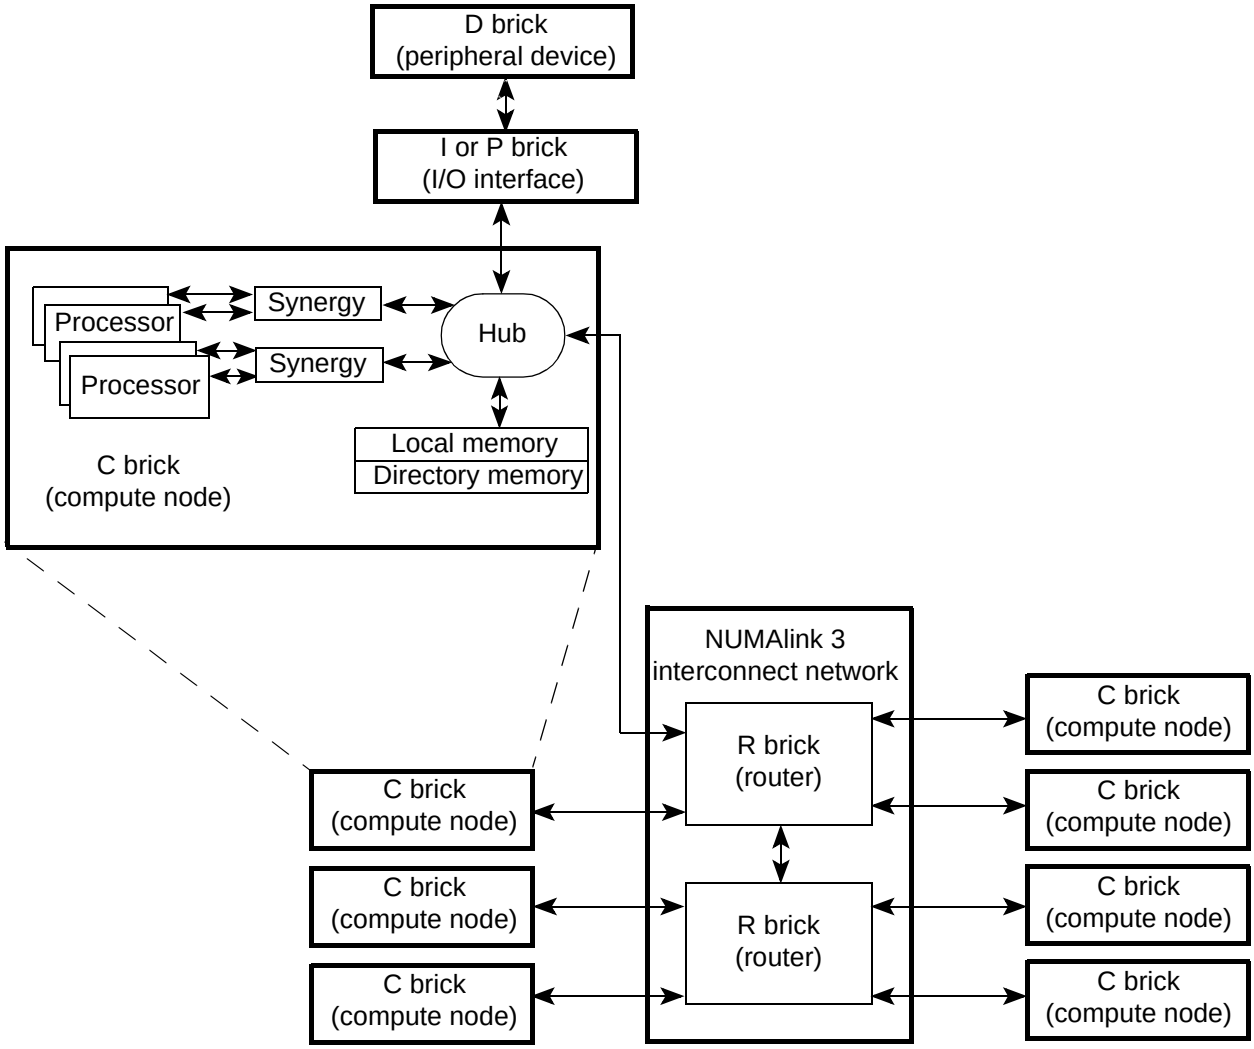

Figure 1-8 SGI SNIA 3000 System Block Diagram

1.2.1 SGI SNIA 3000 Series Server Configurations

Initially, the SGI SNIA 3000 series servers only support 2 to 8 processor in a single-system image (SSI). Eventually, the SGI SNIA 3000 series product line will include the following three single-system image (SSI) classifications:

- The SGI SNIA 3200 server consists of 2 to 8 processors.
- The SGI SNIA 3400 server consists of 4 to 32 processors.
- The SGI SNIA 3800 server consists of 16 to 64 processors.

In addition, the SGI SNIA 3000 series servers will support partitioning of 64-, 128-, and 512-processor systems. The SGI SNIA 3000 series servers will also support clustering. (The maximum number of clustered processors is not defined yet.)

Note: A partition has an independent operating system, can access files from other partitions, and requires an I brick. A cluster consists of standalone systems that are connected to other standalone systems via a high-speed network.

The SGI SNIA 3000 series servers can be expanded by increasing the number of bricks or memory size. Table 1-2 lists the minimum and maximum SSI system configurations.

Table 1-2 Minimum and Maximum SSI Configurations

System Component	Minimum	Maximum
Processors	2	64
Memory	512 Mbytes (1 C brick that has 2 banks of memory)	128 Gbytes (16 C bricks; each C brick has 8 banks of memory)
I/O channels	1	16 (1 I/O channel per C brick)

1.2.1.1 SGI SNIA 3200 Server

An SGI SNIA 3200 server consists of 2 to 8 processors and uses one 17-unit (U) rack. (One U is 1.75 in.) The 17-U rack contains (refer to Figure 1-9):

- 1 or 2 C bricks
  - Minimum amount of memory - 512 Mbytes (1 C brick that has 2 banks of memory)
  - Maximum amount of memory - 16 Gbytes (2 C bricks; each C brick has 8 banks of memory)
- 1 I brick that has a PCI VGA card for Linux installation from a CD-ROM
- 1 P brick
- 1 power bay that contains 4 power supplies

Note: The SGI SNIA 3200 server does not use an R brick.

The SGI SNIA 3200 server might also include:

- Additional D bricks in separate racks (maximum of 9 D bricks per rack)
- A rack (L2) controller or system (L3) controller (both are optional)

An SGI SNIA 3200 server is upgraded by individual bricks or by increments of 2 processors. For example, if the SGI SNIA 3200 server has 1 C brick that contains 2 processors, the system can be upgraded to a 4-processor system by adding 2 processors to the C brick. An SGI SNIA 3200 system that has one 4-processor C brick can be upgraded to an 8-processor system by adding another 4-processor C brick.

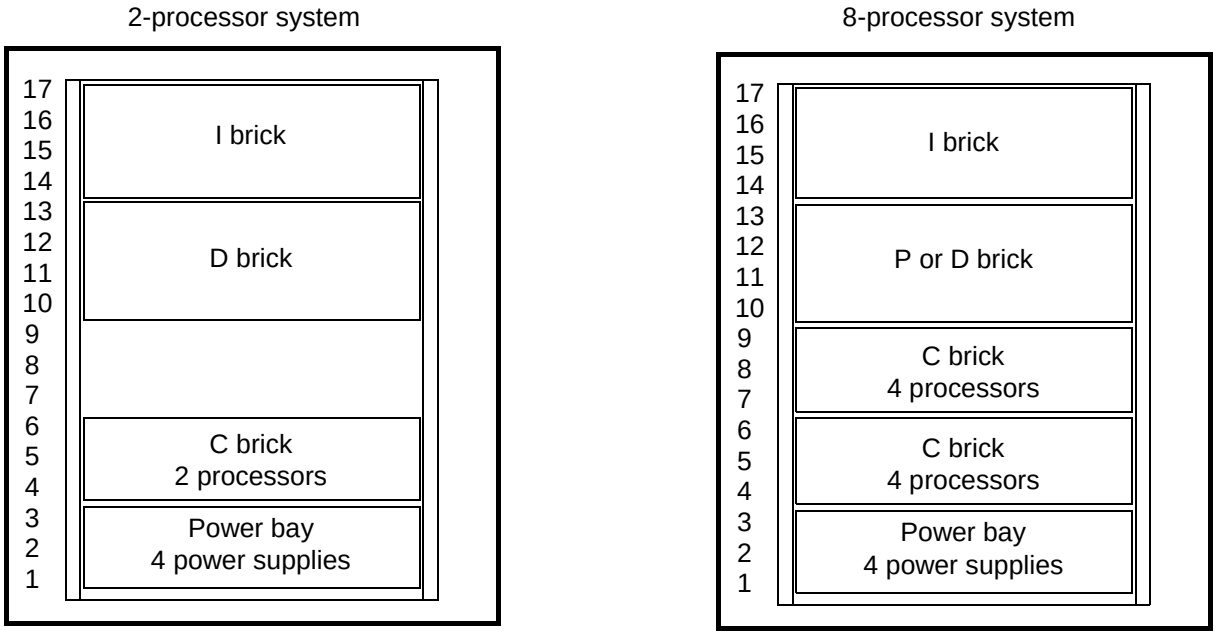

Figure 1-9 SGI SNIA 3200 Servers

1.2.1.2 SGI SNIA 3400 Server

An SGI SNIA 3400 server consists of 4 to 32 processors (refer to Figure 1-10) and uses one 39-U compute rack that contains:

- 1 to 8 C bricks
  - Minimum amount of memory - 512 Mbytes (1 C brick with 2 banks of memory)
  - Maximum amount of memory - 64 Gbytes (8 C bricks; each C brick has 8 banks of memory)
- 1 I brick
- 2 R bricks
- 1 or 2 power bays; each power bay contains 6 power supplies
- 1 L2 controller

The SGI SNIA 3400 server might also include:

- Additional D bricks in separate racks (maximum of 9 D bricks per rack)
- P bricks or additional I bricks in the compute rack
- An L3 controller (optional)

An SGI SNIA 3400 server is upgraded by individual bricks and by increments of 4 processors. For example, to upgrade a 8-processor system to a 16-processor system, two 4-processor C bricks are added to the system.

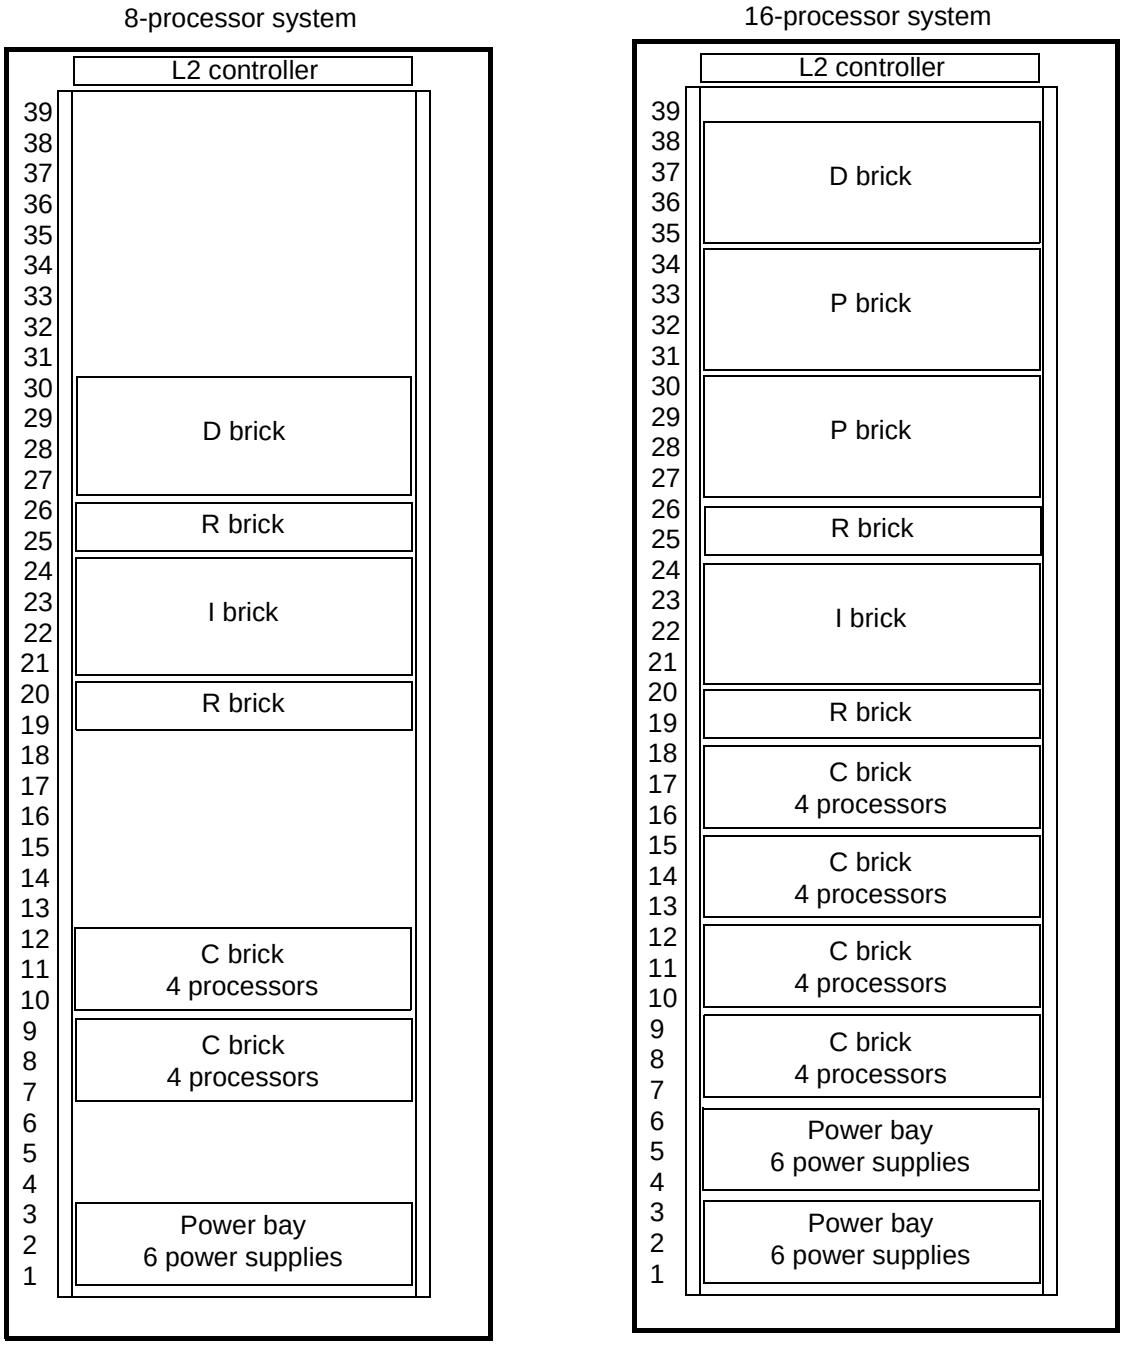

Figure 1-10 SGI SNIA 3400 Servers (D-brick and I/O racks not shown)

1.2.1.3 SGI SNIA 3800 Server

An SGI SNIA 3800 server consists of 16 to 64 processors (refer to Figure 1-11 and Figure 1-12) and uses 1 or 2 compute racks that contain:

- 4 to 8 C bricks per compute rack
  - Minimum amount of memory - 2 Gbytes (1 rack that contains 4 C bricks; each C brick has 2 banks of memory)
  - Maximum amount of memory - 128 Gbytes (2 racks; each rack contains 8 C bricks)
- 2 R bricks per compute rack
- 1 or 2 power bays per compute rack; each power bay contains 6 power supplies
- 1 L2 controller per compute rack

The SGI SNIA 3800 server also consists of 1 or 2 I/O racks that contain:

- 1 to 8 I or P bricks per I/O rack

  Note: One of the I/O racks must contain an I brick. An I/O rack can also contain D bricks.
- 1 power bay that contains 4 power supplies per I/O rack

The SGI SNIA 3800 server also includes:

- D-brick racks
- An L3 controller (optional)

An SGI SNIA 3800 server is upgraded by individual bricks and by increments of 4 processors. For example, to upgrade a 16-processor system to a 20-processor system, a 4-processor C brick is added to the system.

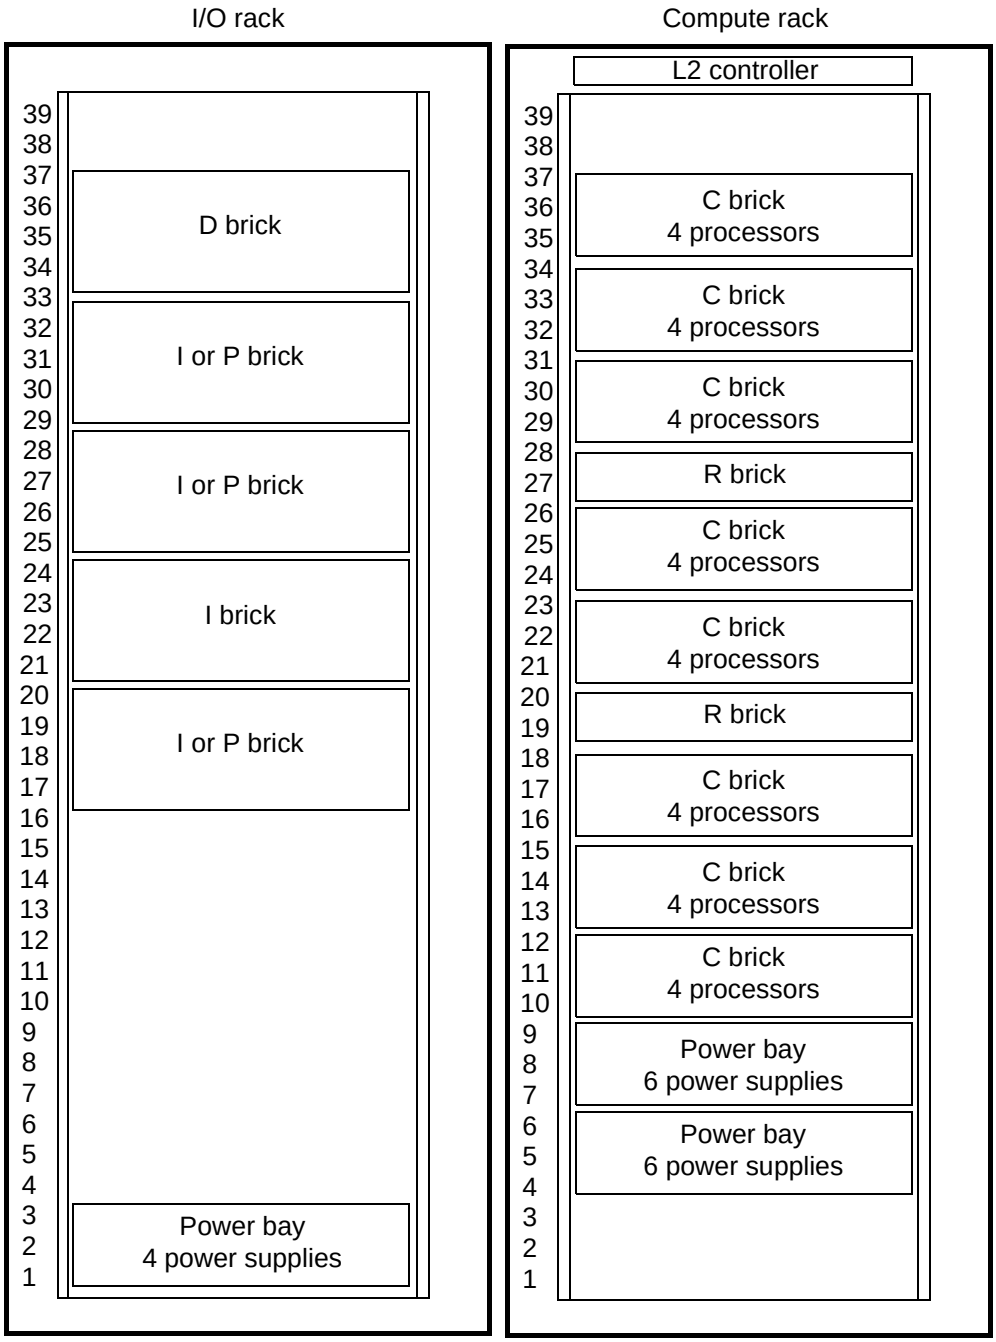

Figure 1-11 32-processor SGI SNIA 3800 Server (D-brick racks not shown)

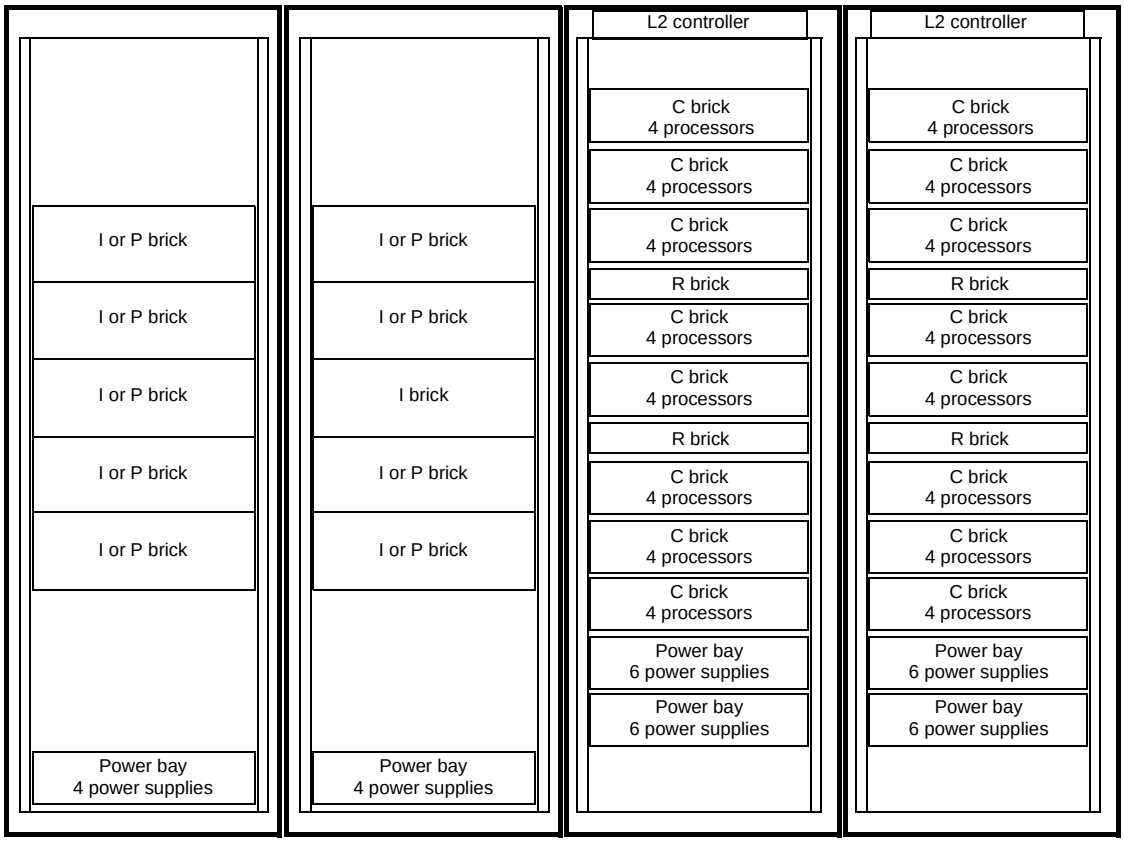

Figure 1-12 64-processor SGI SNIA 3800 Server (D-brick racks not shown)

1.3 SGI Origin 3000 and SGI SNIA 3000 Series Server Differences

Table 1-3 lists the key differences between the SGI Origin 3000 and SGI SNIA 3000 series servers.

Table 1-3 SGI Origin 3000 and SGI SNIA 3000 Series Server Differences

Server	Supported Operating System	Supported Bricks	Supported Processor(s)	Supported Caches
SGI Origin 3000 series	IRIX	C*, R, I, P, X, G, and D	2 to 512 MIPS processors: MIPS R12000, MIPS R12000A, MIPS R14000, and their successors	32-Kbyte primary (L1) cache for data 32-Kbyte L1 cache for instructions 4- or 8-Mbyte secondary (L2) cache **Note:** The size of the L2 cache depends on the processor type.
SGI SNIA 3000 series	Linux	C*, R, I, P, and D	2 to 64 Intel Itanium processors (single-system image) 64 to 512 Intel Itanium processors (partitioning)	16-Kbyte L1 cache for data 16-Kbyte L1 cache for instructions 96-Kbyte L2 cache 2- or 4-Mbyte tertiary (L3) cache 64-Mbyte quaternary (L4) cache

* The SGI Origin 3000 series servers and the SGI SNIA series servers use different C bricks.

The SGI SNIA 3000 series server also has an additional ASIC: the Synergy ASIC. This ASIC is the interface between the Intel Itanium processors and the hub ASIC.

Reader Comment Form

Title: Product Overview
SGI™ Origin™ 3000 and
SGI™ SNIA 3000 Server Series

Number: 108-xxxx-xxx

Your feedback on this publication will help us provide better documentation in the future. Please take a moment to answer the few questions below.

For what purpose did you primarily use this document?

_____Troubleshooting
_____Tutorial or introduction
_____Reference information
_____Classroom use
_____Other - please explain

_____

Using a scale from 1 (poor) to 10 (excellent), please rate this document on the following criteria and explain your ratings:

_____Accuracy _____
_____Organization _____
_____Readability _____
_____Physical qualities (binding, printing, page layout) _____
_____Amount of diagrams and photos _____
_____Quality of diagrams and photos _____

Completeness (Check one and explain your answer)

_____Too much information _____ Too little information _____Correct amount

_____
_____
_____

You may write additional comments in the space below. Mail your comments to the address below, fax them to us at +1 715 726 4353, or e-mail them to us at *spt@sgi.com.* When possible, please give specific page and paragraph references. We will respond to your comments in writing within 48 hours.

_____
_____
_____

NAME _____
JOB TITLE _____
E-MAIL ADDRESS _____
SITE/LOCATION _____
TELEPHONE _____
DATE _____

[or attach your business card]

Attn: Service Publications and Training
890 Industrial Boulevard
P.O. Box 4000
Chippewa Falls, WI 54729-0078
USA