Origin2000™ and Onyx2™ Deskside and Rackmount Installation Instructions

Document Number 108-0155-003

Contributors

Written by Pablo Rozal, Carl Strasen and Tom Dye
Illustrated by Cheri Brown and Dan Young
Production by Kirsten Johnson
Engineering contributions by David Alexander, Jim Ammon, Carolyn Curtis, Ed Freige, Michael Galles, Darrin Goss, Charles Hart, Joe Heinrich, John Hwa, Kirk Law, Bob Marinelli, Phil Montalban, Bob Murphy, Reuel Nash, Carlin Otto, Dave Parry, Bharat Patel, Michael Ramirez, Mark Schwenden, Sam Sengupta, Richard Singer, Carl Strasen, Steve Whitney, Ted Wong, Neill Wood

FCC Warning

This equipment has been tested and found compliant with the limits for a Class A digital device, pursuant to Part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

Attention

This product requires the use of external shielded cables in order to maintain compliance pursuant to Part 15 of the FCC Rules.

Origin2000 and Onyx2 Deskside and Rackmount Installation Instructions
Document Number 108-0155-003

Silicon Graphics, Inc.
Mountain View, California

Contents

Figures

Tables

Introduction

This manual describes the hardware and software information necessary to install, configure an Origin2000 and Onyx2 deskside, rackmount, or multirack system. The Origin2000 is Silicon Graphics® next-generation *server* system that provides revolutionary improvement and expandabilty over the previous-generation CHALLENGE® system. The Onyx2™ is Silicon Graphics next-generation *graphics* system and also provides significant performance capabilities over the previous-generation Onyx™ systems.

This manual is written for Silicon Graphics® system support engineers (SSEs) and third-party field support groups responsible for product installation and testing.

What This Document Contains

This manual is divided into the following chapters and appendixes:

Related Documents

This section lists related Silicon Graphics documents that you may need to consult for additional information.

Site Preparation for Origin Family and Onyx2 (P / N 007-3452-xxx)

Origin2000 Deskside Owner's Guide (P / N 007-3453-xxx)

Origin2000 Rackmount Owner's Guide (P / N 007-3456-xxx)

Onyx2 Deskside Owner's Guide (P / N 007-3454-xxx)

Onyx2 Rackmount Owner's Guide (P / N 007-3457-xxx)

Note: See also the Technical Publications Library Web page at http: / / techpubs.engr.sgi.com

Safety Information

Be sure to read the following information before you begin installation.

Warning: Installation of these products requires specific training and technical knowledge. These instructions are provided for use by Silicon Graphics system support engineers (SSEs) or other Silicon Graphics-trained personnel only. This equipment utilizes electrical power internally that is hazardous if the equipment is improperly disassembled.

Caution: This equipment is extremely sensitive and susceptible to damage caused by electrostatic discharge (ESD), the build-up of electrostatic potential on clothing and other materials. Use the following proper ESD measures.

- Connect a ground strap to your wrist when installing and removing peripherals.
- Be sure that you and all the electrical equipment you handle during this installation remain at a ground potential of zero to avoid damage from ESD.
- Remove a board from its antistatic bag only when you are properly grounded with a ground strap and only when you are working on the board or installing it.
- Do not use an ohmmeter on a board.

Typographic Conventions

These type conventions and symbols are used throughout this manual:

Italics Filenames, variables, IRIX command arguments, command flags, titles of publications, icon names.

`Screen type` Code examples, file excerpts, and screen displays (including error messages).

`Bold Screen type`

User input.

(Parentheses) Following IRIX commands, they surround the reference page (man page) section where the command is described.

[Brackets] Surround optional syntax statement arguments.

IRIX shell prompt for the superuser (*root*).

<Angle brackets>

Nonprinting keys such as <Enter> or command-line <variables>.

Chapter 1

Product Overview

This chapter provides an overview of the Origin2000 and Onyx2 system features, architecture, major components, and new terminology.

1.1 Origin2000 System

The Origin2000 server system is a revolutionary follow-on to the CHALLENGE®-class symmetric multiprocessing (SMP) system. It provides a highly scalable and configurable architecture that can employ a number of powerful server modules. A module can be configured as a standalone unit or grouped together through a high-speed *CrayLink*™ *Interconnect* link to form a supercomputing platform.

The base unit or module is called the Origin2000 deskside system (see Figure 1-1). The deskside system consists of one to eight CPUs with 64 MBs to 4 GBs of memory and can provide a wide variety of I/O interfaces (see Appendix A, "System Specifications"). The Origin2000 configuration also comes in a single rackmount system that contains up to 16 processors (see Figure 1-2) and a multirack system configuration with up to 128 processors (see Figure 1-3).

1.2 Onyx2 System

The Onyx2 graphics supercomputer provides next-generation visualization hardware in a scalable and modular system architecture. The deskside system (see Figure 1-1) has up to four CPUs and either a Reality™ or an InfiniteReality™ board set. The rackmount system (see Figure 1-2) contains up to eight CPUs and two InfiniteReality board sets (or graphics pipes). The multirack system configuration (see Figure 1-3) has up to 24 CPUs or up to four graphics pipes.

The Oynx2 system combines high-performance Reality or InfiniteReality graphics into the revolutionary architecture of the Origin2000 system to provide powerful and scalable graphics workstations. As Figure 1-1, Figure 1-2, and Figure 1-3 show, the Onyx2 and Origin2000 deskside and rackmount systems share the same chassis types. As a result, a number of boards, peripherals, and other hardware components can be swapped between the different systems (see Section 1.6, "Hardware Overview," for additional information).

1.3 Features

This section describes the new features of the Origin2000 and Onyx2 systems.

1.3.1 Origin2000 Capabilities

The Origin2000 system is ideal for evolving applications requiring expansion capability. The Origin2000 system also provides upgradability to larger rack-based configurations. Here is a list of additional features:

- Significantly lower entry system costs (with pay-as-you-grow expandability).
- Support of a large number of processors (up to 128).
- High bandwidth I/O connectivity (up to 800 MB/sec transfer rate, per direction).
- High total memory bandwidth (up to 256 GB of memory in a 128-processor configuration).
- Optional connectivity to peripheral connector interface (PCI) boards.
- Modular architecture that allows service personnel to independently power off portions of the system to swap and maintain system components without disabling the entire system.
- Superscalar R10000™ CPU (in the IP27 Node board) supports advanced memory latency tolerance features such as out-of-order execution and advanced branch prediction to address real-world application demands.
- Large variety of peripheral connectivity options such as Ultra SCSI and Fibre Channel. The Origin2000 I/O subsystem also supports 100Base-T Ethernet, Serial HIPPI, and ATM to provide high-performance communications in a heterogeneous networked environment.

1.3.2 Onyx2 Capabilities

The Onyx2 systems provide most of the capabilities listed in Section 1.3.1 in addition to these other graphics capabilities.

- Available in one to four graphics-pipe configurations
- Nearly twice the bandwidth of RealityEngine (400 MB vs. 220 MB)
- Some configurations come with a super-wide 1920 × 1200 (24-inch) monitor
- Baud rates up to 115200
- Full-speed volume rendering

1.4 Major Differences From Earlier Systems

The Origin2000 and Onyx2 system architectures differ from previous-generation Silicon Graphics systems in the following major ways.

- The Origin2000 and Onyx2 systems do not use a centralized system bus (such as the Ebus). Instead, the system architecture can use the high-speed (800 MB/sec) switched-based CrayLink Interconnect (see Figure 1-9), along with specialized I/O ASICs to allow main memory to be accessed by other CPUs in the configuration and to enable processor-to-device communications.
- The Origin2000 and Onyx2 systems do not use a bulkhead with internal cabling for external I/O connections. External connections are made through the rear of the system module, directly into the board connectors.
- The Origin2000 and Onyx2 systems provide optional connection to the industry-standard peripheral connector interface (PCI) bus for third-party board expansion.
- The Origin2000 and Onyx2 systems use a modular, standard 19-inch rackmountable chassis.
- The Origin2000 and Onyx2 systems use a compression connector along with a cam mechanism to help install many boards correctly and accurately into the chassis. The earlier Silicon Graphics systems used metal pins, which were more susceptible to potential damage and to possible board misalignment problems.
- The Origin2000 and Onyx2 systems boards have a different board form factor. Unlike previous-generation Silicon Graphics systems, the Origin2000 XIO cardcage boards do not have a standard 9U or 6U form factor. The half-height boards are $10 \times 6.5 \times 1$ inches. The full-height boards measure $10 \times 13 \times 1$ inches, and the CHALLENGE and Onyx boards measure $25 \times 20 \times 32$ inches. Therefore, you *cannot* install boards from earlier Silicon Graphics systems into an Origin2000 or Onyx2 system.
- The Origin2000 and Onyx2 systems use single-connector attachment (SCA) drives and a new sled design to help simplify drive installation.
- The Origin2000 and Onyx2 systems have a new (“one-size-fits-all”) power supply. There are no longer separate power boards or online switching supplies (OLSs) to add or swap when power requirements change.
- The system controller messages have been simplified. Messages are now fewer and easier to interpret.
- The Origin2000 and Onyx2 systems have a combinational node board, which houses the R10000 CPUs, main memory, and some of the I/O subsystem. There are no separate CPU or memory boards.

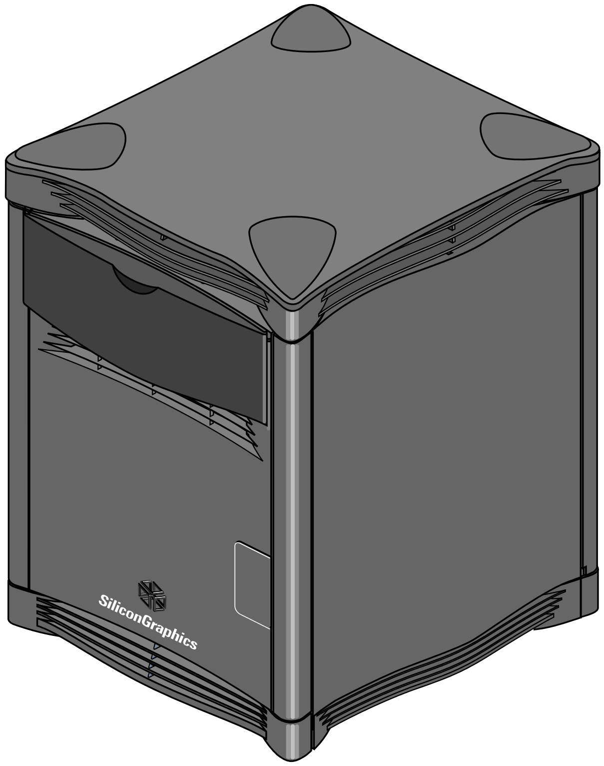

Figure 1-1 Origin2000 or Onyx2 Deskside System

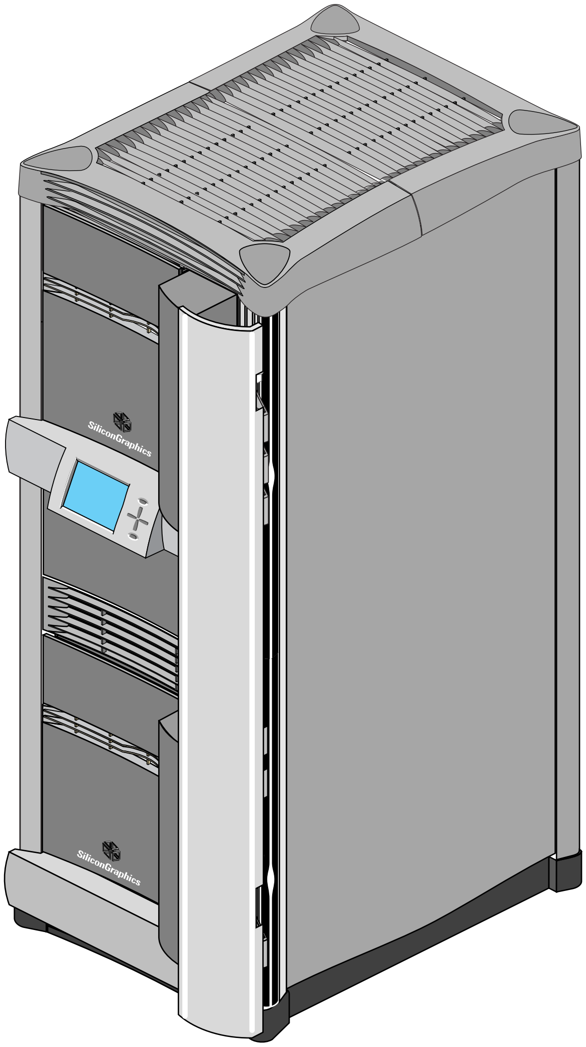

Figure 1-2 Origin2000 or Onyx2 Rackmount System

Figure 1-3 Origin2000 or Onyx 2 Multirack Configuration (Two Rackmount Systems)

Figure 1-4 Another Example of an Origin2000 Multirack Configuration (Four Racks Shown)

1.5 Evolution of the Origin2000 and Onyx2 Systems

Figure 1-5 shows the development path of the Origin2000 system and Figure 1-6 shows the development of the Onyx2 system. The scalable shared-memory multiprocessing (S2MP) architecture can support more than three times the number of processors and can house 16 times more memory than the previous generation CHALLENGE system. The Onyx2 can support more graphics pipes than the earlier Onyx system.

In addition, a single bidirectional CrayLink Interconnect running at 1600 MBs / sec can sustain more bandwidth than the entire CHALLENGE or Onyx Ebus. The S2MP architecture uses distributed shared memory (DSM) to enable read and write access into main memory (see Section 1.5.3 for more information). This shared memory is accessible to all processors through the CrayLink Interconnect and can be accessed with low latency.

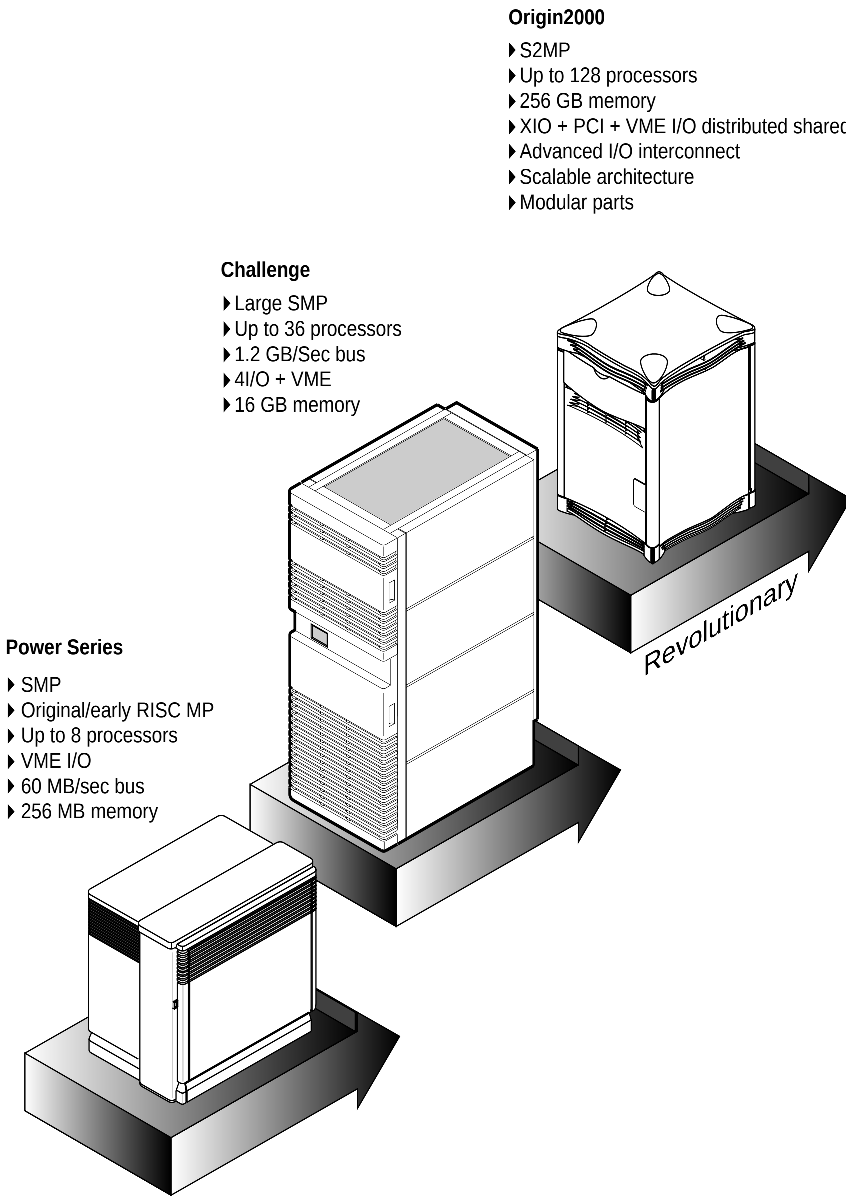

Figure 1-5 SGI Multiprocessing Architectures

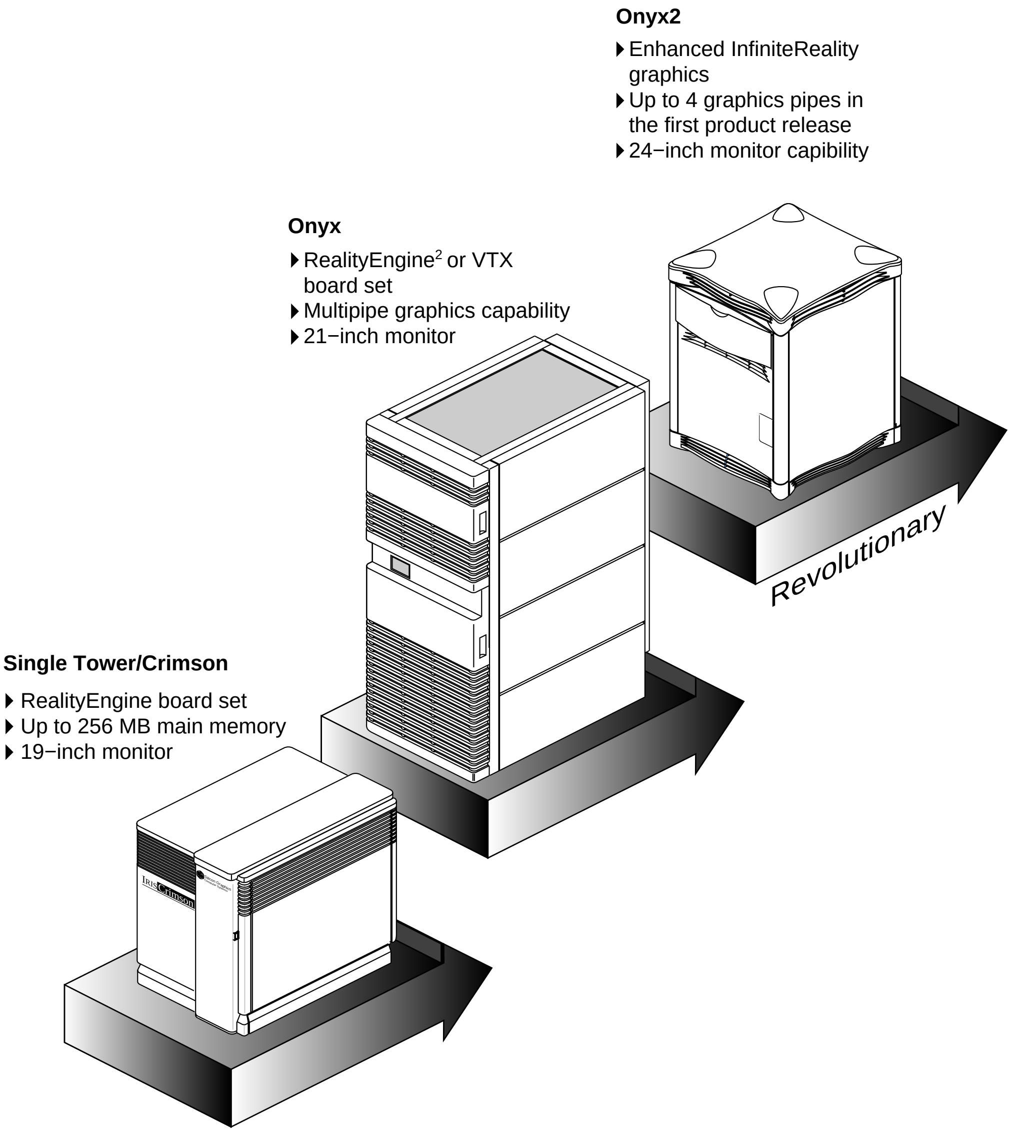

Figure 1-6 Onyx2 Evolutionary Path

As illustrated in Figure 1-7, the Origin2000 and Onyx2 systems are a number of processing modules linked together by the CrayLink Interconnect. Each processing node board contains either one or two processors, a portion of shared memory, a directory for cache coherence, and two interfaces: one that connects to I/O devices and another that links system modules through the CrayLink Interconnect.

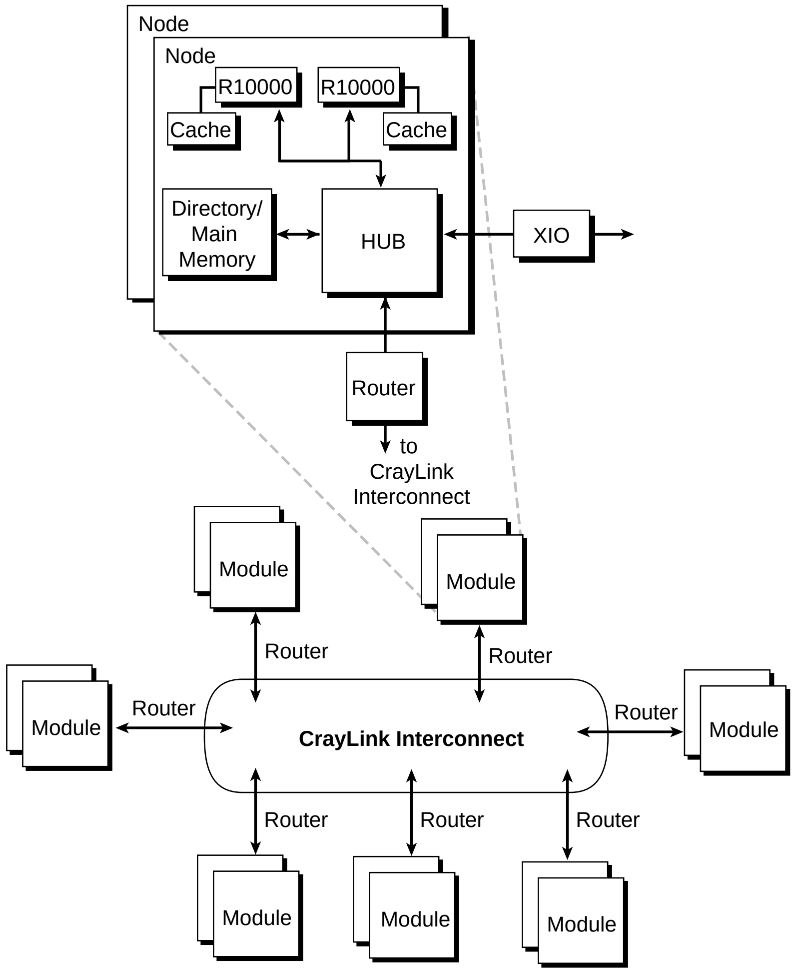

Figure 1-7 Modules in an Origin2000 or Onyx2 System

The CrayLink Interconnect links modules to one another. It may appear to be a type of super data bus, but it differs from a bus in important ways. A bus is a resource that can be used by only one processor at a time. The CrayLink Interconnect is based on an interconnection fabric technology whereby a a *mesh* of multiple, simultaneous, dynamically allocable connections are made from processor to processor, or processor to I/O for example, as they are needed. This web of connections differs from a bus in the same way that multiple dimensions differ from a single dimension: if a bus is a *one*-dimensional line, then the CrayLink Interconnect is a *multi*dimensional mesh.

The system is said to be scalable because it can range from 1 to 128 processors in an Origin2000 configuration and from 1 to 24 in an Onyx2 multirack configuration. As you add modules, you add to and scale the system bandwidth. The system is also modular because it can be increased in size by adding standard modules to the CrayLink Interconnect (see Section 1.5.1). The CrayLink Interconnect is implemented on cables outside these modules.

1.5.1 CrayLink Interconnect

Origin2000 and Onyx2 modules are connected by the S2MP architecture. The S2MP is a set of switches, called *routers*, that are linked by CrayLink Interconnect cabling in various configurations, or *topologies*. The CrayLink Interconnect differs from a standard bus in the following important ways:

- The CrayLink Interconnect is a mesh of multiple point-to-point links connected by the routing switches. These links and switches which are based on the interconnection fabric technology, allow multiple transactions to occur simultaneously.
- The links permit extremely fast switching (800 MBs / sec in each direction, 1,600 MBs / sec bidirectionally).
- The CrayLink Interconnect does not require arbitration and is not as limited by contention, whereas a bus must be contested for through arbitration.
- More routers and links are added as modules are added, increasing the CrayLink Interconnect's bandwidth. A shared bus has a fixed bandwidth that is not scalable.
- The topology of the S2MP is such that the bisection bandwidth grows linearly with the number of modules in the system.

The CrayLink Interconnect provides a minimum of two separate paths to every pair of Origin2000 modules. This redundancy allows the system to bypass failing routers or broken fabric links. Each fabric link is additionally protected by a CRC code and a link-level protocol, which retry any corrupted transmissions and provide fault tolerance for transient errors.

Figure 1-8 illustrates an eight-Node board hypercube with its multiple datapaths. R1 can simultaneously communicate with R0, R2 with R3, R4 with R6, and R5 with R7, all without having to interface with any other Node board.

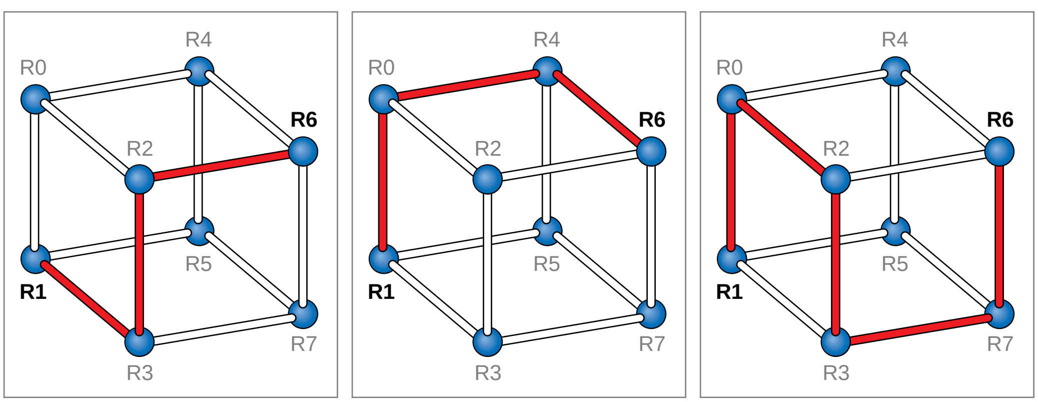

Figure 1-8 Datapaths in the CrayLink Interconnect

1.5.2 Crossbar

Several of the ASICs (Hub, Router, and Crossbow) use a crossbar for linking on-chip inputs with on-chip output interfaces. For instance, an eight-way crossbar is used on the

Crossbow (XBOW) ASIC; this crossbar creates direct point-to-point links among one or more modules and multiple I/O devices. The crossbar switch also allows peer-to-peer communication in which one I/O device can speak directly to another I/O device.

1.5.3 Distributed Shared Address Space (Memory and I/O)

The Origin2000 and Onyx2 systems employ a *distributed shared memory* system architecture where main memory is split among the Node boards. Rather than appearing as one fast memory, main memory is “distributed” over the configuration, with a little piece of the memory near each processor, thus the name distributed shared memory. A directory memory keeps track of the information necessary for hardware coherency and protection.

This architecture differs from a CHALLENGE- or Onyx-class system in which memory is centrally located on and only accessible over a single shared bus. By distributing the Origin2000 and Onyx2 memory among processors, memory latency is reduced; accessing memory near a processor take less time than accessing remote memory. Although physically distributed, all main memory is available to all processors.

The Origin2000 and Onyx2 memory is located in a single shared address space. Memory within this space is distributed among all the processors and is accessible over the CrayLink Interconnect. I/O devices are also distributed within a shared address space; every I/O device is universally accessible throughout the system.

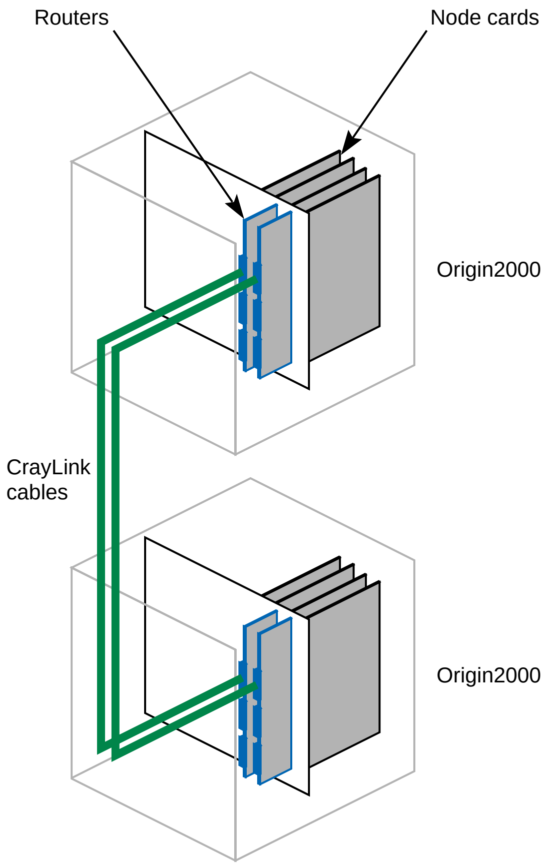

Figure 1-9 Example of CrayLink Interconnect Link in a 16P Origin2000 Configuration

1.6 Hardware Overview

This section provides an overview of the server and graphics hardware for the Origin2000 and Onyx2 systems.

1.6.1 Origin2000 System

Figure 1-10 provides an overall block diagram of an Origin2000 base module. The major hardware components include the

- Node board
- Crossbow (XBOW) ASICs on the midplane
- Router board
- BaseIO board
- XIO slots

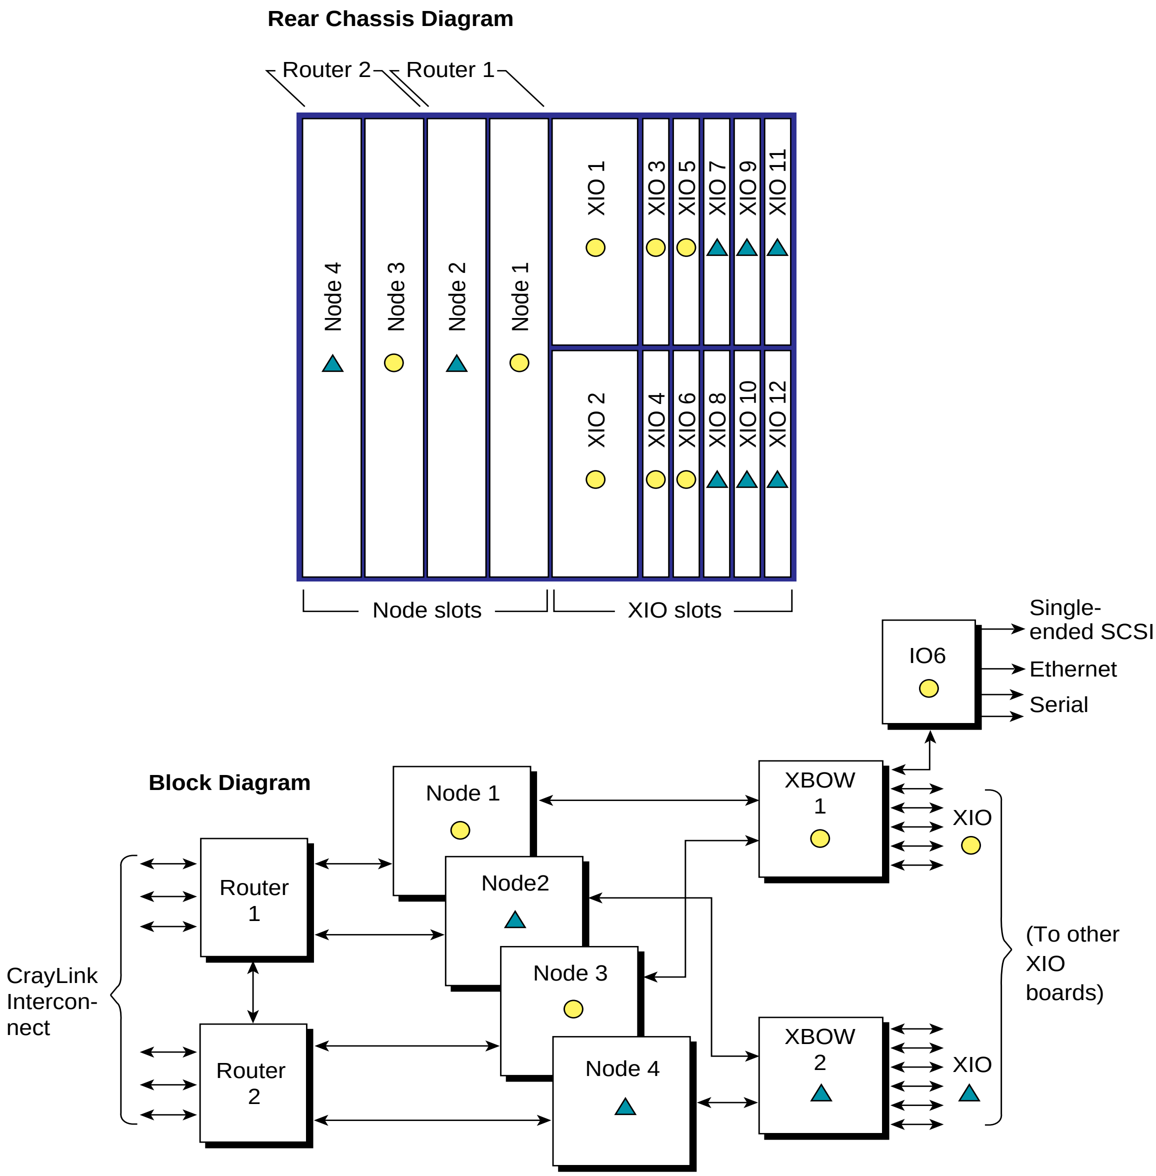

Figure 1-10 Origin2000 Overall Block Diagram (Base Unit)

These components are standard in all system configurations. Table 1-1 briefly describes these components as well as other major new Origin2000 terms.

Communication between the major components takes place through a midplane. This enables the system controller, which is tied to the midplane through a dedicated bus, to monitor systemwide activity.

In Figure 1-10, corresponding Node boards, crossbows, and XIO slots are indicated by circles and triangles. For example, when Node 1 is present, IO slots 1 through 6 (designated by a circle) are activated. When Node 2 is also present, then IO slots 7 through 12 (designated by a triangle) are activated.

The XIO cardcage section provides a nearly fourfold increase in I/O performance in comparison to the CHALLENGE and Onyx I/O subsystem (800 MB/sec for the Origin2000 system versus 220 MBs/sec for CHALLENGE and Onyx). This enables the Origin2000 system to meet demanding I/O requirements such as the multiple streams of video and audio used in multimedia and postproduction environments, virtual reality, and distributed computing.

1.6.2 Onyx2 Graphics Systems

This section provides a brief hardware description of the Onyx2 system. There are two types of deskside configurations— Onyx2 Reality and Onyx2 InfiniteReality. There are also two rack configurations— single-rack and multirack system.

The Onyx2 systems have most of the same basic hardware as the Origin2000 system as well as the specialized graphics components. The base Onyx2 system includes

- Router board
- Node board
- Crossbow
- BaseIO (graphics version)
- XIO cardcage
- Reality or InfiniteReality graphics board set

See Table 1-1 for additional information on these parts.

1.6.2.1 Onyx2 Deskside

Figure 1-11 and Figure 1-12 provide block diagrams of the Onyx2 Reality and InfiniteReality deskside systems. The Onyx2 Reality deskside system, which is the entry-level graphics configuration, uses a GE14-4 board with 4 graphics engines (GEs), the RM8 board, and DG5. The RM8 board combines up to 40 MB of raster memory along with either 16 or 64 MB of texture memory on a single board. The DG5 board provides up to two or eight independent video channels, depending on how the board is configured.

The Onyx2 InfiniteReality deskside system uses a GE14-8 board (with 8 GEs) and 2 RM7 boards that can house two optional TM7 daughterboards, respectively, along with a DG5 board. The RM7 board provides 80 MBs of raster memory, and the TM7 provides either 16, 64, 128, 256 MBs of texture memory. The DG5 board provides up to two or eight independent video channels, depending on how the board is configured.

The GE14 board contains an XG chip that interfaces directly with one of the six XIO slots on the Crossbow ASIC. The XG chip converts the graphics input/output (GIO) signal protocol to XIO signal protocol and vice versa (for additional information, see Chapter 2, "Theory of Operations").

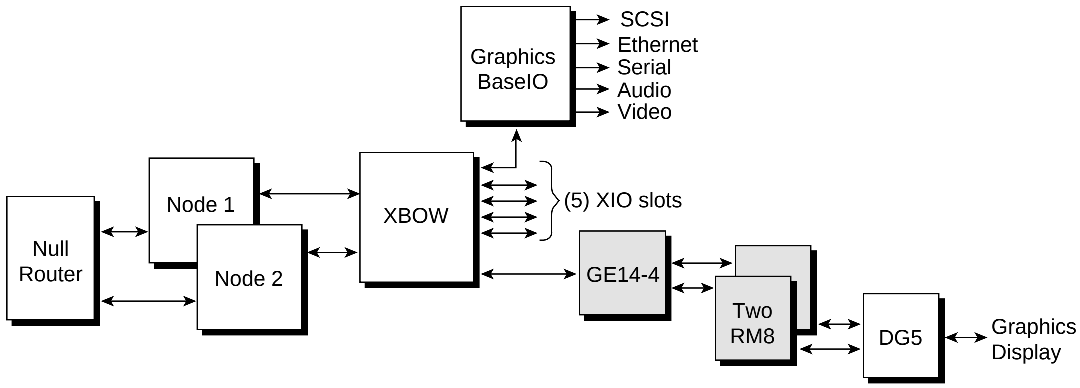

Figure 1-11 Onyx2 Reality Deskside Block Diagram

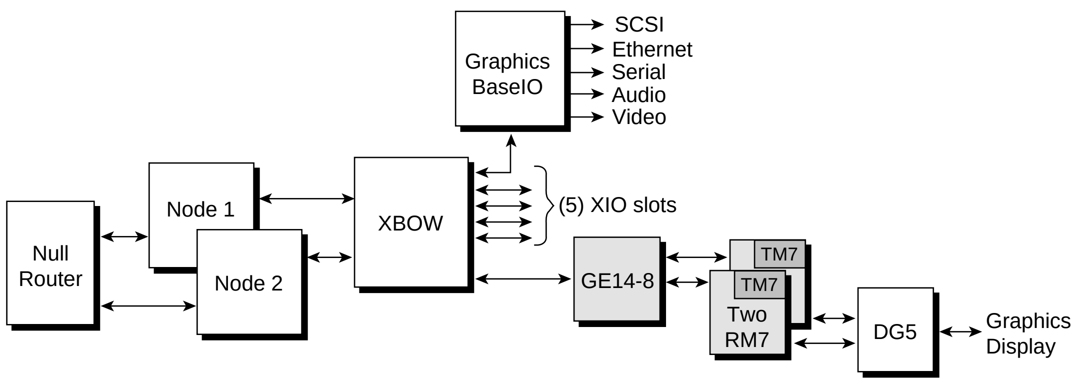

Figure 1-12 Onyx2 InfiniteReality Deskside Block Diagram

1.6.2.2 Onyx2 Rackmount

Figure 1-13 provides an overall block diagram of the Onyx2 rackmount system. The diagram divides the system in two halves, the server side and the graphics side. The server subsystem, located in the bottom half of the Onyx2 rack, is an Origin2000 module with a graphics BaseIO board. The graphics hardware is contained in the top half of the rack. Each subsystem or module has its own midplane, an entry-level System Controller (ELSC), power supply, fan tray, and separate power switches. The modules transmit data using a 1,600 MB bidirectional Crosstown link.

The graphics module contains a Crosstown board that converts GIO to XIO protocol and vice-versa. The Crosstown board provides two sets of ports to interface the InfiniteReality graphics board sets to the Origin2000 server module. The module can contain up to two InfiniteReality2 board sets or pipes. One pipe has a GE14, up to *four* RM7s / TM7s, and a DG5. The second pipe has a GE14, up to *two* RM7s / TM7s, and a DG5.

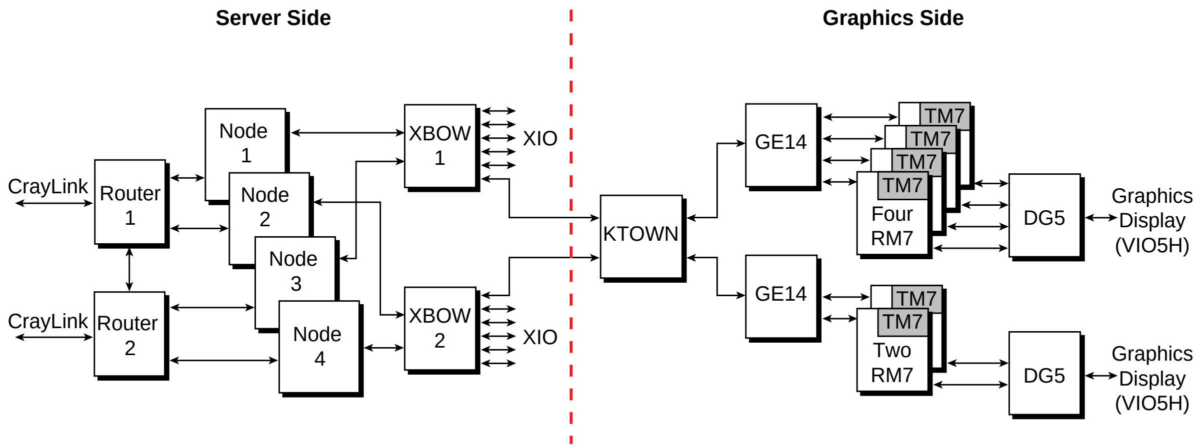

Figure 1-13 Onyx2 Rackmount Block Diagram

Table 1-1 Major Origin2000 and Onyx2 Terms

<table>
<tr><th>Product/Component Name</th><th>Description</th></tr>
<tr><td>Rackmount Modules (Origin2000 Shown)
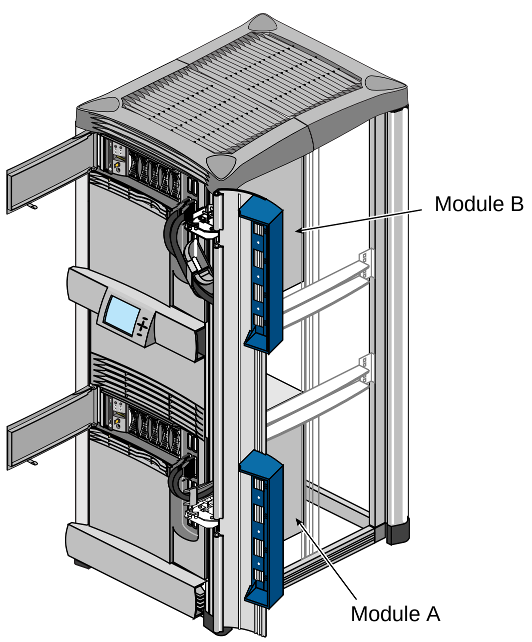
</td><td>The modules in an Origin2000 rackmount system are independent, self-contained subsystems that can house separate sets of hard disks, CPUs, I/O connections and memory.
The Origin2000 modules communicate using the high-speed (800 MBs/sec) CrayLink Interconnect link.
In an Onyx2 rackmount system, the top module is the graphics subsystem and the bottom module is the server subsystem. These modules communicate using a Crosstown cable link.</td></tr>
</table>

Table 1-1 (continued) Major Origin2000 and Onyx2 Terms

<table>
<tr><th>Product/Component Name</th><th>Description</th></tr>
<tr><td>Node Board Location (Origin2000 Deskside Shown)

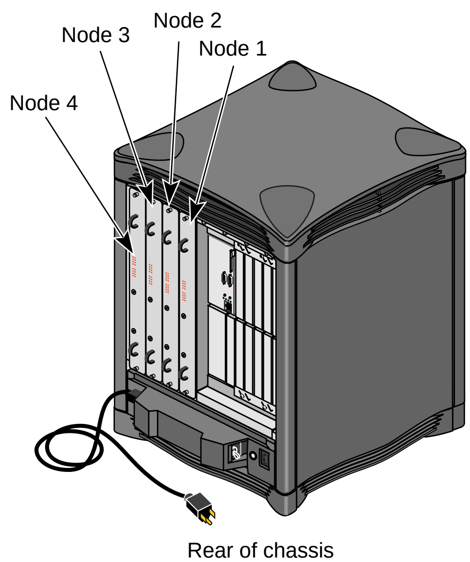
</td><td>The IP27 Node board is the main processing board in the Origin2000 and Onyx2 systems. It contains one or two R10000 processors, the Hub (which provides an interface to the I/O subsystem and the CrayLink Interconnect), and a portion of main (or global) memory, as well as directory memory. Each Node board can support from 64 MBs to 4 GBs of memory. A base system Origin 2000 module can have one, two, or four Node boards. A base system Onyx2 module can have one or two Node boards.</td></tr>
<tr><td>Views of the Node Board

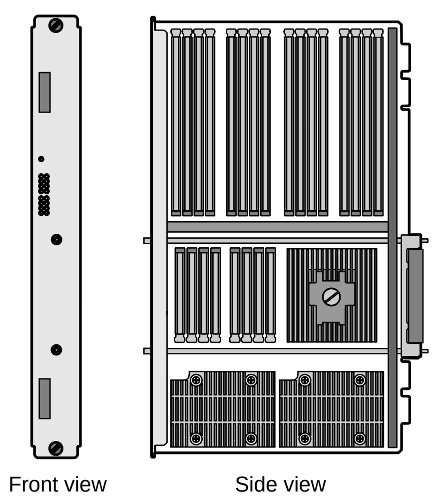
</td><td>Node board (simplified block diagram)

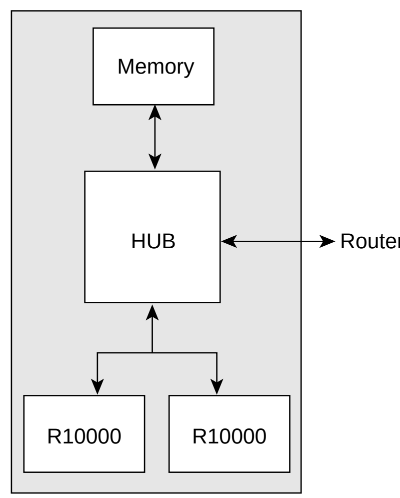
</td></tr>
</table>

Table 1-1 (continued) Major Origin2000 and Onyx2 Terms

<table>
<tr><th>Product/Component Name</th><th>Description</th></tr>
<tr><td>R10000 on Node Board

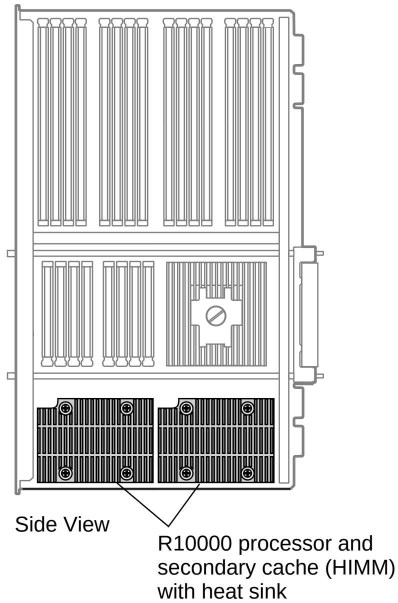

</td><td>The 180-MHz or 195-MHz R10000 CPU can perform at more than 390 million floating point operations per second (MFLOPS). The R10000 has a four-way superscalar architecture that fetches and decodes four instructions per cycle. This processor contains 64 KBs of primary cache and is supported by 1 or 4 MBs of secondary cache.

The R10000 and secondary cache are mounted on a horizontal inline-memory module (HIMM) assembly with heatsink (see below).

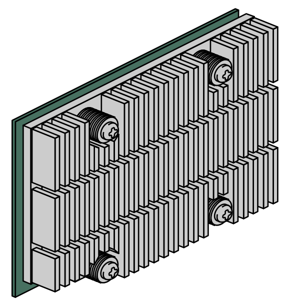
</td></tr>
<tr><td>Hub Location on Node Board

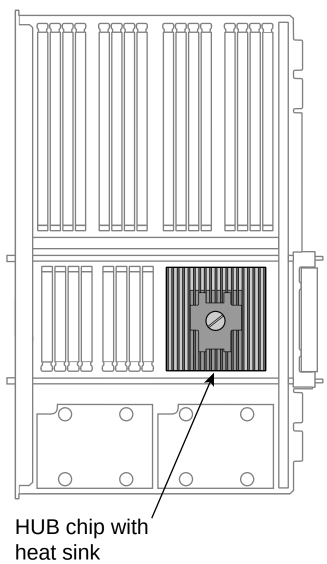

</td><td>The Hub ASIC is the communication infrastructure between the R10000(s), I/O subsystem, main memory, and CrayLink Interconnect (see diagram below). The Hub basically takes the place of the Everest bus interface chips and the CC chips on the MC3 board. The Hub interfaces with directory memory, which is responsible for maintaining cache coherency

The Hub also supports a high-speed 16-bit XIO port (800 MBs/sec input and output bandwidth). The Hub's I/O port can interface the Hub directly to one XIO device or can be used with the Crossbow chip to interface up to seven XIO devices.

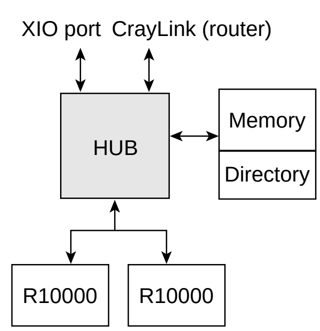

</td></tr>
</table>

Table 1-1 (continued) Major Origin2000 and Onyx2 Terms

Product/Component Name	Description
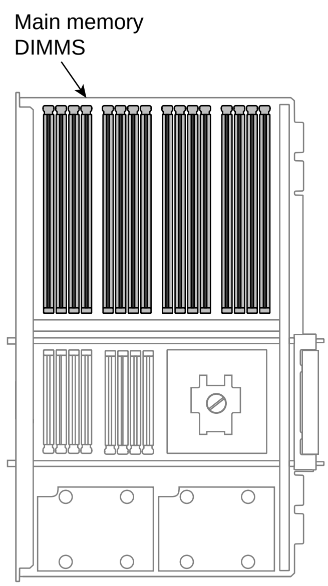	The Origin2000 and Onyx2 main memory uses SDRAMs mounted on dual inline memory modules. A node board can have from 64 MBs to 4 GBs of main memory. Memory upgrades are available in 64-, 128-, and 512- MB increments.
Premium directory memory 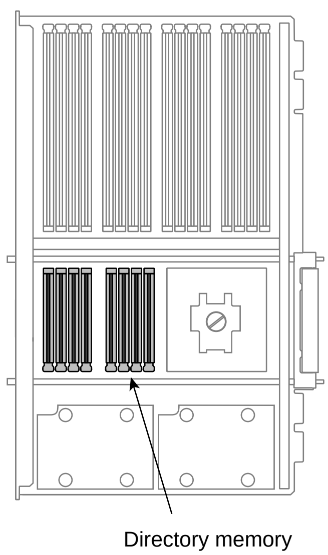	The optional premium directory memory DIMMs are required *only* in configurations with more than 16 Node boards. The directory DIMMs maintain cache coherency in these larger configurations. Note: The premium directory memory is not needed in any Onyx2 configuration.

Table 1-1 (continued) Major Origin2000 and Onyx2 Terms

Product/Component Name	Description
BaseIO Location in Chassis 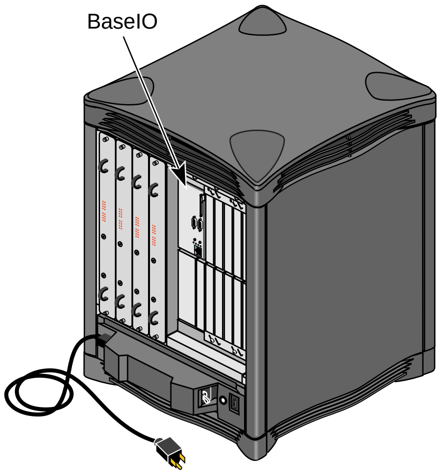	This board provides basic I/O functions for the system such as serial ports, (fast) Ethernet, and single-ended wide SCSI. A dedicated slot (IO1) in the XIO cardcage houses the BaseIO. This board cannot be installed in any of the other XIO expansion slots. (See Chapter 3 for additional information.) An MIO daughter board augments the BaseIO board with video, audio, keyboard, and mouse connections, and a parallel port.
BaseIO board (P/N 030-0734-xxx) 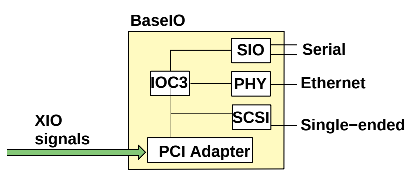	Major ASICS on the BaseIO board include the PCI adapter (or bridge), which converts signals to XIO protocol, and the IOC3, which interfaces with serial and Ethernet devices.
Midplane 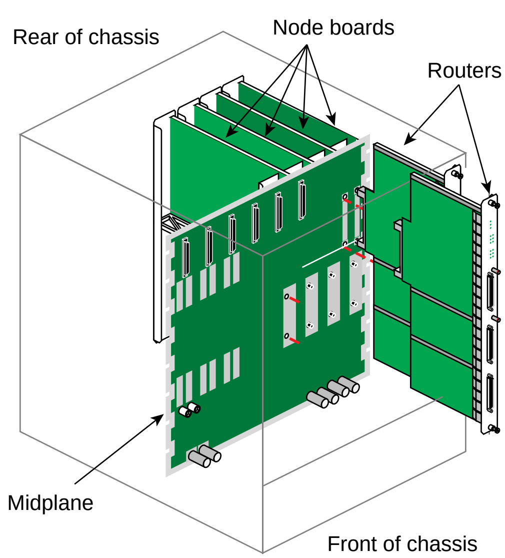	The midplane interfaces with all the major components in the Origin2000 and Onyx2 chassis. The midplane distributes power to all boards, drives, and fans in a module and provides system clock distribution. Connections are made to both the front and the back of this board. Note: There are two types of midplanes in the Onyx2 rackmount configuration—the server subsystem midplane (for system I/O, main memory, and processing) and a graphics midplane.

Table 1-1 (continued) Major Origin2000 and Onyx2 Terms

Product/Component Name	Description
Crossbow (Xbow) on Midplane 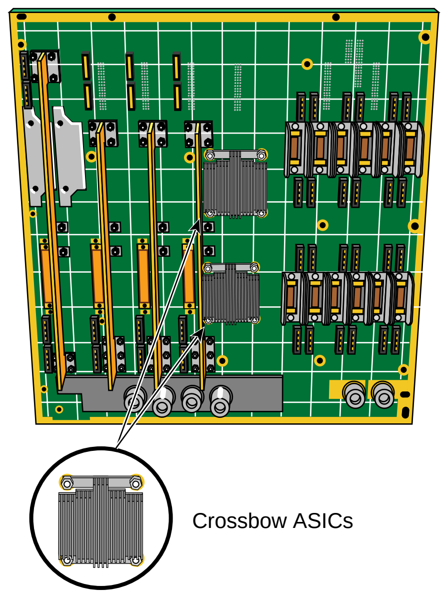 	The Crossbow (or Xbow) ASIC enables XIO devices to efficiently communicate with the host and one another (see diagram below). The Crossbow contains an eight-way crossbar switch that provides point-to-point links between one or more modules and multiple I/O devices. Six of the ports are used for XIO and two connect to Node boards. 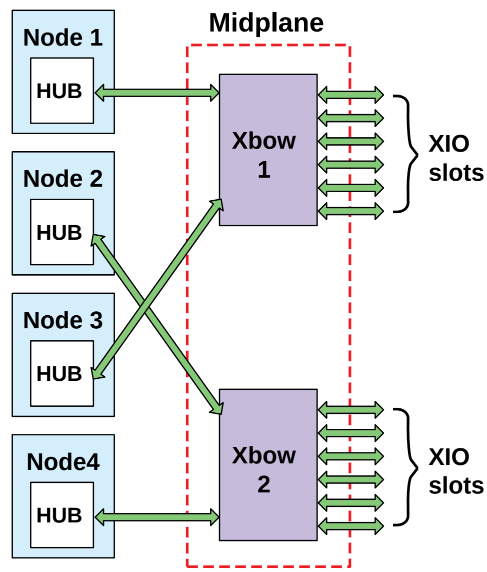
Crossbar Switch 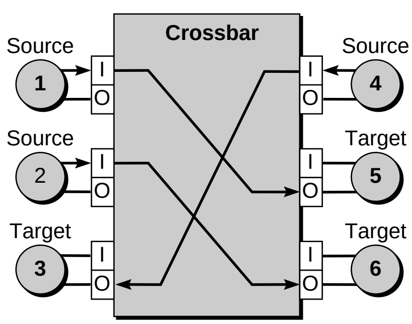 	A number of crossbar switches are used in various ASICs for the Origin2000 and Onyx2 systems (see Section 1.5.2 and Section 2.3 in Chapter 2). The Hub, for example, uses a two-way crossbar; the router chip in a Rack Router board uses a six-way crossbar. The crossbar provides switches between data streams and between ports.

Table 1-1 (continued) Major Origin2000 and Onyx2 Terms

Product/Component Name	Description
System NIC on Midplane 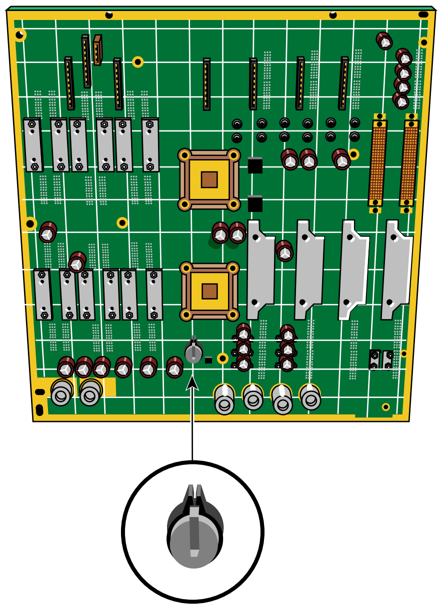	The system number in a can (NIC) contains valuable data such as system ID (serial number) and licensing information. The system NIC resembles a coin-shaped battery and is located in the front portion of the midplane. This NIC can be removed and reinstalled in a replacement midplane. NICs are also used in other boards, such as Node boards, Router boards, and IO6, and other XIO boards. On these boards, the NICs are usually embedded in a tiny chip that is soldered into the board. For more information on NICs, see "NIC Reporting" in Chapter 7.

Table 1-1 (continued) Major Origin2000 and Onyx2 Terms

Product/Component Name	Description
Routers 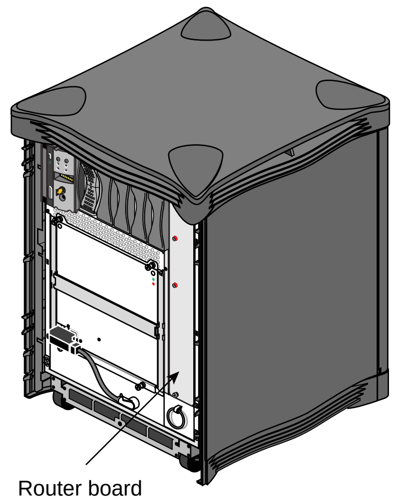	The routers enable communication between two or more Node boards. There are three router boards types and a router chassis. Null Router board Star Router board Rack Router board (pictured to the bottom left) Cray Router chassis The Null and Star Routers cannot be used for external connections. For additional information on Router boards, see Chapter 3, "Chassis Tour." On a deskside system, the router board pictured below is located behind a panel as shown in the upper left figure.
Rack Router board 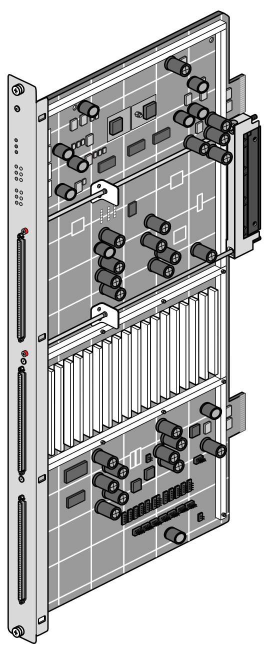	

Table 1-1 (continued) Major Origin2000 and Onyx2 Terms

Product/Component Name	Description
CrayLink Interconnect cable 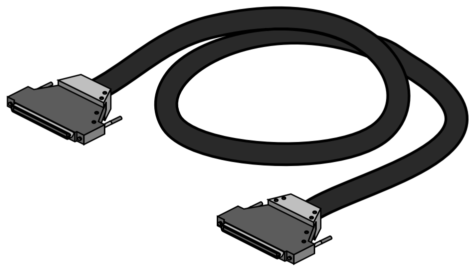	This is the cabled interface that runs between Rack Router boards. The Rack Router boards connect to the Node boards within the same module or to other Router boards in another module. The CrayLink Interconnect provides a high-speed, scalable interconnection between modules within the system. Its topology minimizes average hop counts and provides physical link redundancy (so that if one link fails, another link can take its place).
Router Chip 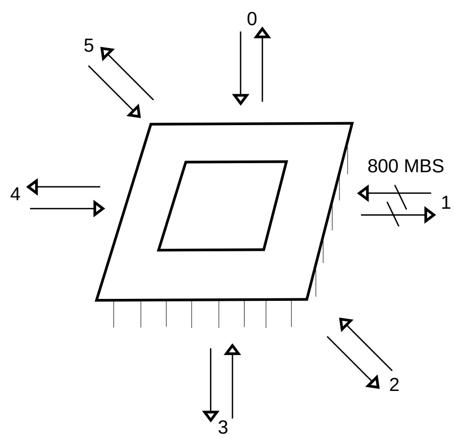	This is the primary active component of the CrayLink Interconnect. The router chip is a bidirectional packet-switched controller that can transport up to 4.8 GBs / sec through one chip (or up to 800 MBs / sec peak per port). Each module chassis can have one or two routers. The router interfaces the two hubs on two Node boards. The Router board contains a router chip that has six ports. Two of the ports connect to the Hubs; a third port is for internal router-to-router connections. The fourth, fifth, and sixth ports connect to external routers on other rackmount modules.
XIO Cardcage (Rear of Chassis)	The XIO cardcage allows you to install up to 6 full-height or 12 half-height XIO boards for networking and peripheral support into the Origin2000 systems or the Onyx2 rackmount chassis. The Onyx2 deskside chassis can accommodate up to three full-height or five half-height XIO boards.

Table 1-1 (continued) Major Origin2000 and Onyx2 Terms

Product/Component Name	Description
Compression Connector 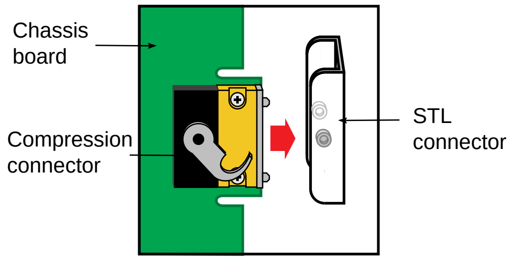	The compression connector mechanically latches to the sensitive STL (Silicon Graphics transistor logic) connector located on the midplane. The compression connector uses a cam mechanism to cushion the Origin2000 and Onyx2 board connections to the midplane gently. This mechanism is generally good for 100 connection attachments.
Single Connector Assembly (SCA) drives	The Origin2000 and Onyx2 systems use SCA drives to replace the bulkier and multiple-connector drives from the CHALLENGE and Onyx systems. The SCAs combine the SCSI interface plus SCSI ID and power pins all in one connector.
Power Distribution Unit (PDU)—Racks only	The PDU is the central power source that accepts 220-VAC input. It also provides a common power and grounding source for the rack to support CrayLink Interconnection between modules.
Peripheral Connector Interface (PCI) Adapter	The optional PCI adapter enables installation of up to three PCI boards into the Origin2000 base module.
InfiniteReality graphics in Onyx2 Rack 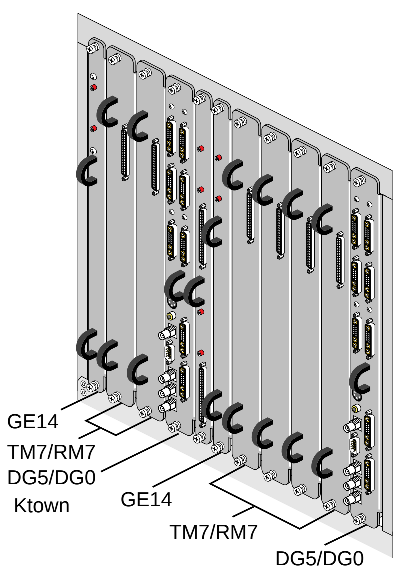	This is a modification of InfiniteReality for Onyx. The illustration on the left shows two sets of boards (called graphic pipes) in the Onyx2 rack system. The GE has been engineered to talk to a Xbow ASIC instead of an FCI (flat-cable interface) as in the Onyx products.

Table 1-1 (continued) Major Origin2000 and Onyx2 Terms

Product/Component Name	Description
Crosstown 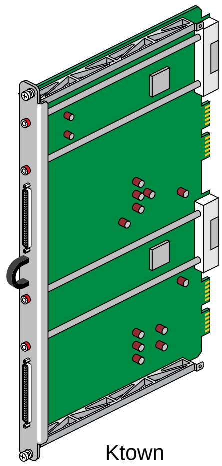	The Crosstown board converts Crosstown data from the server module to GIO bus signals for processing by the GE14.

1.7 Related Product Line

Table 1-2 and Figure 1-14 describe the entire related product line including the graphics configurations and peripheral chassis systems.

Table 1-2 Related Product Line and CrayLink Interconnect Capability

Product Name	Description	CrayLink Interconnect Capability?
Origin2000 Deskside	This is the base module. The Origin2000 deskside contains one to eight R10000 processors on up to four Node boards (one or two processors per Node) and 12 half-size or 6 full-size XIO slots. In addition, it contains a BaseIO board, a five -disk drive bay, and system controller. The Origin2000 deskside module has a 19-inch rackmountable form factor (after the skins are removed) and can be installed in standard 19-inch racks.	No (see Note below)
Origin2000 Rackmount	This is the Silicon Graphics 19-inch rackmount version. It contains up to two Origin2000 modules, a multimodule System Controller (MMSC), and CrayLink Interconnect cable management hardware.	Yes
Onyx2 Deskside	This system has InfiniteReality graphics and up to two Node boards.	No

Table 1-2 (continued) Related Product Line and CrayLink Interconnect Capability

Product Name	Description	CrayLink Interconnect Capability?
Onyx2 Rackmount	This is the graphics rackmount system. The Onyx2 rackmount provides two InfiniteReality graphics pipes and an Origin2000 module.	Yes
Origin200	This product provides a low-end entry into the Origin™ family product line. Two Origin200 systems can be connected together using CrayLink Interconnect cables, but not to other system types.	Yes (See comments in adjoining column to the left.)
Origin Vault	This “Vault L-size” peripheral housing module can accommodate up to six 3.5”and two 5-1/4”(half height) drives that can be either Single-ended or differential SCSI depending on the chassis type..	No

Note: CrayLink Interconnect capability for the Origin2000 and Onyx2 deskside modules is not supported. However, it is possible to install a deskside module into a rack to provide CrayLink Interconnect capability (see Chapter 9, “Installing a Deskside or Graphics Module Into a Rack” for more information).

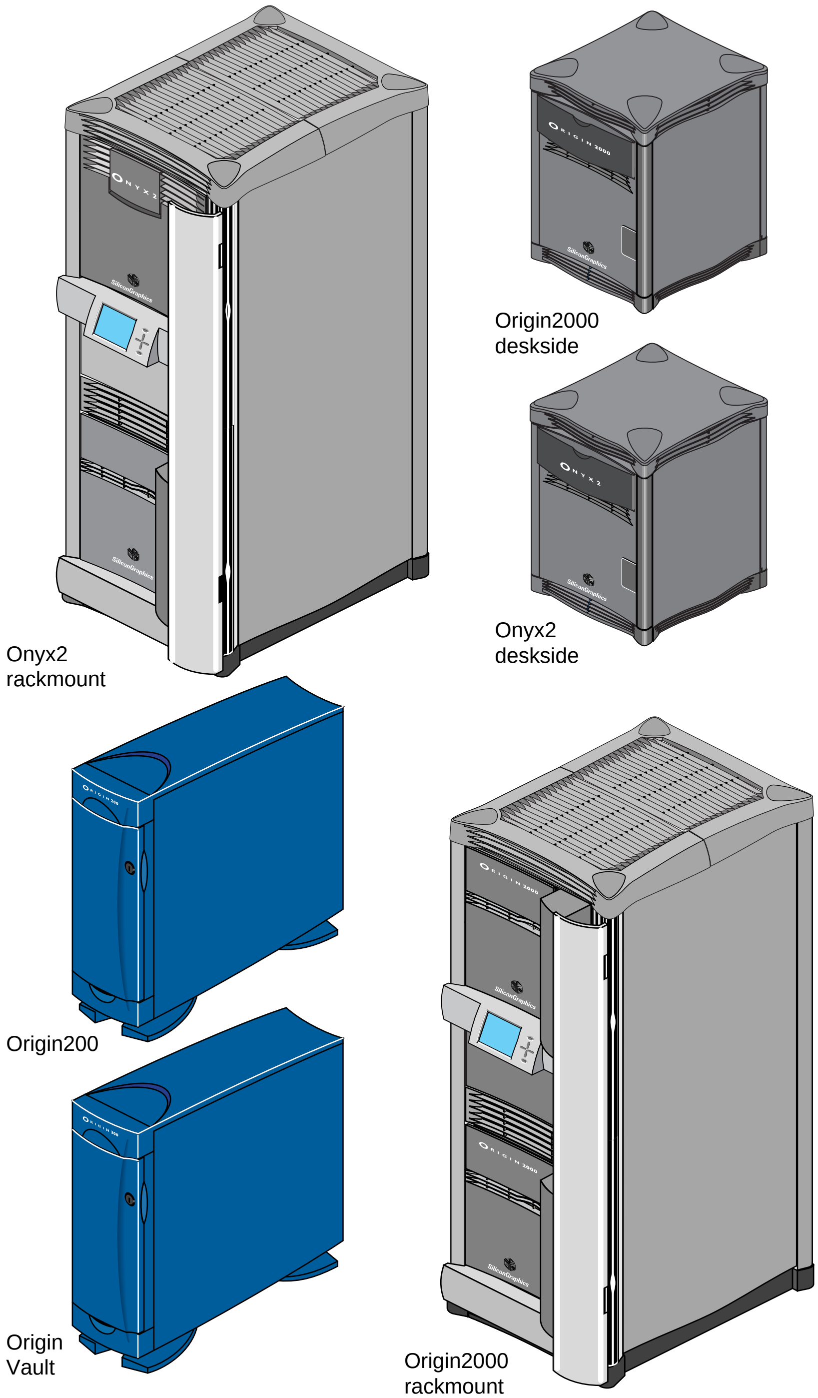

Figure 1-14 Related Product Line

Chapter 2

Theory of Operations

This chapter provides a detailed description of the Origin2000 and the Onyx2 board sets, major ASICs, and the XIO subsystem. See Chapter 1, “Product Overview,” for an Overall Block Diagram of the Origin2000 and Onyx2.

2.1 System Board Set

The base Origin2000 and Onyx2 system configurations generally include these common boards:

- IP27 Node board
- Router board
- BaseIO board
- Midplane

In addition to the system board set, the Onyx2 system ships with either the Reality or the InfiniteReality graphics board set, depending on the configuration.

2.1.1 Node Board

The basic building block of an Origin2000 or an Onyx2 system is the IP27 Node board, which plugs into the rear cardcage of a deskside or rackmount enclosure. The Node board contains the Hub ASIC with interfaces to the processor(s), memory, I/O, and CrayLink Interconnect. Figure 2-1 shows a block diagram of the Node board, with its central Hub ASIC.

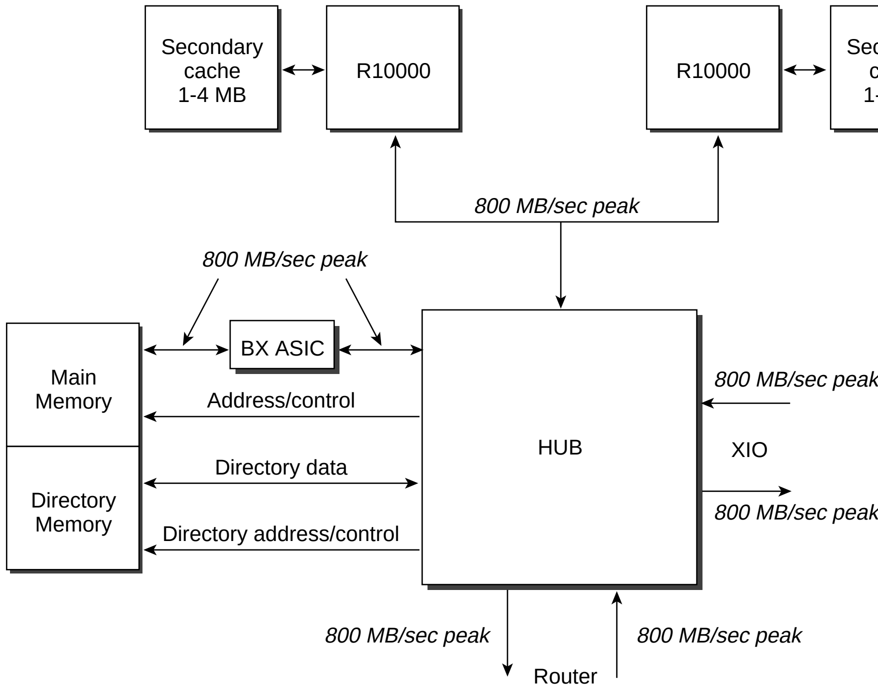

Figure 2-1 Block Diagram of the Node Board

2.1.1.1 Hub Interfaces

The Node board has a central Hub ASIC, which can be connected to

- either one or two processors
- either main memory and its associated directory memory or the flash PROM
- the system interconnection fabric, through a dedicated Router port
- I/O crossbar interconnect, through a dedicated XIO port (single port) or Crossbow ASIC (eight ports)

2.1.1.2 Processors

Each Node board is capable of supporting either one or two R10000 processors. Each processor is mounted on a horizontal inline memory module (HIMM) together with its primary cache, and either one or four MBs of secondary cache, as shown in Figure 2-2.

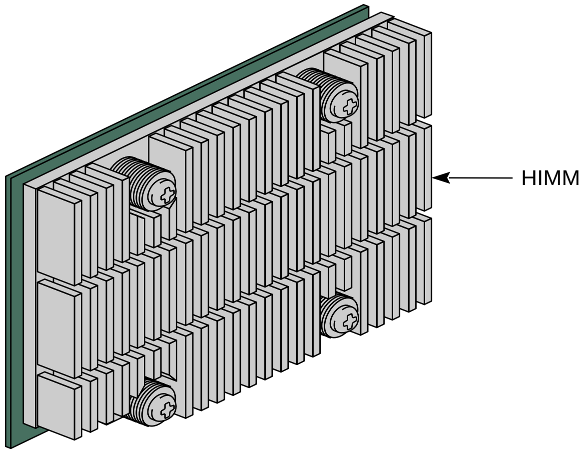

Figure 2-2 Horizontal Inline Memory Module

2.1.1.3 Memory

Main and directory memory are implemented using synchronous DRAM (SDRAM) parts mounted on dual inline memory modules (DIMMs). Each Node board has 8 banks of main memory. Each bank contains two DIMMs. For configurations with up to 16 Node boards, directory memory is included in the main memory DIMMs; for configurations with more than 16 Node boards, premium directory memory must be added in separate slots (see Chapter 8, "Memory Upgrades" for more information).

2.1.1.4 Distributed Shared Memory

The Origin2000 and Onyx2 systems use distributed shared memory (DSM). With DSM, main memory is partitioned among processors but is accessible to and shared by all the processors. The Origin2000 and Onyx2 systems divide main memory into two classes: local and remote. Memory on the same Node board as the processor is labeled *local*, with all other memory in the system labeled *remote*.

To a processor, main memory appears as a single addressable space containing many blocks, or pages. Each Node board is allotted a static portion of the address space—which means there is a gap if a Node board is removed. Figure 2-3 shows an address space in which each Node board is allocated 4 GBs of address space and Node 2 is removed, leaving a hole from address space 4G to 8G.

2.1.1.5 Flash PROM

The 1-MB flash PROM contains the nonvolatile system log and NVRAM. This PROM also supplies the code to set up or initialize the CPUs, memory, console I/O port, and CrayLink Interconnect. In addition, the IP27 flash PROM provides the command monitor and POD.

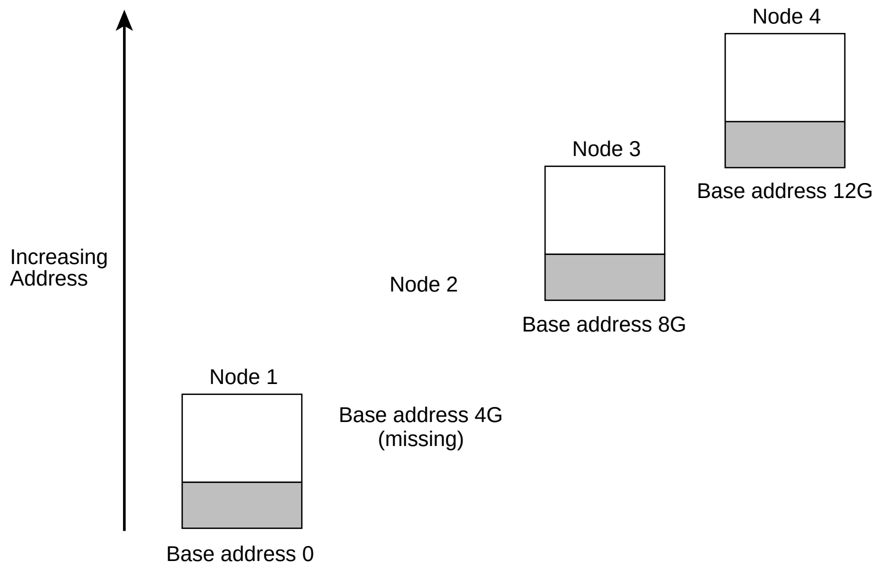

Figure 2-3 Origin2000 and Onyx2 Address Space

2.1.2 Router Board

Router boards physically link the Hub ASIC on the Node board to the CrayLink Interconnect.The CrayLink Interconnect provides a high-bandwidth, low-latency connection between all the Node boards. Central on a Router board is the Router ASIC, which implements a full six-way nonblocking crossbar switch.

The location of the Router board is shown in Figure 2-4.

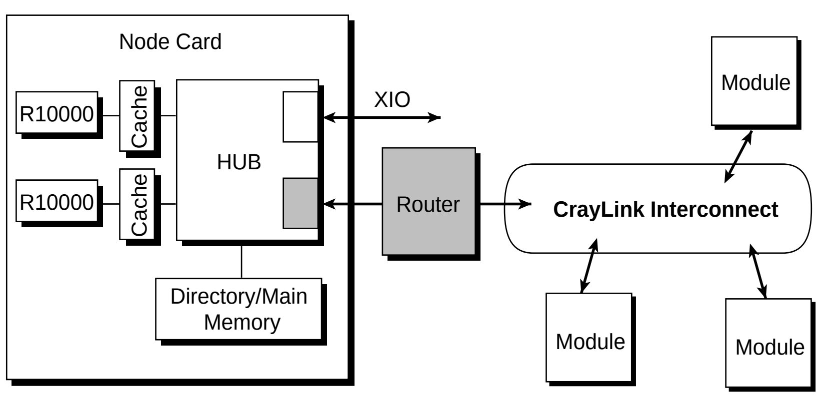

Figure 2-4 Location of a Router Board in Origin2000 and Onyx2 Systems

The Router crossbar allows all the Router ports to operate simultaneously at full duplex; each port consists of two unidirectional data paths. The Router board also includes a set of protocols that provides a reliable exchange of data even in the face of transient errors on links, manages flow control, and prioritizes data so that older data is given precedence over newer data.

2.1.2.1 Types of Router Boards

There are four types of routers. Three are boards and the forth is a chassis.

- The *Null Router* connects two Node boards in the same deskside enclosure (up to four R10000 processors). The Null router configuration cannot be expanded.
- The *Star Router* connects up to four Node boards. Three- and four-Node configurations are created using a Rack router with a Star router.
- The *Rack Router* cables together 2 to 32 Node boards (from 2 to 64 CPUs), physically located in from one to eight enclosures.
- The *Cray Router* chassis is used in conjunction with Rack routers to expand the system from 33 to 64 Node boards (65 to 128 CPUs).

Programmable tables within the Router ASICs control packet routing through the CrayLink Interconnect. These tables allow for partial configurations (system sizes that are not 2^n) and reconfigurations around broken links or inoperative modules.

2.1.2.2 Xpress Links

Additional CrayLink Interconnect cables, called Xpress links, can be installed between Standard Router ports for increased performance (reduced latency and increased bandwidth). Each Xpress link provides additional bandwidth of 800 MBs/sec each way (1.6 GBs bidirectionally).

2.1.3 BaseIO Board

The BaseIO board (also called the IO6 board) contains I/O functions that are standard in each base system enclosure. There are two types of BaseIO boards—one for the Origin2000 server system and another for the Onyx2 graphics system.

2.1.3.1 Origin2000 BaseIO Board

The BaseIO board for the server system provides these functions:

- one 10/100-Base-TX fast Ethernet link, with autonegotiation (compliant with 802.3u)
- two 460-Kbaud serial ports, composed of dual, independent UARTS
- one external, single-ended wide SCSI port (compliant with X3.131-1994)
- one internal fast-20 SCSI port (compliant with X3.131-1994)
- one real-time interrupt output, for frame sync
- one interrupt input
- flash PROM
- NVRAM
- time-of-day clock

The logical location of a BaseIO board, connected to a Crossbow ASIC, is illustrated in Figure 2-5.

2.1.3.2 Onyx2 BaseIO Board

The graphics BaseIO board (sometimes called the IO6G) is augmented by the Media I/O (MIO) daughterboard. The MIO is primarily used in graphics systems to provide additional audio, serial, keyboard/mouse, and parallel ports. Specifically, the MIO board adds the following to a BaseIO board:

- one IEEE 1284 parallel port
- four 460-Kbaud serial ports, dual independent UARTS
- one audio analog stereo input port
- one audio analog stereo output port
- one audio AES3/AES11/SPDIF digital input port
- one audio AES3/AES11/SPDIF digital output port
- one audio Alesis ADAT/Optical SPDIF digital fiber input port
- one audio Alesis ADAT/Optical SPDIF digital fiber output port
- two keyboard ports
- two mouse ports

The majority of audio functions are provided by the RAD ASIC.

Due to its size, there can be only one BaseIO board per deskside or module enclosure. Because an MIO board is mounted as a daughterboard on the IO6G board, there can be only one MIO board in a deskside or module enclosure. If more than one MIO is desired, it must be added to a multimodule system in which additional BaseIO boards can be installed as well.

The BaseIO and MIO are treated as an atomic module when MIO functions are used. A block diagram of the MIO board is shown in Figure 2-6.

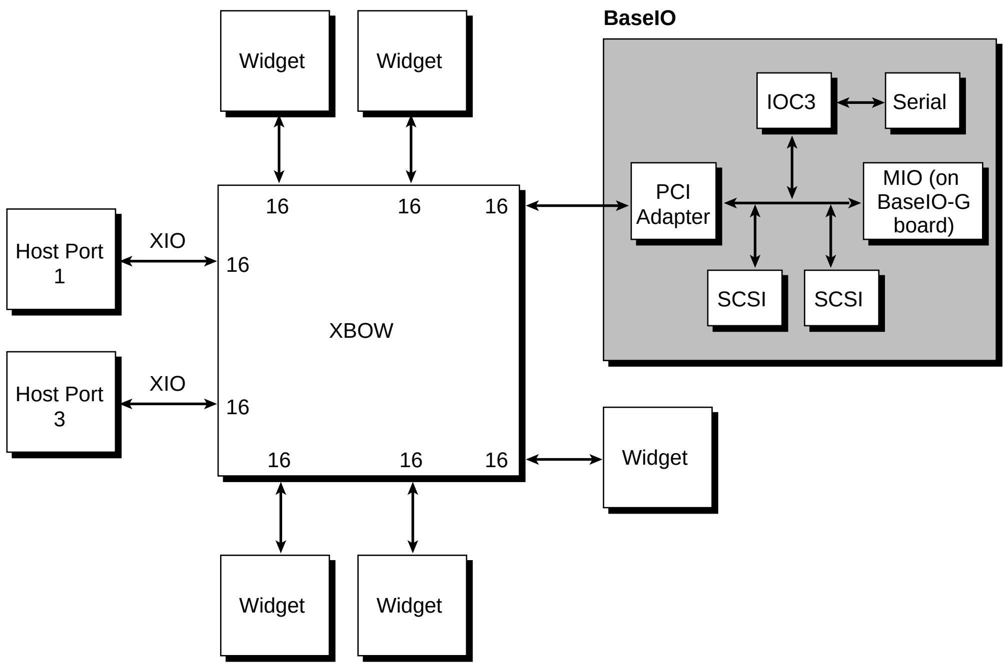

Figure 2-5 Logical Location of an BaseIO Board in an Origin2000 and Onyx2 Systems

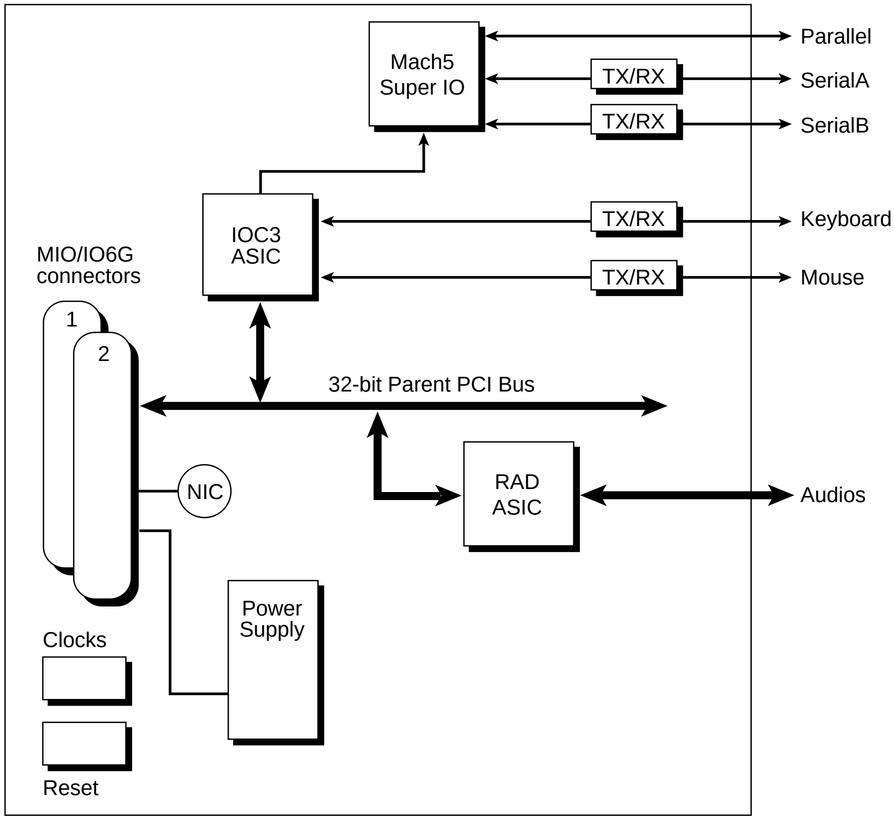

Figure 2-6 MIO Board Block Diagram

2.1.4 Midplane

Functionally, the Origin2000 and Onyx2 midplanes provide the following:

- a standard system clock for both the XIO and the CrayLink Interconnection
- STL links for the CrayLink Interconnection and XIO links within the module
- power distribution throughout the system
- system control signals and additional real-time clock distribution
- digital media sync
- SCSI connections

Physically, the components located on the midplane are

- two or four 300-pin STL connectors for Node boards (two in the Onyx2 system and four in the Origin2000 system)
- one or two 300-pin connectors for Router boards (one in the Onyx2 system and two in the Origin2000 system)
- 96-pin connectors for XIO boards
- five connectors for wide, single-ended ultra-SCSI disk drives
- one connection for a single-ended SCSI CD-ROM
- one connection for a System Controller
- two system clocks:
  – Systems containing 180 MHz node boards have 90 MHz hub chips, and should have the 360 MHz clock oscillator on the midplane activated
  – Systems containing 195 MHz node boards have 97.5 MHz hub chips, and should have the 390 MHz clock oscillator on the midplane activated

2.2 XIO Subsystem

The XIO subsystem or cardcage, which is located in the rear of the chassis, provides 800-MB (peak transfer rate) performance for XIO boards. The XIO subsystem is implemented primarily through one or two eight-ported Crossbow ASICs (see Section 2.3.3).

2.2.1 XIO Protocol

The Origin2000 and Onyx2 systems use an advanced input-output (I/O) subsystem, consisting of a number of high-speed XIO links. XIO supports a wide range of Silicon Graphics and third-party I/O devices.

2.2.1.1 Distributed I/O

XIO is distributed, with an I/O port on each Node board. As with the Origin2000 and Onyx2 systems' distributed memory, each I/O port is accessible by every CPU. I/O is controlled either through the single-port XIO protocol link on the Node board or through an intelligent crossbar interconnect on the Crossbow (XBOW) ASIC.

2.2.1.2 Crossbow Expansion

A Crossbow ASIC expands the single XIO port to a total of eight ports: six are used for I/O and two are connected to Node boards. Ports using the XIO protocol can be programmed for either 8- or 16- bit communication. The electrical interface for XIO is the same as that used by the CrayLink Interconnect

2.2.1.3 XIO Devices — Widgets

The form factor for XIO widgets* may vary. Typically a widget is a single board, either half-size (10 x 6.5 x 1 inch) or full size (10 x 13 x 1 inch); however, it is possible for a widget to include a daughterboard.

XIO can also run outside an enclosure using the Crosstown protocol, which is described in Section 2.2.2, “Crosstown (XTOWN) Board.”

2.2.2 Crosstown (XTOWN) Board

The XIO protocol can also run outside an enclosure using the Crosstown (XTOWN) conversion board. For example, an XIO slot may be occupied by a Crosstown board, which contains an STL-to-3.45V PECL converter and a Crosstown cable attachment; the result is to convert the XIO STL link to differential signal levels. This connection can support XIO devices up to three meters away. Crosstown is primarily used to support graphics configurations.

2.3 Origin2000 and Onyx2 ASICs

Within an Origin2000 and Onyx2 systems, each ASIC is responsible for a critical element of information transfer. The ASICs and their locations in a system are

- Hub chip located on each Node board
- Router chip located on each Router board in the CrayLink Interconnect
- Crossbow chip located on the midplane as part of the XIO interconnect
- Bridge, LINC, and IOC3 chips located on XIO boards linked to the I/O interface

Note: Graphics-specific ASICs are discussed in Section 2.4, “Graphics Subsystem.”

Figure 2-7 shows the interconnections between these ASICs and the communication protocols that run on the interconnections from the Hub ASIC.

---

* I/O devices are pseudonymously referred to as “widgets.”

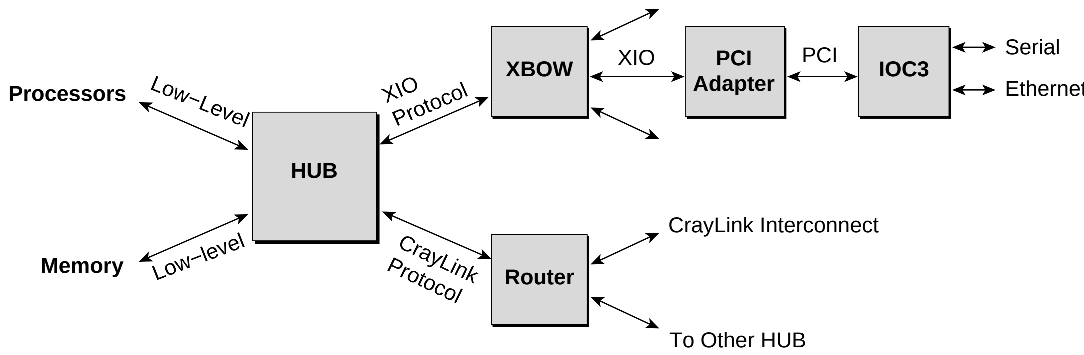

Figure 2-7 ASIC Protocols

The Hub ASIC is the core of the Node board. It interconnects to processors, memory, the CrayLink Interconnect, and the XIO interconnect.

- From its processor interface, the Hub chip supports a low-level bus protocol similar to the shared-bus protocol on the CHALLENGE Ebus.
- To the interconnection fabric, the Hub chip supports a high-level messaging protocol called the CrayLink Interconnect, which implements the cache-coherent distributed shared memory.
- To the I/O interconnect, the Hub chip supports a separate high-level protocol called XIO, optimized for I/O devices that are not using the cache-coherence protocol.
- A memory/directory protocol interfaces with the high-speed, interleaved memory system. It supports directory memory and page migration counters.

Figure 2-7 also shows Router chip and Crossbow chip interconnections. Note that the Crossbow chip uses XIO protocol at both its input ports and its output ports. Similarly, the Router chip uses the CrayLink Interconnect protocol at both its input and its output ports.

Finally, the figure shows how a Bridge chip interconnects with the Crossbow and a communications controller on an XIO board; in this case, the controller is the IOC3 chip. The bridge takes Crossbow as input and interfaces it with the PCI protocol. The controller supported by the bridge (the IOC3) takes PCI as input and interfaces it to Ethernet and serial.

2.3.1 Hub ASIC

The Origin2000 and Onyx2 system architecture defines a distributed shared memory multiprocessor using 1 to 1,024 processors.* The Origin2000 and Onyx2 systems provide a single address space with cache coherence applied across the entire system.

The Origin2000 and Onyx2 systems are organized into a number of Node boards; each Node board contains up to two processors, a portion of the global memory, a directory to maintain cache coherence, an interface to the I/O subsystem, an interface to the other Node boards on the system, and a Hub ASIC, which links all these subsystems through a crossbar.

2.3.1.1 Hub Interfaces

The Hub ASIC can be viewed as the core of the Node board. Physically, it is located on the Node board, and is responsible for connecting all four interfaces of the Node board together:

- Each Node board can have up to two R10000 processors linked to the Hub.
- A portion of the distributed, shared main memory is connected to the Hub, together with directory memory used for cache coherence and page migration counts. This memory can range in size up to 4 GB.
- Either one or two Hubs can be connected to each Crossbow ASIC.
- Either one or two hubs can be connected to a Router board, which in turn links the Hub to other Node boards connected to the system wide-interconnection fabric.

2.3.1.2 Communication Control

The four interfaces listed in Section 2.3.1 are interconnected by an internal crossbar. The interfaces on the Hub communicate by sending messages through the crossbar.

The Hub controls *intranode* communications between the Node board's subsystems, as well as *internode* communications with other Hub ASICs on other Node boards. The Hub converts internal messages, using a request/reply format, to and from the external message format used by the XIO or CrayLink Interconnect port. All internal messages are initiated by processors and I/O devices.

---

* The initial release of the product supports 128 CPUs.

2.3.2 Router ASIC

The Router ASIC is a six-port dynamic switch that forms the interconnection fabric connecting the Node boards. Physically, the Router chip is located on a Router board, which plugs into the deskside midplane opposite the XIO and Node boards. A block diagram of the Router ASIC is given in Figure 2-8.

Functionally, the Router ASIC does the following:

- Determines the most efficient connection of receive-to-send ports, given the set of received messages, and dynamically switches connections between any of the six pairs of ports through the six-way crossbar
- Communicates reliably by using the CrayLink Interconnect link-level protocol (LLP) to other routers and hubs
- May route different messages to the same destination through different paths, for greatest speed and efficiency (adaptive routing)
- To reduce latency, routes messages without having to receive the entire message (wormhole routing)
- Buffers CrayLink Interconnect messages
- Communicates with a total peak bandwidth of 9.6 GB/sec

Figure 2-8 shows the source-synchronous drivers (SSD) and receivers (SSR) that multiplex and demultiplex the high-speed external connection (400 MHz) to the router internal frequency (100 MHz).

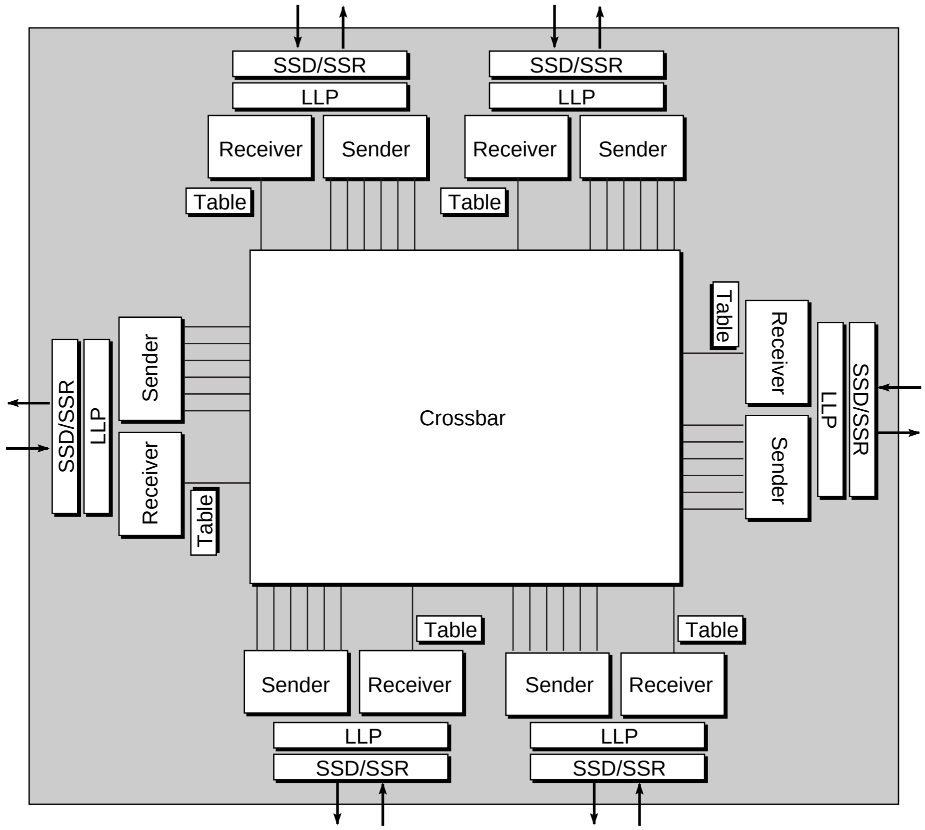

Figure 2-8 Block Diagram of the Router ASIC

Three of the Router ports are STL, for rapid communication between ASICs inside the enclosure. The other three Router ports are differential PECL, for communication between modules over cables outside the enclosure. The three internal router ports use single-ended STL signaling to minimize pin count. The three external router ports use differential PECL, which provides better noise immunity on external links.

2.3.3 Crossbow ASIC

The Crossbow ASIC has a dynamic crossbar switch that expands the dual-host XIO port to six 16-bit I/O ports. Each I/O port can run in either 8- or 16-bit mode, with rate-matching buffers to decouple 16-bit to 8-bit ports. At least one Crossbow port must connect, and a maximum of two Crossbow ports can connect to a host.

The Crossbow ASIC uses the XIO protocol at all its ports. There are two Crossbow ASICs on each deskside midplane.

A functional view of a Crossbow ASIC, with dual hosts and six half-size XIO boards, is shown in Figure 2-9. (CrayLink Interconnect cabling is not shown in this illustration.)

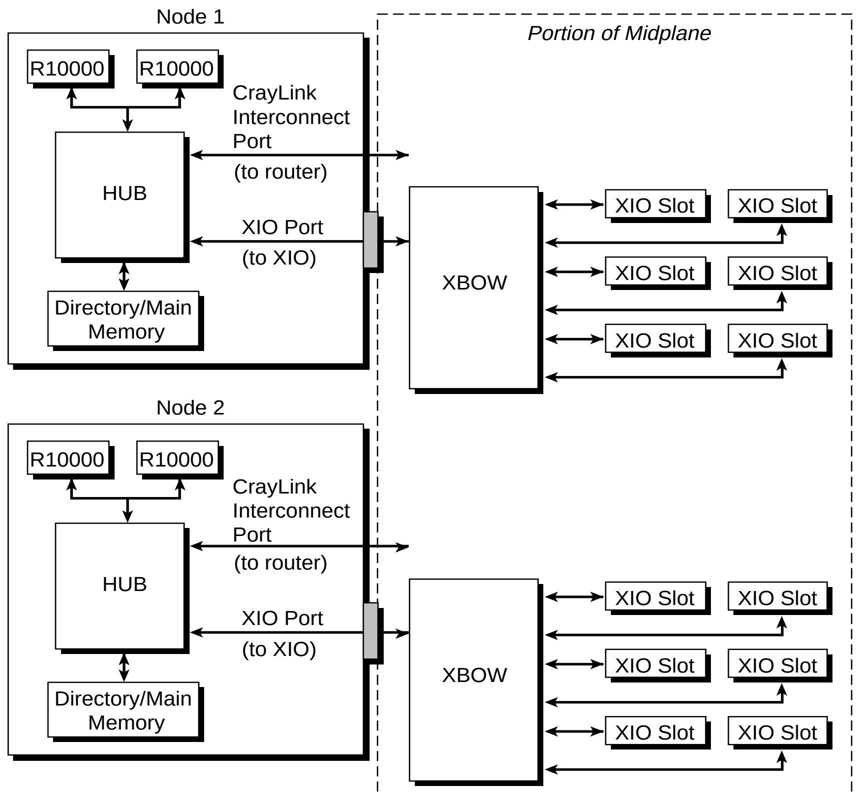

Figure 2-9 Functional Location of Crossbow ASIC

In this dual-host configuration, the six remaining XIO ports are statically partitioned between the two Node boards; if one of the host Node boards is inoperable, the second host Node board can be programmed to take control of all the XIO ports. Note that the Node boards connect to two ports of the Crossbow (1 and 3) and that the remaining six ports connect to the XIO widgets.

As described earlier, a *widget* is a generic term for any device connected to an XIO port. The Crossbow ASIC contains a crossbar switch that dynamically connects individual ports to particular I/O widgets (host, graphics board, serial I/O), as shown in Figure 2-10. The Crossbow decodes fields in XIO messages to determine control and destination information.

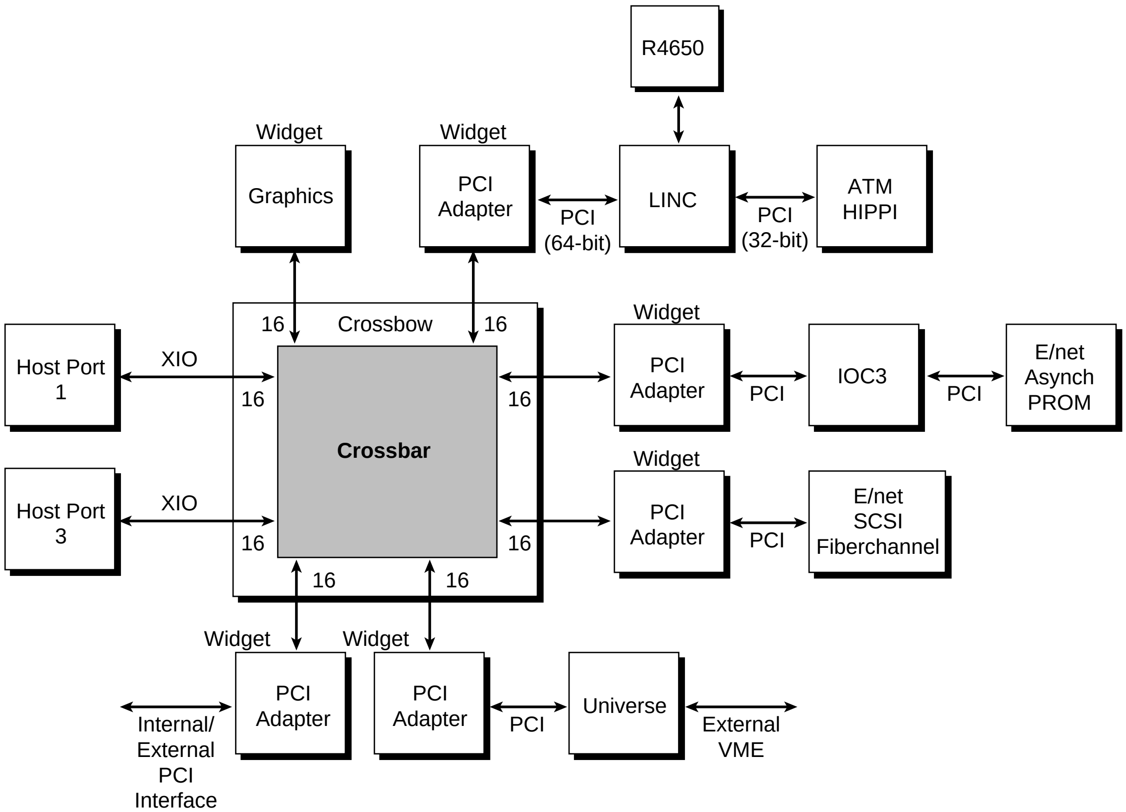

Figure 2-10 Block Diagram of a Crossbow ASIC, Showing Eight Ports Connected to Widgets

2.3.4 Bridge ASIC

The Bridge ASIC is physically located on an XIO board. It converts the XIO link to the PCI bus protocol. The Bridge ASIC also provides address mapping, interrupt control, read prefetching, and write gathering.

Peak bandwidth of the Bridge ASIC is 800 MB/sec on an XIO link and 266 MB/sec on the PCI link. There is a Bridge ASIC on every I/O widget board.

An illustration of an I/O subsystem with a Bridge ASIC is given in Figure 2-11.

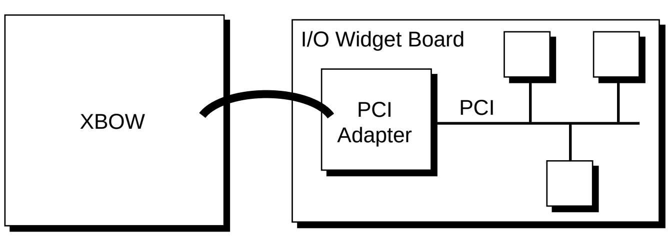

Figure 2-11 Bridge ASIC

2.3.5 IOC3 ASIC

The IOC3 ASIC takes PCI output from the Bridge ASIC and converts it to standard I/O protocols for Ethernet, parallel I/O, and serial I/O. The IOC3 ASIC is physically located on the BaseIO and Media IO (MIO) boards.

2.3.6 LINC ASIC

The LINC ASIC is designed to support intelligent controllers and optimize throughput with a variety of scatter and gather functions.

The LINC ASIC is located on the HIPPI serial and ATM boards; the ASIC converts the 64-bit PCI protocol to 32-bit PCI protocol. Using an IDT R4650 MIPS processor running at 132 MHz, the LINC has a 64-bit host side “parent” PCI bus (PPCI), which provides

- 64-bit addressing that is used for system DMA access
- 32-bit addressing that is used for peer-to-peer DMA and PIO access

The LINC ASIC also supports a 32-bit “child” PCI bus (CPCI) for attaching interface devices; the CPCI provides request/grant and interrupt support for up to two devices.

A block diagram of an I/O subsystem with both a Bridge ASIC and two LINC ASICs is given in Figure 2-12.

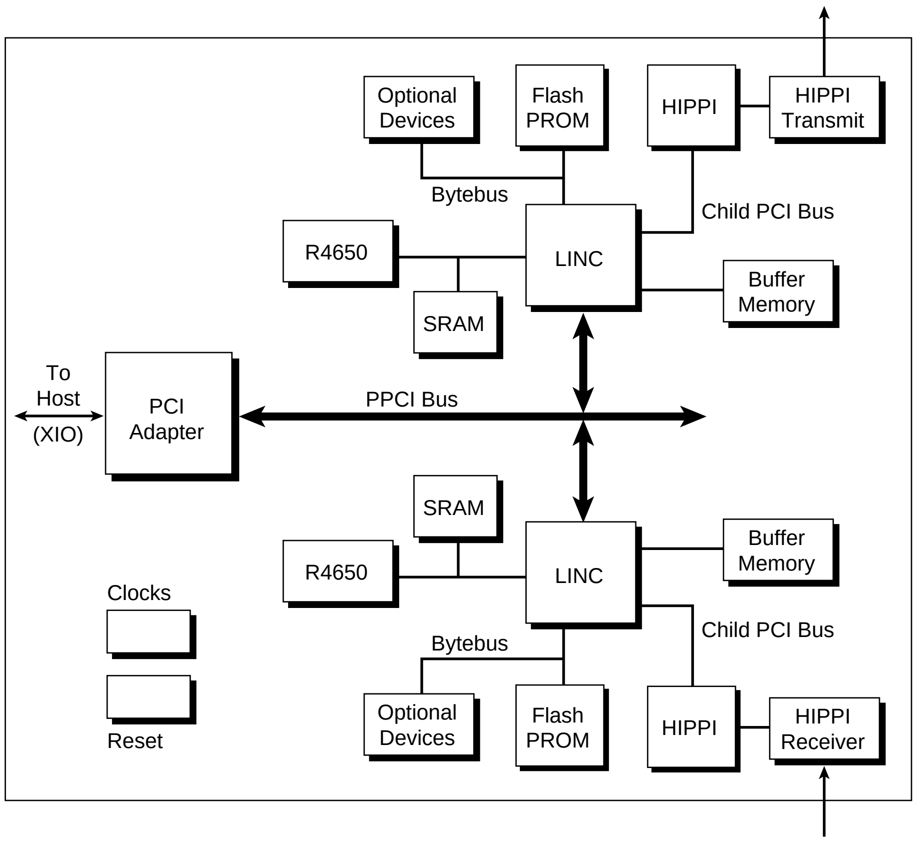

Figure 2-12 Block Diagram of LINC ASICs With Bridge ASIC

2.4 Graphics Subsystem

The Reality and InfiniteReality visualization subsystem consists of three boards: the Geometry Engine® (GE), the Raster Manager (RM), and the Display Generator (DG). The Reality board set is the entry-level graphics for the Onyx2 and installs into deskside systems *only*. The Reality graphics consists of the GE14-4, RM-8, and DG5 board (see Figure 2-13).

The InfiniteReality board set can be used in both deskside and rack systems. This graphics set consists of the GE14-4, RM7/TM7, and DG5 boards (see Figure 2-13).

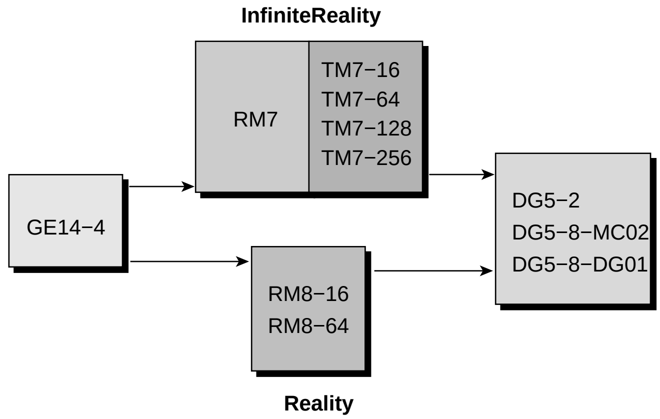

Figure 2-13 Onyx2 InfiniteReality and Reality Board Sets

2.4.1 The GE14 Board

The GE14 board (P / N 030-0681-00x) processes OpenGL® commands and data from the Node board and is the first stage of the graphics pipeline.

The GE14 receives vertex data defining the location, orientation, color, and texture-mapping coordinates of a polygon. If the polygon has more than three vertices, the GE14 subdivides the polygon into triangles. Triangles are the basic polygonal working units of the system. Data output from the GE14 geometry subsystem passes over the Triangle bus to the raster memory subsystem (RM7).

Primary ASICs on the GE14 include the

- crosstalk to graphics (XG)
- host interface processor (HIP)
- Geometry Engine distributor (GED)
- Geometry Engines (GE11s)
- back-end FIFO (BEF)

The HIP communicates with the DG5 (or optional DG5-8) board over the video control (VC) bus. The BEF drives the triangle bus (Tbus).

The graphics processor interface and the readback bus (Rbus) feed pixel data to the GED ASIC. The GED distributes pixel and vertex data to the GE11 ASICs on the GE14 board.

The GE11 processors each contain three floating-point unit (FPU) data paths, and four internal 512 x 32 blocks of support RAM. The GE11 is also supported externally by 64K x 72-bit words of microcode and a 64K x 32-bit-wide data SRAM. The GE11’s microprocessor can perform 480 million floating point operations per second (MFLOPS).

The GE11 ASICS also have these internal features:

- three floating point MUL / ALU cores
- address and data registers for each FPU core
- float-to-fix converters for each FPU core
- a 256 x 32 output FIFO for each FPU that buffers information going to the Tbus

The readback bus (Rbus) provides a path for pixels flowing from the RM7 framebuffer to the GED.

The Video Control bus provides access to the color maps, window display modes, and cursor control modes. The bus is a direct access from the host Node board to the HIP ASIC that acts as a gatekeeper on the GE14.

See Figure 2-14 for a simplified block diagram of the GE14 board.

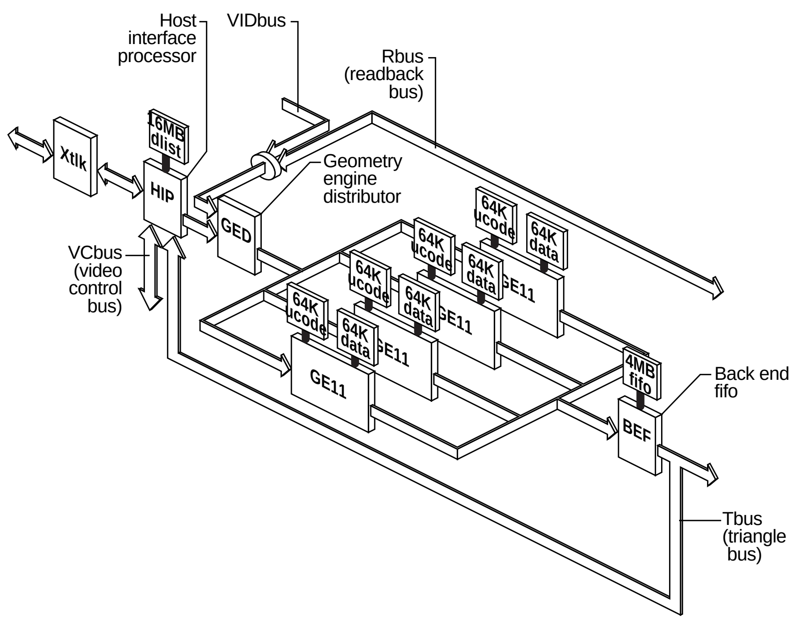

Figure 2-14 GE14 Functional Block Diagram

2.4.2 The Raster Memory Board

There are two versions of the Raster Memory board—the Onyx2 Reality RM8, which provides 40 MBs of raster memory and either 16 or 64 MBs of texture memory, and the Onyx2 InfiniteReality RM7 / TM7, which provides 80 MBs of raster memory and either 16 or 64 MBs of texture memory.

The Raster Memory board scan-converts triangle data from the triangle bus (Tbus) into pixel data. The triangles and other primitives are converted to 2 x 2 pixel quads that are optionally textured. The pixel quads are then sent to the framebuffer. Following the transfer to the framebuffer, the DG5 is given access to the color data stored in the completed frame. There are 16, 64, or 256 MBs of texel storage capacity (using SDRAM) on each TM7. The Texture Memory (TM7 board) is a daughterboard that resides on the RM7

Note: You cannot mix 16-MB, 64-MB, or 256 MB-TM7s in the same graphics pipe.

An RM7 / TM7 functional block diagram is shown in Figure 2-15. Note that a maximum of six RM7 / TM7s (two on one pipe and four on the other pipe) are supported in the Onyx2 InfiniteReality rackmount graphics subsystem.

The Raster Memory board is composed of the following main components:

- pixel generator (PG) ASIC
- texel generator (TG) ASIC
- texture manager (TM) ASIC
- SDRAM-based texture memory
- texture filter (TF) ASICs
- 20 image memory processors (IMPs) with two SGRAMs for each IMP ASIC in the Reality graphics RM8 board or 20 MPs with four SGRAMs for each IMP ASIC in the InfiniteReality graphics RM7 board

To get pixels into the framebuffer, the PG must scan-convert each primitive sent over the Tbus to find the parameter values for each of the polygon’s interior pixels. The PG distributes these values in 2 x 2 pixel blocks among the four TFs.

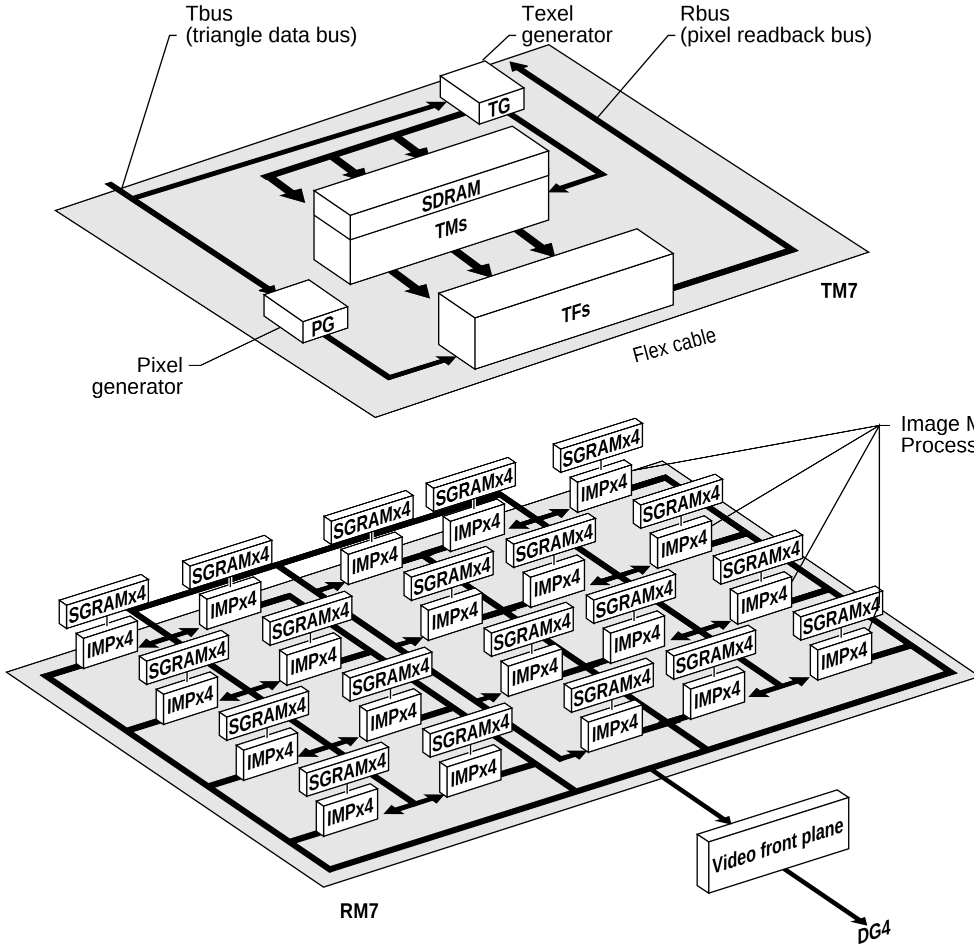

Figure 2-15 RM7 / TM7 Functional Block Diagram

Two data paths from the Tbus facilitate processing performance. One path processes the pixels for texture mapping via the texel generator and texture managers. The other path through the pixel generator produces RGB information that is passed directly to the texture filter ASICs.

The texel generator ASIC receives both “raw” triangles and generated X and Y pixel coordinates from the pixel generator. The texel generator then scan-converts each triangle to generate texture coordinates that are passed to the texture managers (TMs). The TMs perform lookups, formatting, and initial filtering of the texel quads. The resulting texture RGB information is passed to the TF ASICS, where it is combined with the basic RGB values from the pixel generator.

The TFs also perform fog calculations and texture table lookups before sending pixel quads to the IMPs. Each IMP ASIC contains four image engines, each handling a one-pixel-wide by two-pixels-high block at a time.

The IMP array receives the final color values and texture results from the texture filter.

The main responsibilities of the IMPs are to

- filter the subpixel-rendered image into the actual displayable framebuffer
- time-share the drawing and video display functions
- decide whether to write the pixel based on the z-buffer value
- perform the alpha blending of each new pixel with the pixel value that has already been rendered at the same location
- send the digital pixels to the display generator board (DG5) over the FP1 video front plane

Communication between the RM7(s) and the DG5 board is over the video front plane only.

The InfiniteReality board set can have up to four RM boards. As more RM boards are added, vertical display spans are interleaved, providing higher resolution and increased pixel fill rate.

2.4.3 The DG5 Board

There are two versions of the DG5 and both are multichannel capable. The DG5-2, which is the standard version, has two independent video channels. The DG5-8, which is the high-performance version, has eight independent video channels. Each version can drive separate RGB or virtual reality (VR) stereoscopic displays.

The display generator (DG5) subsystem requests and receives digital framebuffer pixel data from the RM board over the video front plane (see Figure 2-16). The DG5 processes the pixel data through an XMAP ASIC that sorts the pixels and streams them onto the video packet bus.

The DG5 also handles all pixel clocking and genlocking functions. The FM2 ASIC handles the role of the functional manager.

The XMAP ASIC handles cursor display functions.

Once the processed video data leaves the XMAP, it enters the packet bus and can be sent to one of three possible outputs:

- one of the video output channels
- an NTSC encoder
- a PAL encoder (VTR channel)

The video output controller (VOC) ASIC assembly consists of

- a video output formatter for generating video-timing signals
- a FIFO buffer

Each VOC ASIC supplies data to a 3-DAC array that feeds the analog RGB signals out.

NTSC or PAL circuitry signals come from the VOC through encoder and field buffer RAMs.

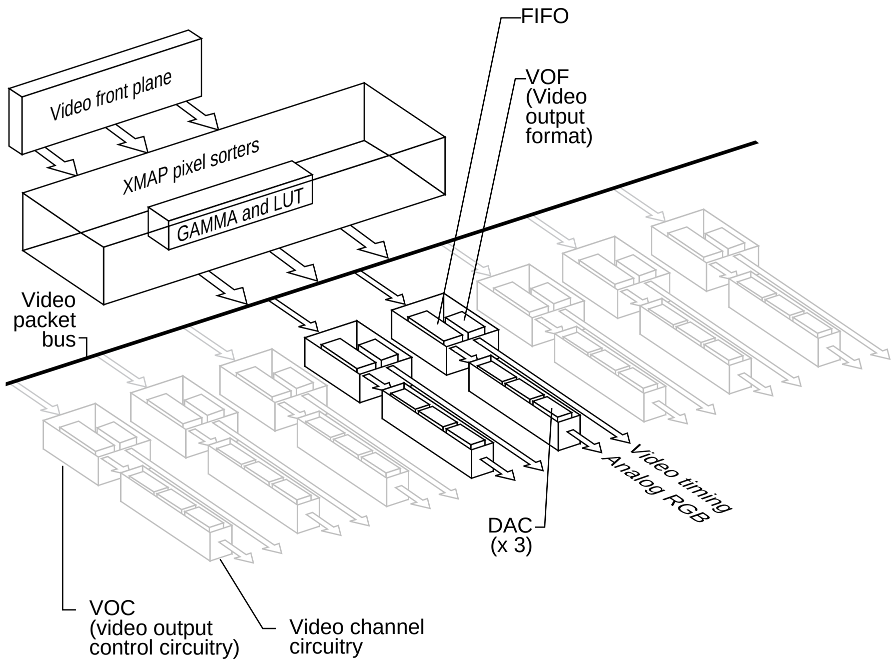

Figure 2-16 DG5 Board Functional Block Diagram

The video output on the DG5 can support two or (optionally) eight monitor connections. The default monitor resolution supported by InfiniteReality is 1,280 x 1,024 at 72 Hz. DG5 total pixel output support peaks at 300M pixels per second.

With two monitors x 1,280 x 1,024 x 72 Hz, only 188, 743, 680 (188M) pixels per second are used. However, with eight monitors x 1280 x 1024 x 72 Hz, 754, 974, 720 (755M) pixels per second is 455M pixels above the support limit. Therefore, to use eight monitor connections, you must use a combination of lower and higher resolution monitors within the limit of 300M pixels per second.

Chapter 3

Chassis Tour

This chapter provides an overview of the chassis for the Onyx2 and Origin2000 deskside and rackmount systems; a description of the controls, connectors, and indicators; a functional description of the midplane; and parts identification for the major components.

3.1 Key Points in This Chapter

Here are key points of information for this chapter.

- The standard Origin2000 and Onyx2 drive bay uses single connector assembly (SCA) drives *only*.
- The system disk is oriented differently from the other optional drives in the Origin2000 systems and the Onyx2 rack system. In the Onyx2 deskside system, all the drives are oriented in the same direction, with the handle on the right (as you face the system).
- If only one IP27 Node board is present in either the Origin2000 deskside and rack system or the Onyx2 rack system, then only 6 of the 12 XIO slots can be activated.
- Neither a Null Router nor a Star Router board can be used for CrayLink Interconnections.
- The Origin2000 and Onyx2 systems use PC-style serial cables. You cannot use earlier Silicon Graphics serial cables.

3.2 Origin2000 Deskside System

Figure 3-1, Figure 3-2, Figure 3-3, and Figure 3-4 provide various views of the Origin2000 deskside system, and Table 3-3 describes the numbered components in the figures.

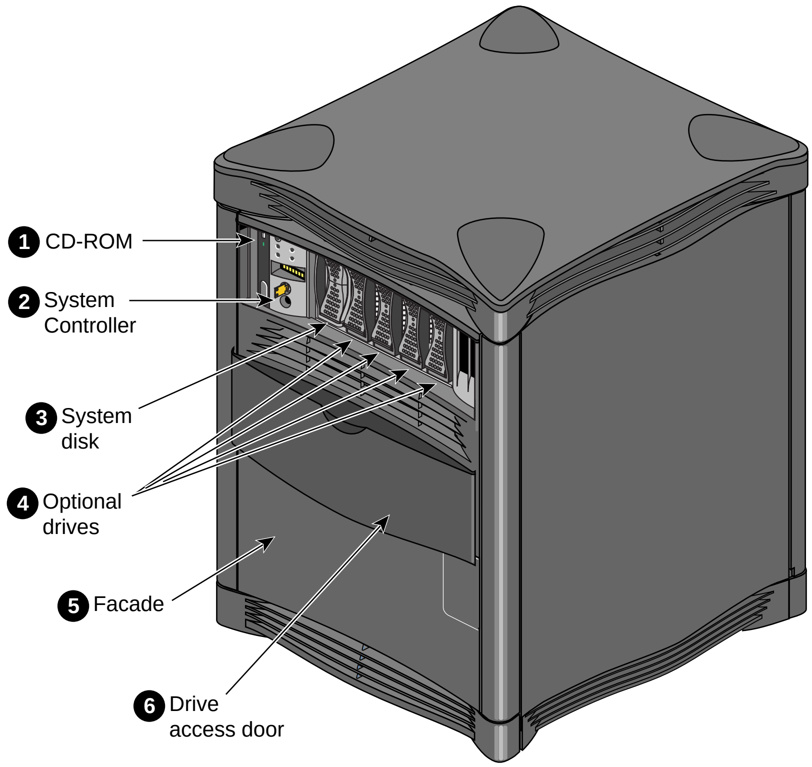

Figure 3-1 Origin2000 Deskside (Front View)

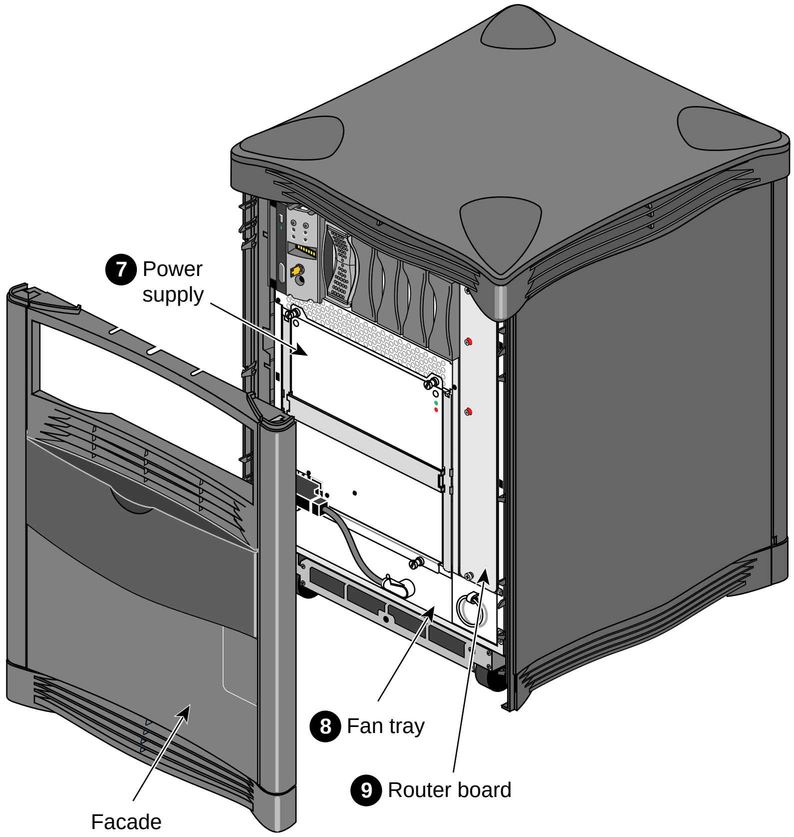

Figure 3-2 Origin2000 Deskside—Facade Removed

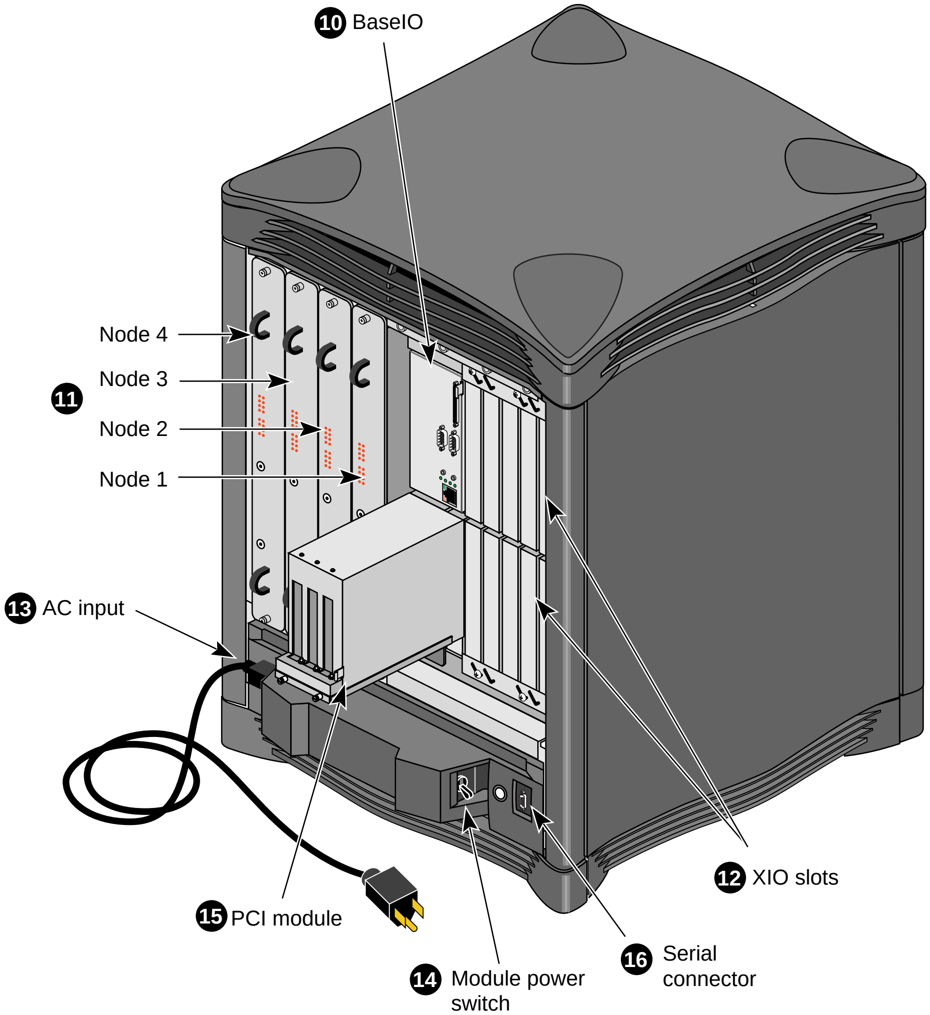

Figure 3-3 Origin2000 Deskside (Rear View)

Table 3-1 Origin2000 Deskside Hardware Components

Component	Description
1. CD-ROM drive	The CD-ROM is standard with each module. The CDs are installed vertically into the drive (see Chapter 6, “Installation”).
2. Module System Controller and display	The module System Controller is an independent, microprocessor-controlled device that powers up and helps boot the system. It also contains a key switch for turning the module off and on.
3. System disk	The system disk contains the operating system and other key software directories. The system disk must be installed in the drive position 1shown in Figure 3-5. Note that the system disk is oriented differently from the other drives in the system.

Table 3-1 (continued) Origin2000 Deskside Hardware Components

Component	Description
4. Optional SCA disk or tape drives	Each module can house up to five single connector assembly (SCA) drives (including the system disk). The disk drives are single-ended, ultra SCSI drives with a peak transfer rate of 40 MBs / sec.
5. Facade	The removable facade covers the power supply and fan tray for a module chassis.
6. Drive access door	The drive access door slides up or down to provide access to the drives. The door should be kept closed to prevent dust and other contaminants from affecting drive performance.
7. Power supply	This 1,900-watt auto ranging (110- to 220-VAC) power supply provides power for all deskside configurations.
8. Fan tray	The fan tray pulls in air from the top of chassis and exhausts the air through side vents in the bottom of the chassis.
9. Router board	The Router board ports provide high-speed (800 MBs / sec) connectivity between Node boards. These boards are sealed behind to help prevent the installation of CrayLink Interconnect cabling. This cabling is not supported in a deskside system. There are three types of Router boards: a Null Router, a Star Router, and a Rack Router board. These router boards are described in Section 3.7.6.
10. BaseIO board	This board provides basic I / O functions for the system such as serial ports, fast Ethernet, and single-ended wide SCSI. A dedicated slot in the XIO cardcage houses the BaseIO. This board cannot be installed in any other XIO expansion slot.
11. Node board	The Node board is the main processing board in the Origin2000 and Onyx2 systems. The Node board contains one or two R10000 CPUs, the Hub (which provides an interface to the I / O subsystem and the CrayLink Interconnect), a portion of main memory, as well as directory memory. Each Node board can support from 64 MBs to 4 GBs of memory. A single-rack system can have one to eight Node boards. A multiple-rack system can have up to 64 Node boards.
12. XIO slot cardcage	The XIO slot cardcage allows you to install additional I / O boards into the chassis. The system module can house up to 6 full-height boards or 12 half-height boards
13. AC input	The deskside accepts 110-VAC or 220-VAC input. Varies with configuration.
14. Module power switch	This is the main circuit breaker for the deskside chassis.
15. Peripheral connector interface (PCI) adapter	The optional PCI adapter allows you to install up to three third-party PCI boards. The PCI adapter provides an interface to the XIO slots of the Origin2000 system.
16. Serial connector	Electrically, this is the same connector that appears on the MSC front panel. This port can connect to a multimodule System Controller in a rack or to an external modem.

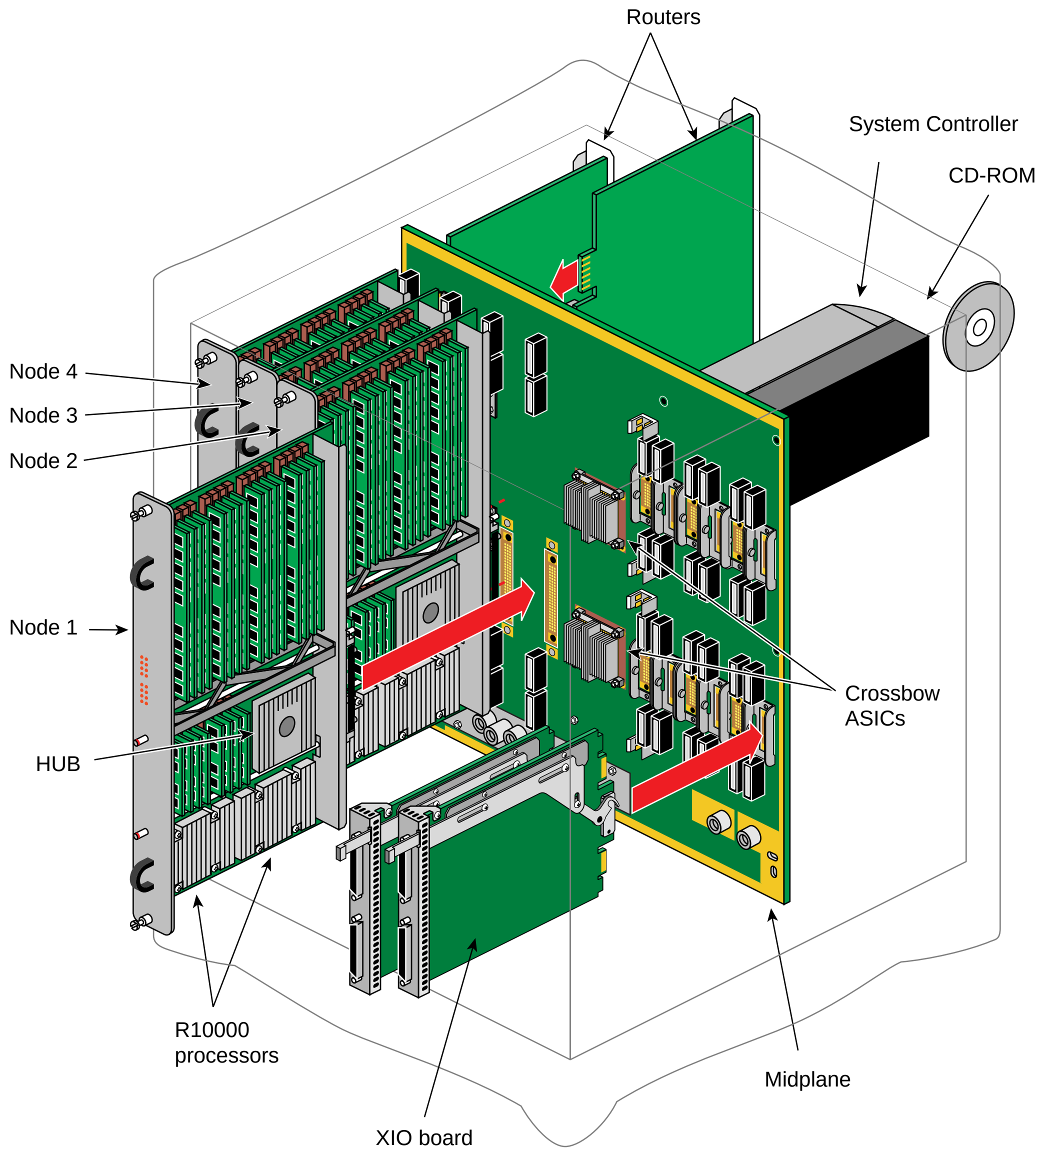

Figure 3-4 Origin2000 Deskside Front and Rear (Partial Exploded View)

3.3 Origin2000 Rackmount System

Figure 3-5 and Figure 3-6 show the major parts of the Origin2000 rackmount system, and Table 3-2 describes the numbered components in the figures.

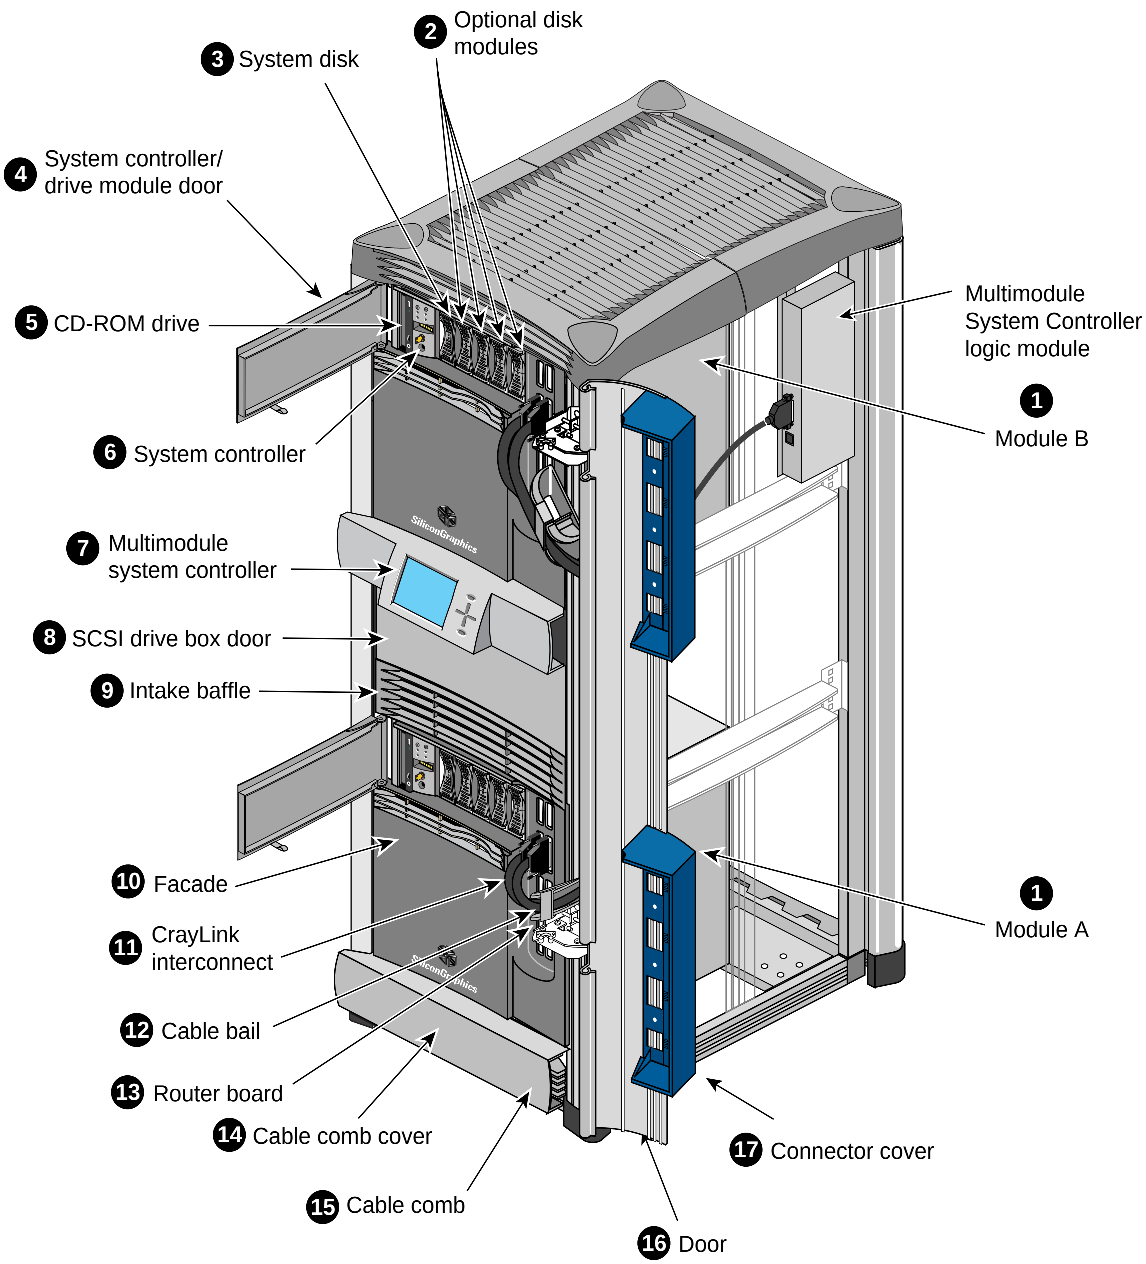

Figure 3-5 Origin2000 Rackmount System

Note: In Figure 3-5, the side panel is removed for clarity.

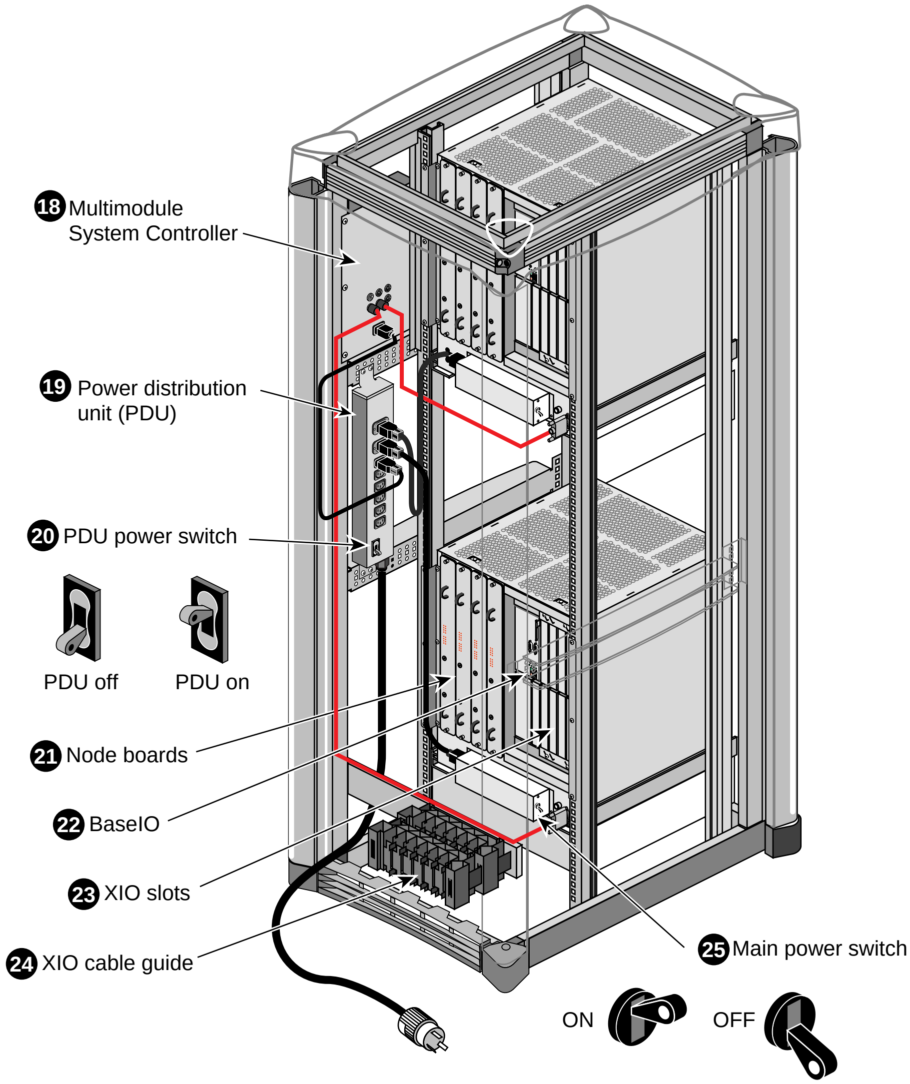

Figure 3-6 Origin2000 Rack Rear View (Bottom Module Shown)

Note: In Figure 3-6, the side panel is removed for clarity.

Table 3-2 Rackmount Hardware Components

Component	Description
1. Module A / Module B	The modules shown in Figure 3-5 are independent computing subsystems with a separate set of CPUs, disks, System Controller, and I / O connections. There can be up to 2 modules in a rack and up to 16 modules in a multirack configuration.
2. Optional SCA drives	Each module can house up to five single connector assembly (SCA) drives (including the system disk). The disk drives are single-ended, ultra SCSI drives with a peak transfer rate of 40 MBs / sec.
3. System disk	The system disk contains the operating system and other key software directories. The system disk must be installed in the drive position shown in Figure 3-5. Note that the system disk is oriented differently from the other drives in the system. See “Loading the Single Connector Assembly (SCA) Drives Into the System” in Chapter 6 for additional information.
4. System controller / drive module door	This door should be kept closed to prevent dust and other possible contaminants from affecting the drive performance.
5. CD-ROM drive	The CD-ROM drive is standard with each system and is attached to the module System Controller (MSC) as a single module (see “Removing the Module System Controller (MSC) and CD-ROM” in Chapter 10).
6. Module System Controller (MSC) and display	The module System Controller is an independent, microprocessor-controlled device that powers up and helps boot the system. It also contains a key switch for turning the module off and on
7. Multimodule System Controller (MMSC) multimodule display	The MMSC display is an intelligent keypad interface that can control all the system modules in a rack. The multimodule display has greater functionality than the MSC display. There can only be one MMSC display per rack configuration.
8. Origin Vault SCSI drive box	This drive box enclosure provides six half-height and two full-height, 5-1 / 4-inch drive slots for single-ended and differential SCSI or fibre channel drives.
9. Intake baffle	This baffle helps proper airflow through the rack. Note that the top of the rack has a vent as well. Airflow is generally pulled in from the top and middle of the rack and exhausted through the back and bottom of the rack.
10. Facade	The removable facade covers the power supply and router boards for a module chassis.
11. CrayLink Interconnect cabling	This is the physical link that enables the modules in a rack to communicate and share resources. The CrayLink Interconnect cable is made up of delicate copper strands. Be careful when handling this cable.
12. Cable bale	The cable bales hold the CrayLink Interconnect cable in place to prevent excessive bending, which can cause damage.

Table 3-2 (continued) Rackmount Hardware Components

Component	Description
13. Router board	The router board ports provide high-speed (800 MBs / sec) connectivity between Node boards. There are three types of Router boards: a Null Router, a Star Router, and a Rack Router board. These router boards are described in Section 3.7.6.
14. Cable comb cover	This aesthetic, removable cover hides the CrayLink Interconnect cable inside the rack chassis.
15. Cable comb	The comb holds the CrayLink Interconnect cable in place when it is tucked into the grooves
16. Cable door	The cable door hides the CrayLink Interconnect routing between modules.
17. Connector cover	The connector cover protects the router board ports and cabling.
18. Multimodule System Controller (MMSC)	The MMSC located in the rear of the chassis is a separate microprocessor-controlled unit that interfaces the individual MSCs and MMSCs from other systems.
19. Power distribution unit (PDU)	The PDU is the central power source for the rack. All the modules and peripherals connect to the PDU. Note that the PDU has a separate power switch.
20. PDU switch	The PDU switch is the main circuit breaker for the entire rack assembly.
21. Node boards	The Node board is the main processing board in the system. It contains one or two R10000 CPUs, the hub (which provides an interface to the I / O subsystem and the CrayLink Interconnect), a portion of main memory, as well as directory memory. Each node board can support from 64 MBs to 4 GBs of memory. A single-rack system can have 1 to 8 Node boards. A multiple-rack system can have up to 64 Node boards.
22. BaseIO board	This board provides basic I / O functions for the system such as serial ports, fast Ethernet, and single-ended wide SCSI. A dedicated slot in the XIO cardcage houses the BaseIO board. This board cannot be installed in any other XIO expansion slot.
23. XIO slot cardcage	The XIO cardcage allows you to install additional I / O boards into the chassis.
24. XIO cable guide	The XIO cable guide management helps ensure proper laying out of cables in the rear of the chassis.
25. Module power switch	This switch powers the individual modules on and off.

3.4 Onyx2 Deskside System

Figure 3-7 through Figure 3-10 provide various views of the Onyx2 deskside system and Table 3-3 describes the numbered components in the figures.

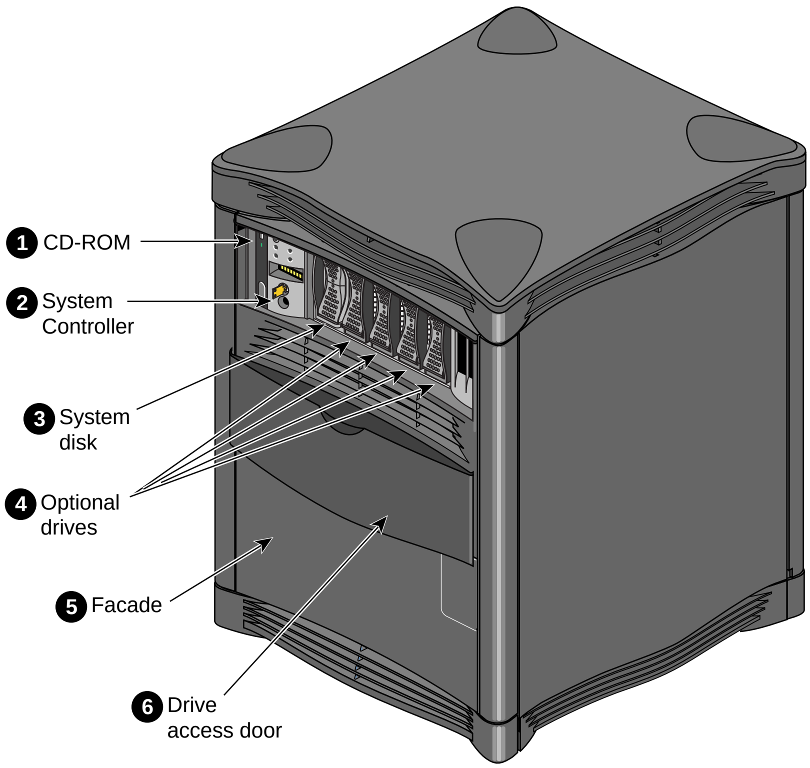

Figure 3-7 Onyx2 Deskside (Front View)

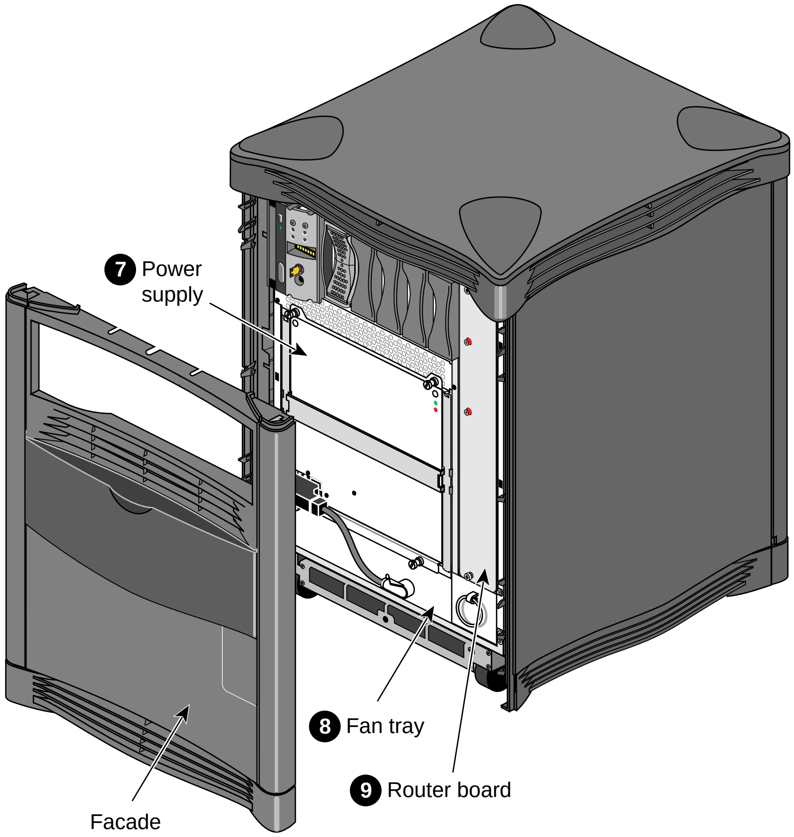

Figure 3-8 Onyx2 Deskside—Facade Removed

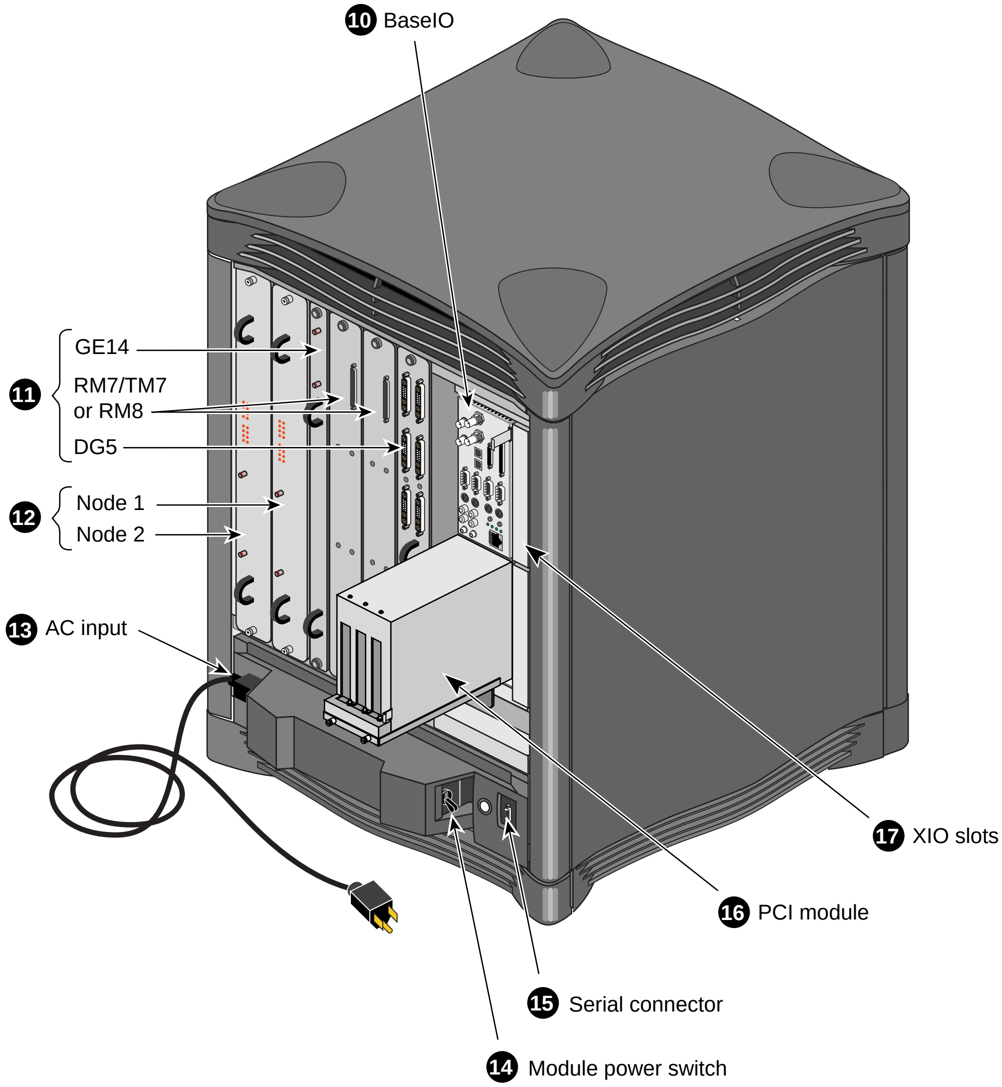

Figure 3-9 Onyx2 Deskside System (Rear View)

Table 3-3 Onyx2 Deskside Hardware Components

Component	Description
1. CD-ROM drive	The CD-ROM is standard with each module. The CDs are installed vertically into the drive (see Chapter 6, “Installation.”).
2. Module System Controller and display	The module System Controller is an independent, microprocessor-controlled device that powers up and helps boot the system. It also contains a key switch for turning the module off and on.
3. System disk	The system disk contains the operating system and other key software directories. The system disk must be installed in the drive position shown in Figure 3-7. Note that the system disk is oriented differently from the other drives in the system.

Table 3-3 (continued) Onyx2 Deskside Hardware Components

Component	Description
4. Optional SCA disk or tape drives	Each module can house up to five single connector assembly (SCA) drives (including the system disk). The disk drives are single-ended, ultra SCSI drives with a peak transfer rate of 40 MBs / sec.
5. Facade	The removable facade covers the power supply and fan tray for a module chassis.
6. Drive access door	The drive access door slides up or down to provide access to the drives. The door should be kept closed to prevent dust and other contaminants from affecting drive performance.
7. Power supply	This 1,900-watt auto ranging (110- to 220-VAC) power supply provides power for all deskside configurations.
8. Fan tray	The fan tray pulls in air from the top of chassis and exhausts the air through side vents in the bottom of the chassis.
9. Router board	The Router board ports provide high-speed (800 MBs / sec) connectivity between Node boards. This board are sealed behind to help prevent the installation of CrayLink Interconnect cabling. This cabling is not supported in a deskside system. There are three types of Router boards: a Null Router, a Star Router, and a Rack Router. These router boards are described in Section 3.7.6.
10. Graphics BaseIO board	This board provides basic I / O functions for the system such as serial ports, (fast) Ethernet, single-ended wide SCSI, audio, video, and a parallel port. A dedicated slot in the XIO cardcage houses the BaseIO. This board cannot be installed in any other XIO expansion slots.
11. Graphics board set	The Onyx2 graphics board set consist of a GE14 board, a Raster Memory board (either an RM7 / TM7 or an RM8), and a DG5 board. Each board occupies a separate slot. There can be up to two Raster Memory boards in an Onyx2 deskside system.
12. Node board	The Node board is the main processing board in the system. It contains one or two R10000 CPUs, the hub (which provides an interface to the I / O subsystem and the CrayLink Interconnect), a portion of main memory, as well as directory memory.
13. AC input	The deskside accepts 110-VAC or 220-VAC input. The chassis configuration determines whether 110-VAC or 220-VAC input is used.
14. Module power switch	This the main circuit breaker for the deskside chassis.
15. Serial connector	Electrically, this is the same connector that appears on the MSC front panel. This port can connect to a multimodule System Controller in a rack or to an external modem.
16. Peripheral connector interface (PCI) adapter	The optional PCI adapter allows you to install up to three third-party PCI boards. The PCI adapter provides an interface to the XIO slots of the system.
17. XIO slot cardcage	The XIO cardcage allows you to install additional I / O boards in the Onyx2 chassis. This chassis can house up to three full-height boards or five half-height boards.

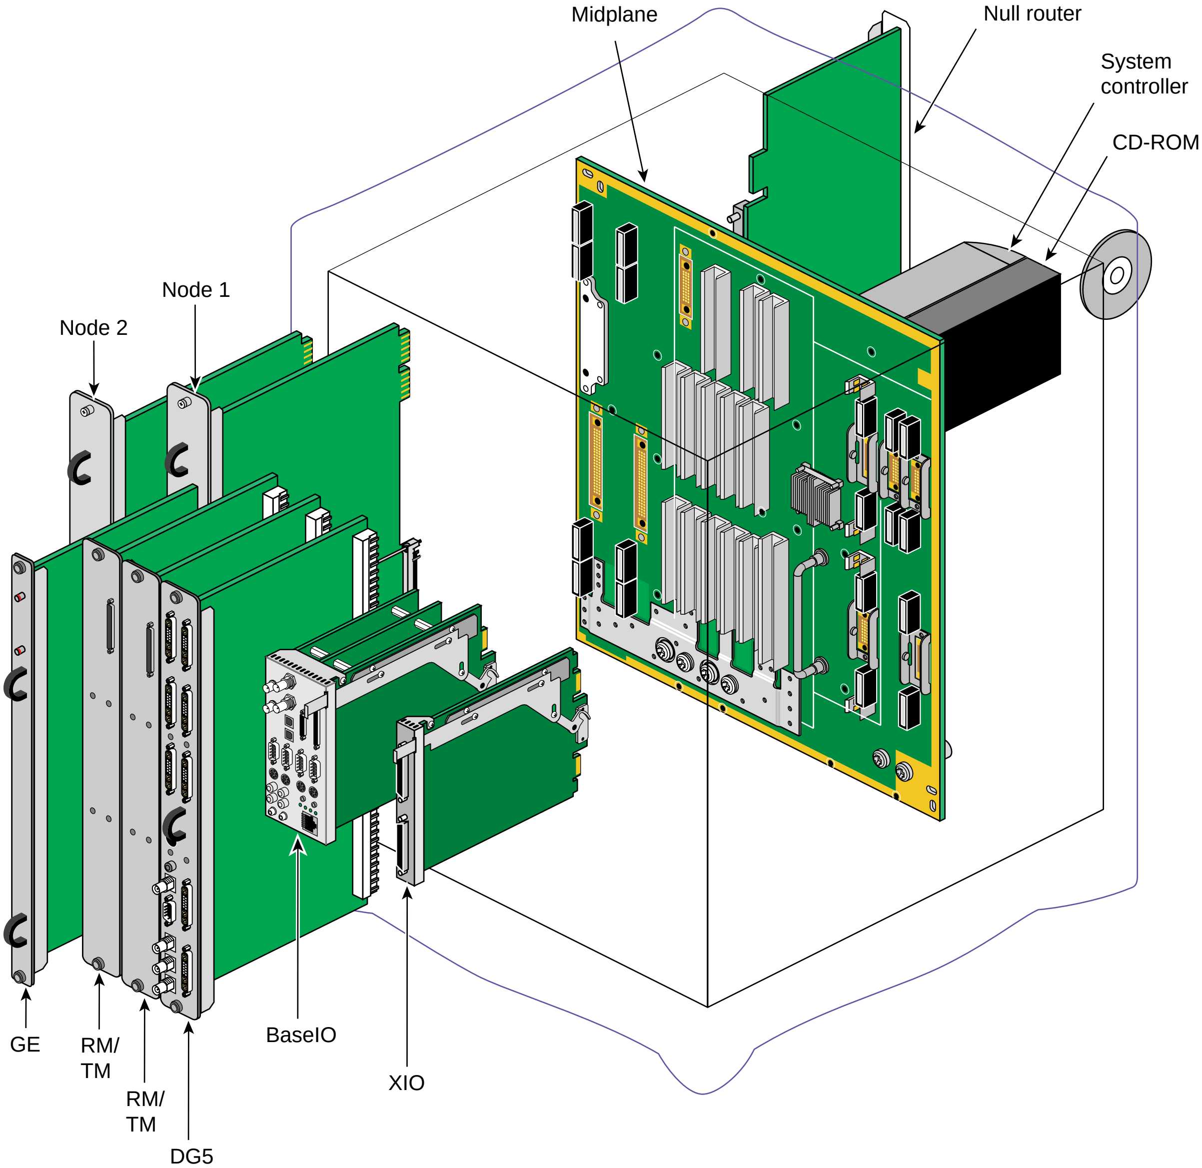

Figure 3-10 Ony2 Deskside Chassis (Partial Exploded View)

3.5 Onyx2 Rackmount System

Figure 3-11 and Figure 3-12 provide various views of the Onyx2 rackmount system, and Table 3-2 describes the numbered components in the figures.

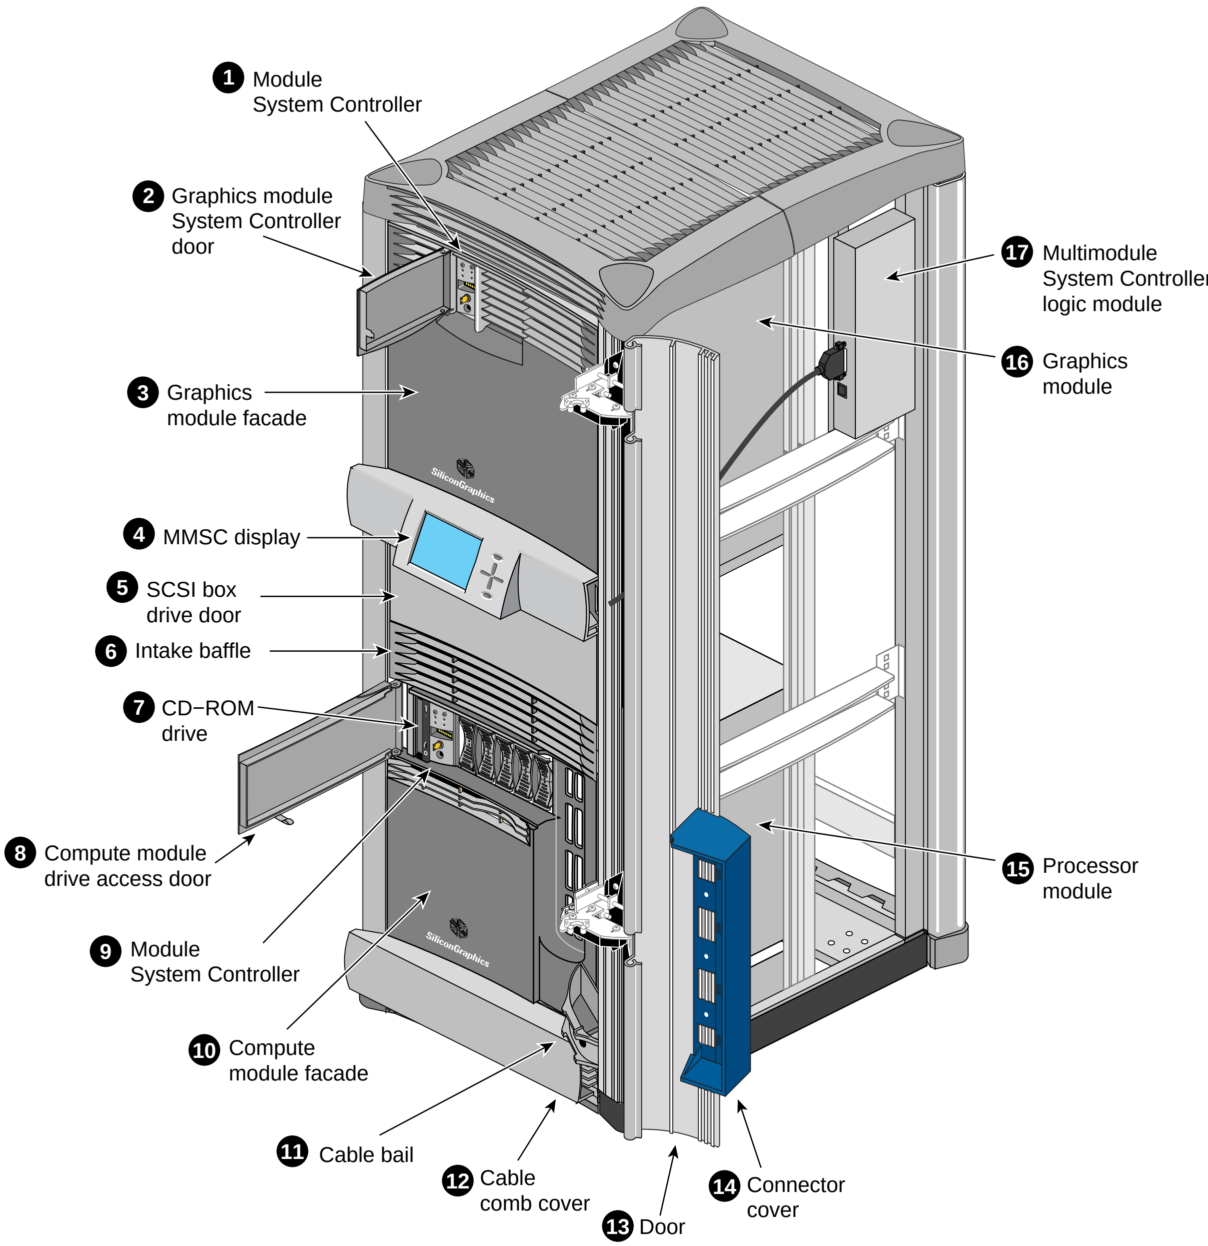

Figure 3-11 Onyx2 Rackmount—Front of Chassis

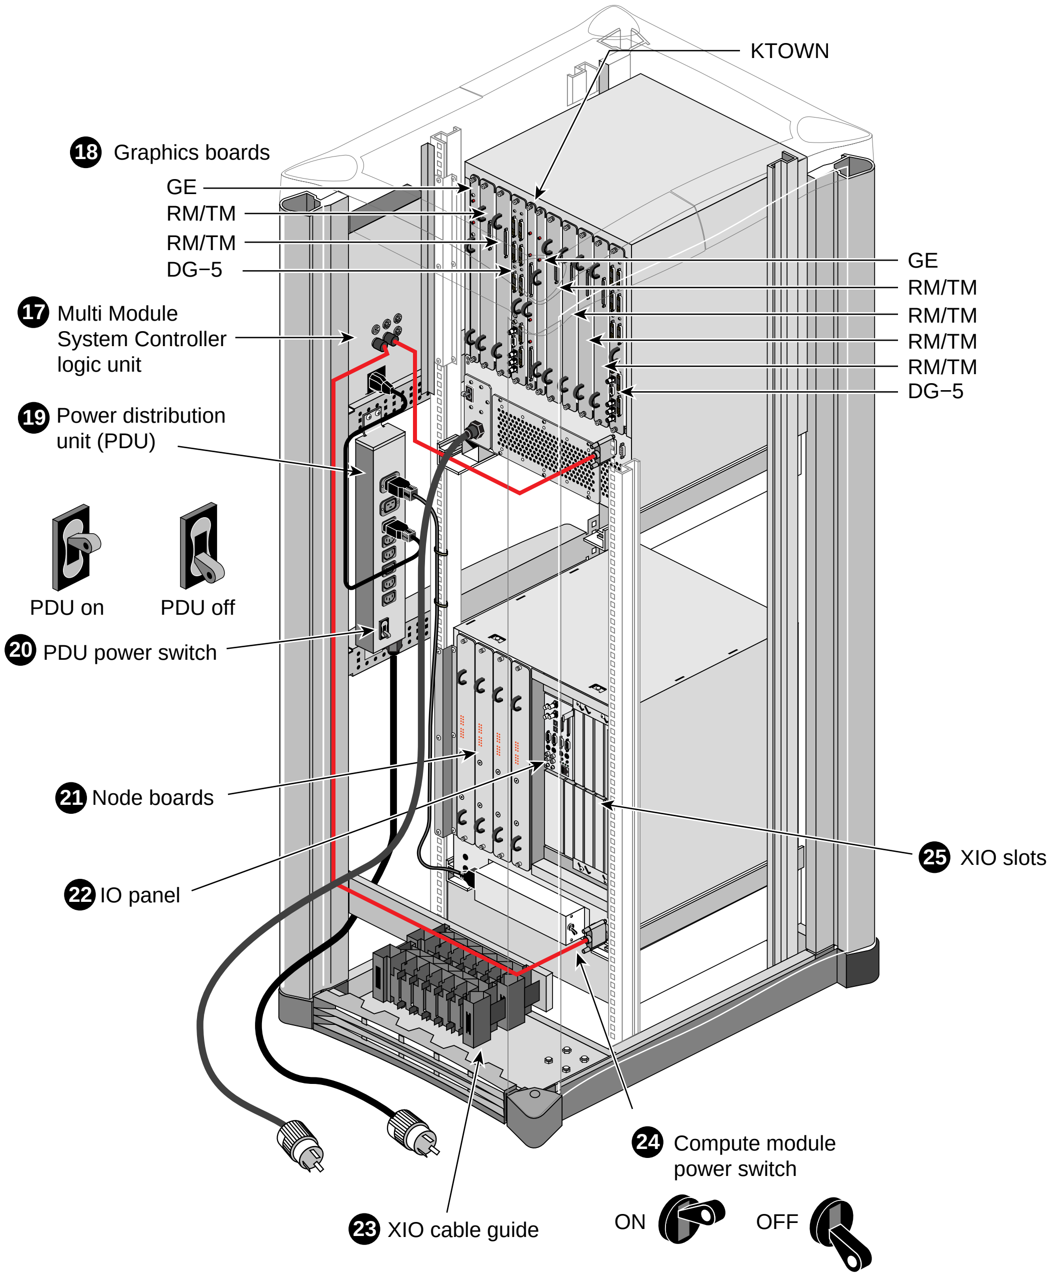

Figure 3-12 Onyx2 Rackmount System—Back of Chassis

Table 3-4 Onyx2 Rackmount Hardware Components

Component	Description
1. Module System Controller (MSC) and display	The module System Controller is an independent, microprocessor-controlled device that powers up and helps boot the system. It also contains a key switch for turning the module off and on.
2. System controller / drive module door	This door should be kept closed to prevent dust and other possible contaminants from affecting the drive performance.
3. Graphics module facade	The removable facade covers the power supply for the graphics module chassis.
4. Multimodule System Controller (MMSC) multimodule display	The MMSC display is an intelligent keypad interface that can control all the system modules in a rack. The multimodule display has greater functionality than the MSC display. There can only be one MMSC display per rack configuration.
5. Origin Vault SCSI drive box	This drive box enclosure provides six half-height and two full-height, 5-1/4-inch drive slots for single-ended and differential SCSI or fibre channel drives.
6. Intake baffle	This baffle helps proper airflow through the rack. Note that the top of the rack has a vent as well. Airflow is generally pulled in from the top and middle of the rack and exhausted through the back and bottom of the rack.
7. CD-ROM drive	The CD-ROM drive is standard with each system and is attached to the module System Controller (MSC) as a single module (see “Removing the Module System Controller (MSC) and CD-ROM” in Chapter 10).
8. Drive access door	The drive access door slides up or down to provide access to the drives. The door should be kept closed to prevent dust and other contaminants from affecting drive performance.
9. Module System Controller (MSC) and display	The module System Controller is an independent, microprocessor-controlled device that powers up and helps boot the system. It also contains a key switch for turning the module off and on
10. Facade	The removable facade covers the power supply and router boards for a module chassis.
11. Cable bale	The cable bales hold the CrayLink Interconnect cable in place to prevent excessive bending, which can cause damage.
12. Cable comb cover	This aesthetic, removable cover hides the CrayLink Interconnect cable inside the rack chassis.
13. Cable door	The cable door hides the CrayLink Interconnect routing between modules.
14. Connector cover	The connector cover protects the router board ports and cabling.
15. Processor module	The processor or server module is an independent computing subsystems with a separate set of CPUs, disks, System Controller, and I/O connections.

Table 3-4 (continued) Onyx2 Rackmount Hardware Components

Component	Description
16. Graphics module	The graphics module contains one or two graphics board sets, a separate power supply, and MSC.
17. Multimodule System Controller (MMSC)	The MMSC located in the rear of the chassis is a separate microprocessor-controlled unit that interfaces the individual MSCs and MMSCs from other systems.
18. Graphics boards	The Onyx2 rack holds up to two graphics pipes. The leftmost pipe (as you face the chassis rear) has one GE board, up to two RM/TM boards, and a DG-5. The other pipe has one GE board, up to four RM/TM boards, and a DG-5.
19. Power distribution unit (PDU)	The PDU is the central power source for the rack. All the modules and peripherals connect to the PDU. Note that the PDU has a separate power switch.
20. PDU switch	The PDU switch is the main circuit breaker for the entire rack assembly.
21. Node boards	The Node board is the main processing board in the system. It contains one or two R10000 CPUs, the hub (which provides an interface to the I/O subsystem and the CrayLink Interconnect), a portion of main memory, as well as directory memory. Each node board can support from 64 MBs to 4 GBs of memory. A single-rack system can have 1 to 8 Node boards. A multiple-rack system can have up to 64 Node boards.
22. BaseIO board panel	This board panel provides basic I/O functions for the system such as serial ports, fast Ethernet, and single-ended wide SCSI. A dedicated slot in the XIO cardcage houses the BaseIO board. This board cannot be installed in any other XIO expansion slot.
23. XIO slot cardcage	The XIO cardcage allows you to install additional I/O boards into the chassis.
24. Module power switch	This switch powers the individual modules on and off.
25. XIO cable guide	The XIO cable guide management helps ensure proper laying out of cables in the rear of the chassis.

3.6 Board Configuration and Layout

Note: The following information does not apply to the Onyx2 deskside system.

This section describes the location, preferred arrangement, and orientation of XIO boards, Node boards, and Router boards. It describes the mandatory orientation of XIO boards in the midplane, and the connections of Node boards with Router boards.

3.6.1 Location of the XIO and Node Boards

There is a direct correlation between the number of installed Node boards and the number of XIO slots that can be activated.

If a system module has only one Node board, then only 6 of the 12 XIO slots are activated. When two Node boards are installed, all XIO slots can be activated, *if* the Node boards are positioned in the correct slots.

Figure 3-13 diagrams the Node board-to-XIO-board correlation. The Node boards and their corresponding XIO board slots are indicated with either a circle or a triangle.

Node boards 3 and 4 also have corresponding XIO slots indicated by a circle or a triangle. For example, if Node boards are installed in slot 1 and slot 3, then only the corresponding XIO slots (1 through 6) are activated.

Some of the information in Figure 3-13 is provided on a label on the back of the deskside and rackmount chassis in a subsequent release of the product. The alphanumeric designators (for example, IO3 and IO4) are used by diagnostic programs to identify faulty boards. For example, if a diagnostics error message contains the IO3 designation, this indicates a problem with the board installed in slot IO3.

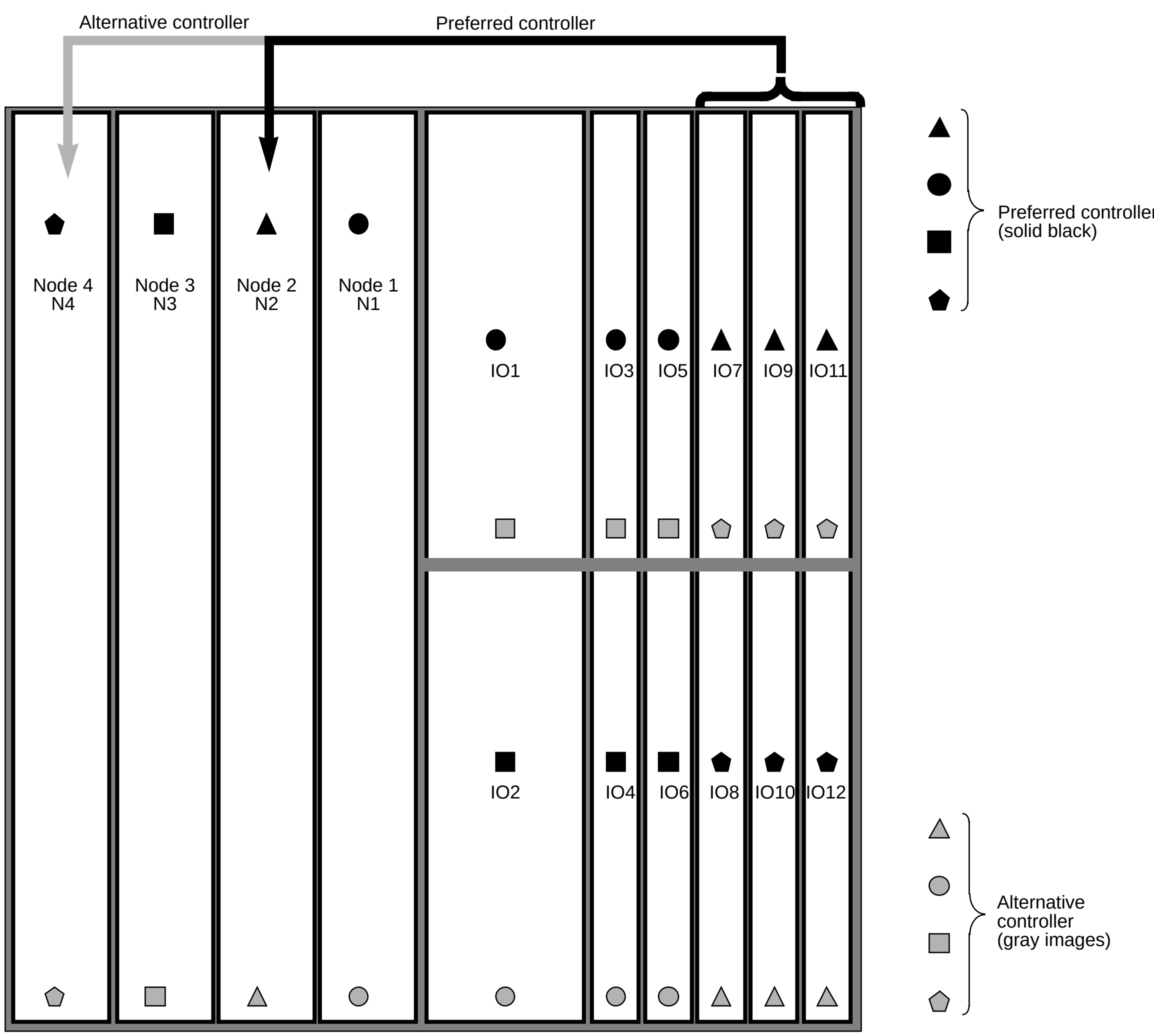

Figure 3-13 Board Configuration and Layout for a System Module

3.6.2 XIO Board Configurations

Two types of tags—black and gray—are used in this document to indicate *preferred* and *alternative* XIO board configurations.

- The solid black shapes in Figure 3-13—circle, triangle, square, pentagon—indicate the **preferred** IO-board-to-Node-board control connections.
- The gray shapes in Figure 3-13 indicate allowable **alternative** configurations.

Note: These labels are used in this document only; they are not on the system.

For instance, it is preferred that IO slot 7 (IO7) be controlled by Node board 2; therefore, both locations (IO7 and N2 in Figure 3-13) are tagged with solid black triangles. The solid black arrow going from IO7 to N2 in Figure 3-13 also indicates this preferred arrangement.

Although it is not the preferred configuration, IO7 can also be controlled by Node 4. This is indicated by the gray pentagons on IO7 and N4 in Figure 3-13. The gray arrow in Figure 3-13 also indicates this alternative configuration.

3.6.3 Orientation of XIO Boards

Figure 3-14 shows the way the XIO boards must be oriented when sliding them into their midplane slots. The component sides of each of the following pairs of XIO boards must face towards each other:

- 3 and 5
- 7 and 9
- 4 and 6
- 8 and 10

Installing them in any other position could damage the boards or the system.

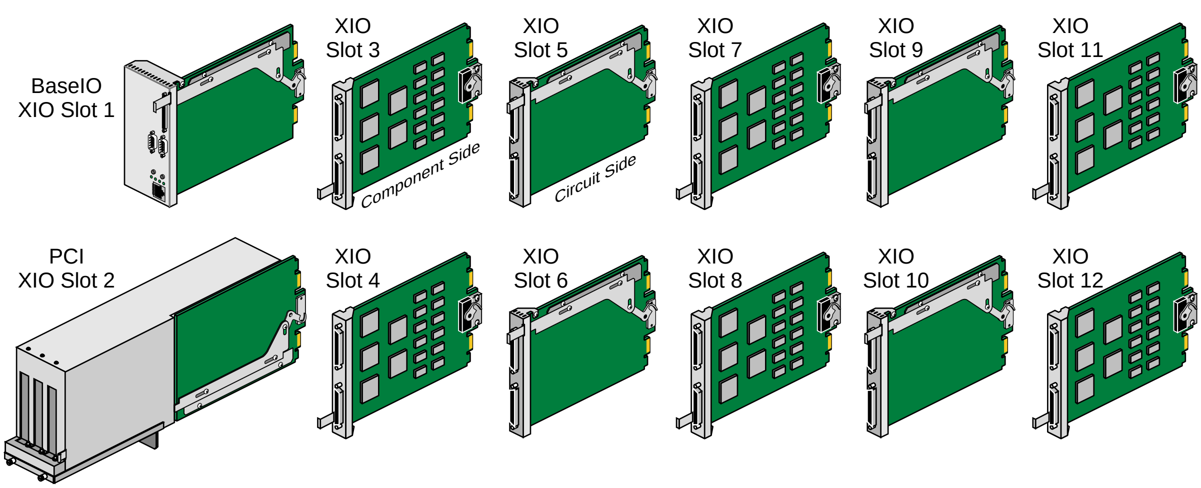

Figure 3-14 Orientation of the XIO Boards in Midplane Slots

3.6.4 Connecting Node Boards to Router Boards

Once they are in the module, Node boards are electrically connected to Router boards and Crossbow (XBOW) ASICs through the midplane. These connections are shown in Figure 3-15 (Router boards are physically located opposite the midplane from the Node boards, and XBOWs are located on the midplane).

- Node boards N1 and N2 connect to Router board R1
- Node boards N3 and N4 connect to Router board R2
- Node boards N1 and N3 connect to XBOW 1
- Node boards N2 and N4 connect to XBOW 2

These connections are also indicated by the black configuration tags (circle, square, etc.) that were first described in Figure 3-13.

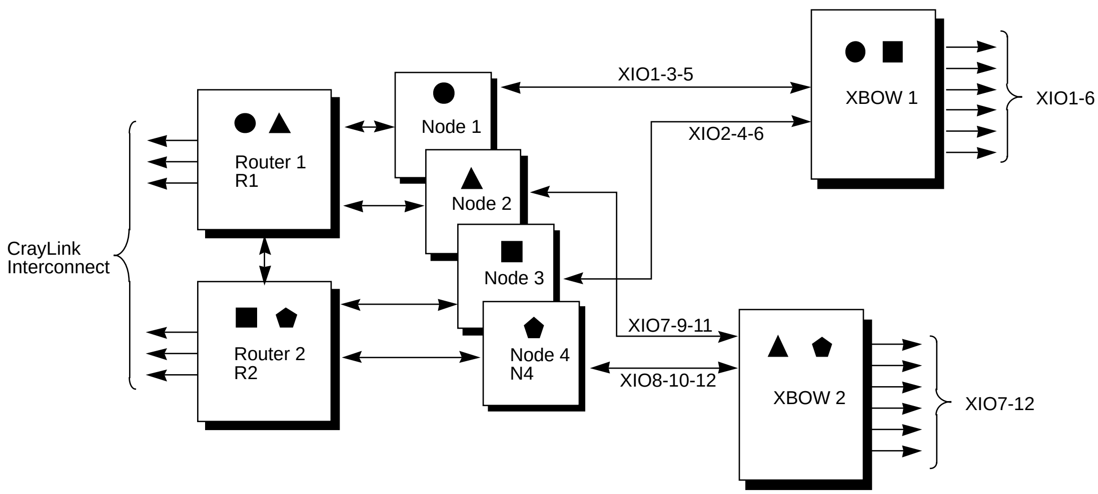

Figure 3-15 Origin2000 Overall Block Diagram

Figure 3-16 shows the connection dependencies of the Node boards, Routers boards, XBOWs, and XIO boards. For instance, Node 3 is connected to XBOW 1 and Router 2; Router 1 is connected to both Nodes 1 and 2.

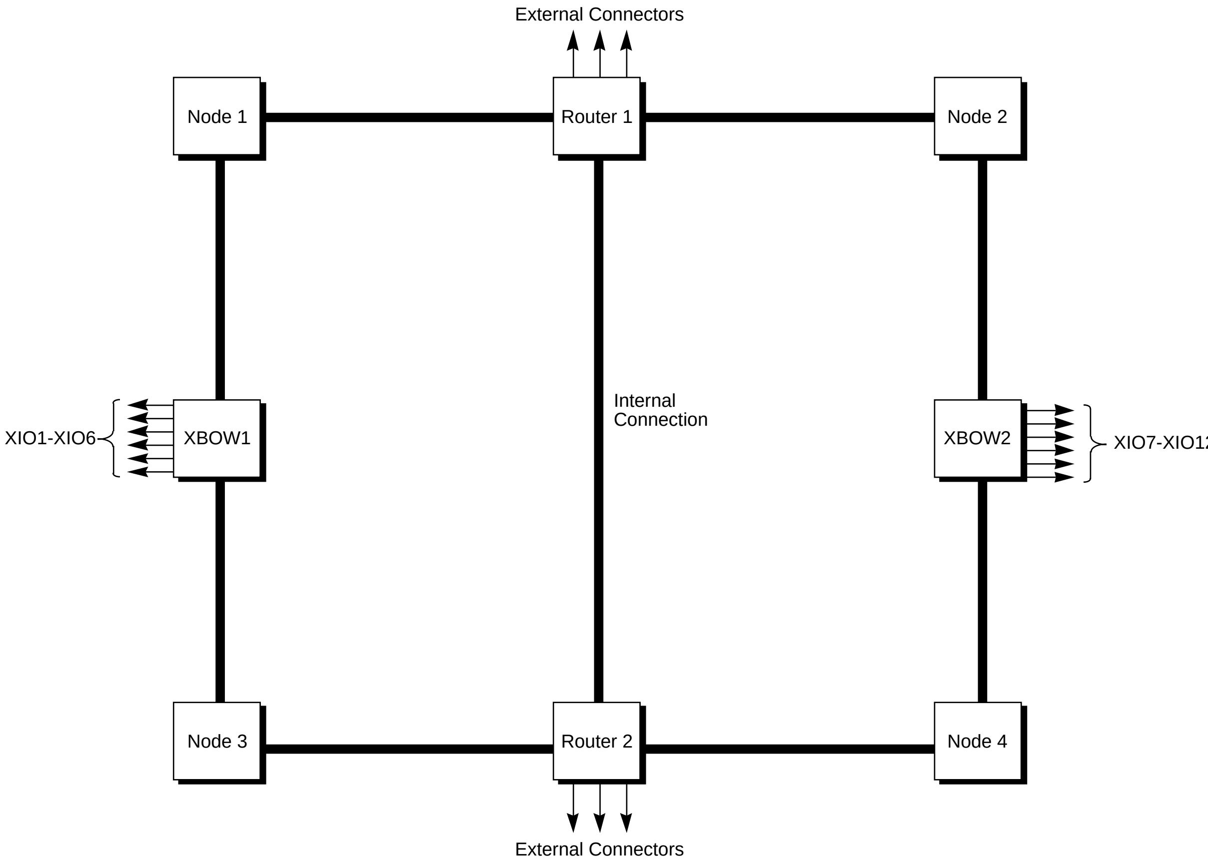

Figure 3-16 Graph of XIO, XBOW, Router, and Node Board Connections

3.6.5 Configuring Single and Multiple-Module Systems

This section lists six basic configuration rules that must be followed when configuring the Origin2000 system. It also provides configuration guidelines for:

- single modules
  - installing Node boards in a single-module system
  - installing XIO boards in a single-module system
  - loading the BaseIO board and the optional PCI expansion module
- multiple modules
  - optimizing the installation of Node boards in more than one module
  - optimizing the installation of XIO boards in more than one module

3.6.5.1 Configuration Rules

There are six rules, for both single-module and multiple-module systems, that *must be followed* when configuring the Origin2000 system:

- Each Origin2000 module must have at least one Node board.
- There must be at least one BaseIO board in the system (whether you have a single-module or multiple-module system).
- The BaseIO board must be placed in slot IO1.
- If present, the PCI expansion must be placed in slot IO2 (beneath the BaseIO board).
- Neither a BaseIO board nor an MSCSI board can ever go into slot IO2.
- To use all 12 of the XIO slots in a module, there must be at least two Node boards installed: one Node board must be placed in an odd slot (either N1 or N3) and one Node board must be placed in an even slot (either N2 or N4).

Caution: Remember, when installing XIO boards 3 through 9 and 4 through 10, each pair of boards must be inserted with their component sides facing **toward** each other, as shown in Figure 3-14.

3.6.5.2 Guidelines for a Single-Module System

This section provides guidelines for adding XIO (including BaseIO and the optional PCI expansion module) and Node boards to a single-module system.

BaseIO and PCI expansion module

- **Loading BaseIO and PCI expansion modules**—Install the BaseIO board in slot IO1. If present, install the PCI expansion module in slot IO2. Refer to "Configuration Rules" on page 25 for additional information.

Node boards

- **Installing Node boards in a single-module system**—When filling a single module, install the Node boards from center outward, starting at slot N1. Then add N2, N3, and finally N4, as shown in Figure 3-17.

  Note: When installing XIO boards in the system (described below), disperse them across all the available Node boards, in order to spread the bandwidth and control as much as possible.

- **Installing XIO Boards in a single-module system**

  XIO boards

  – If there is only one Node board in the module, it is placed in slot N1 and only the leftmost XIO slots (slots 1 through 6) are usable. In this case, fill the XIO slots sequentially, 3 through 6. Remember, slot IO1 must hold the BaseIO board, and slot IO2 holds the PCI expansion module (if the PCI expansion module is included in the system).

  XIO boards

  – If there is more than one Node board in the module, fill the XIO slots as described in Table 3-5 and shown in Figure 3-17, so that XIO boards are spread as evenly as possible across the system.

Table 3-5 gives a recommended installation pattern. In it, XIO boards are installed so they are connected beginning with Node 1, and then progress through Nodes 2, 3, and 4.

Installation then starts over with Node 1 again. In this way the XIO bandwidth and control is equalized over the four Node boards.

Table 3-5 Installing XIO Boards to Equalize Bandwidth and Control

Order of Installation	Start Loading XIO Slots At:		Node Board That Controls the XIO Slot
First XIO Board loaded is	IO1	which connects to	Node 1
Next XIO Board loaded is	IO7	which connects to	Node 2
Next XIO Board loaded is	IO2	which connects to	Node 3
Next XIO Board loaded is	IO8	which connects to	Node 4
Next XIO Board loaded is	IO3	which connects to	Node 1
Next XIO Board loaded is	IO9	which connects to	Node 2
Next XIO Board loaded is	IO4	which connects to	Node 3
Next XIO Board loaded is	IO10	which connects to	Node 4
Next XIO Board loaded is	IO5	which connects to	Node 1
Next XIO Board loaded is	IO11	which connects to	Node 2
Next XIO Board loaded is	IO6	which connects to	Node 3
Final XIO Board loaded is	IO12	which connects to	Node 4

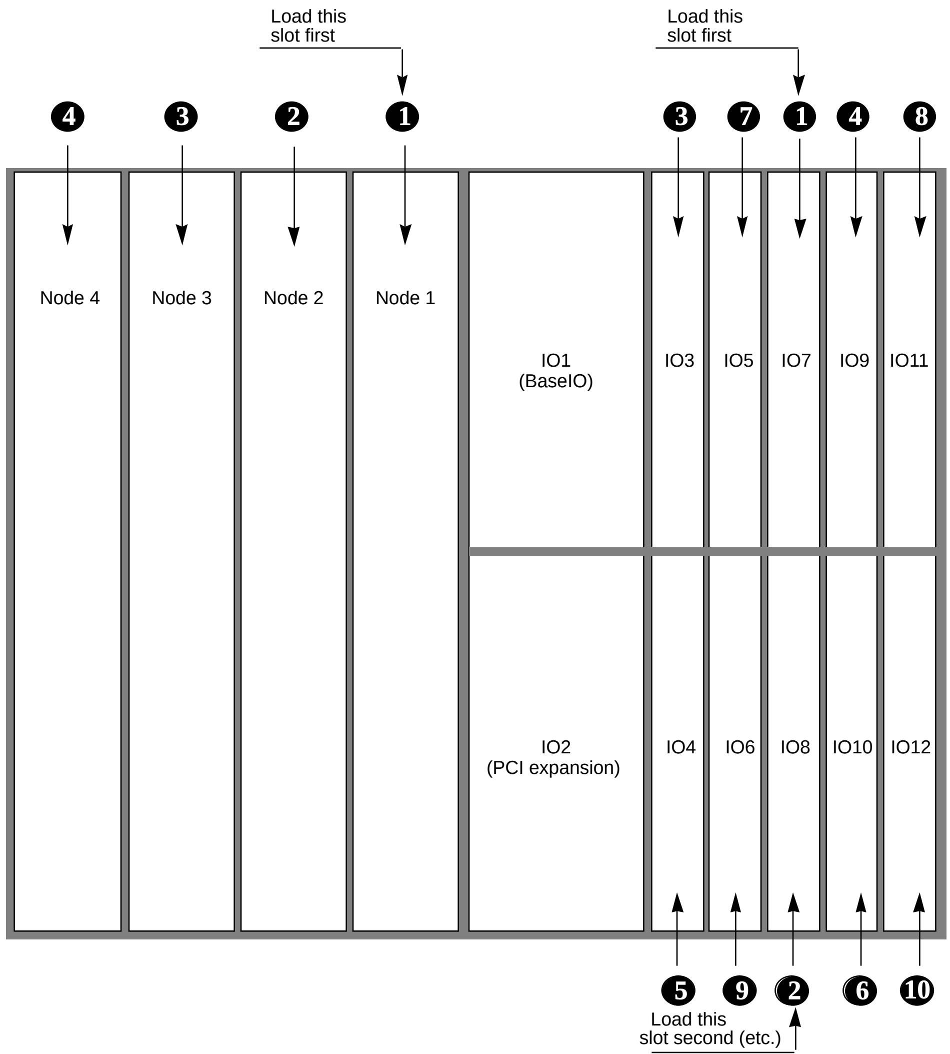

Figure 3-17 Installing Node and XIO Boards in a Module

3.6.6 Guidelines for a Multiple-Module System

This section provides guidelines for adding Node and XIO boards to a system that has more than one module. These guidelines are aimed at providing optimum installation for either Node boards or XIO boards, as follows:

- For most efficient installation of Node boards in a multimodule system, see the section titled, “Optimizing the Installation of Node Boards in More Than One Module” on page 28.
- For most efficient installation of XIO boards in a multimodule system, see the section titled, “Optimizing the Installation of XIO Boards in More Than One Module” on page 28.

3.6.6.1 Optimizing the Installation of Node Boards in More Than One Module

Follow the guidelines in this section for most efficient placement of Node boards in your system.

Node boards

1. Add Node boards to a module until that module is full, starting with slot N1, then N2, N3, and finally N4. This is the order shown in Figure 3-17.

   Note: This step differs importantly from Node board installation guidelines presented in the next section, which describes most efficient placement of XIO boards.

Node boards

2. Only when a module is filled should you add Node boards to another module.

   This means you cannot have a module with empty Node board slots. For instance, a seven Node board system can have only two system modules. Four of the Node boards go in one module, and the remaining three Node boards go in the second module.

XIO boards

3. Add XIO boards following the rules starting with Step 2 in the next section, "Optimizing the Installation of XIO Boards in More Than One Module."

3.6.6.2 Optimizing the Installation of XIO Boards in More Than One Module

Follow the guidelines in this section for most efficient placement of XIO boards in your system.

Node boards

1. Distribute the Node boards among slots N1 and N2 as evenly as possible throughout all modules. That is, fill slots N1 and N2 in all modules before adding any Node boards to slot(s) N3 or N4. This allows optimal distribution of XIO cards in each module: to use all 12 XIO slots, there must be a Node board installed in one of the odd slots (N1, N3) and a Node board installed in one of the even slots (N2, N4). For instance, in a four-module system with seven Node boards, three of the modules have Node boards in both slots N1 and N2, while the remaining module has a single Node board in slot N1 (2N+2N+2N+1N).

   Note: This section describes the most efficient placement of XIO boards. This step differs importantly from Node board installation guidelines presented in the previous section, which describes most efficient placement of Node boards.

BaseIO boards

2. Install all available BaseIO boards into IO1 slots. Remember, there can be only one BaseIO board in a module, and the BaseIO can go only into the IO1 slot.

PCI expansion module

3. Install all available PCI expansion boxes into IO2 slots. Remember, there can be only one PCI expansion box in a module, and the optional PCI expansion can go only into the IO2 slot.

XIO boards

4. Divide the remaining XIO boards evenly among the modules**.** Use the order described in Table 3-5 and shown in Figure 3-17, keeping in mind the number of Node boards installed. (In other words, if you have only Node 1 installed, do not install any XIO boards in slots IO7 through IO12.)

3.7 Major Components

This section identifies key parts that are common to both the Origin2000 and the Onyx2 deskside and rackmount systems.

- Node board
- System Controller
- BaseIO board
- optional XIO boards
- PCI adapter
- Router boards
- Midplane
- power distribution unit (PDU)
- CrayLink Interconnect, Xpress link, and Crosstown cables
- PCI adapter
- Remote serial connector
- DG5 display generator board
- 24-inch monitor
- serial cables

3.7.1 Node Board

Figure 3-18 shows the location of the Node boards in the system. The Node board is the main processing board in the Origin2000 and Onyx2 systems. An individual system module may have up to four Node boards. Figure 3-19 shows the required Node board slot positions in a module. The first Node board must be installed in the rightmost slot (as you face the rear of the chassis). Additional Node boards are added sequentially to the left.

Figure 3-19 and Figure 3-20 identify major components on the Node board. The Node board houses the R10000 CPU, the main memory DIMMs, and the Hub ASIC. As described earlier, the Hub is the central I/O processor for the Node board interfacing. A group of LEDs on the bulkhead of the board provides status information for the board.

3.7.1.1 R10000

The 180 or 195 MHz superscalar R10000 CPU is the main processor for the Origin2000 and Onyx2 systems. There can be one or two R10000 CPUs on each IP27 Node board. You may hear a module or rack referred to as an “8P” or “16P” system. The *P* stands for R10000 *p*rocessor. An 8P system has eight R10000 processors.

3.7.1.2 Main Memory

The Origin2000 and Onyx2 systems use SDRAMs mounted on dual inline memory modules (DIMMs) for main memory. An IP27 Node board can have from 64 MBs to 4 GBs of main memory. Memory upgrades are available in 64-, 128-, and 256-MB increments.

3.7.1.3 Premium Directory Memory

The optional premium directory DIMMs are only required for configurations with more than 16 Node boards. These directory memory DIMMs maintain cache coherence in large system configurations. Cache coherency helps provide data consistency when multiple processors need to access the same piece of memory.

3.7.1.4 Hub

The Hub on the Node board is the primary communication link between the R10000 processor, I/O subsystem, main memory, and CrayLink Interconnect. The Hub also interfaces with directory memory, which is responsible for maintaining cache coherency

3.7.1.5 Status LEDs

Figure 3-19 shows the bulkhead on the front of the Node board. The LEDs provide status information for the individual board. For information on the LED board states, see Chapter 7, “Diagnostic Tools.”

3.7.1.6 Flash PROM

This 1-MB PROM contains code to initialize and set up the R10000s, memory, console port and CrayLink Interconnect. The IP27 PROM also contains NVRAM information and provides the power-on diagnostics (POD) command interpreter.

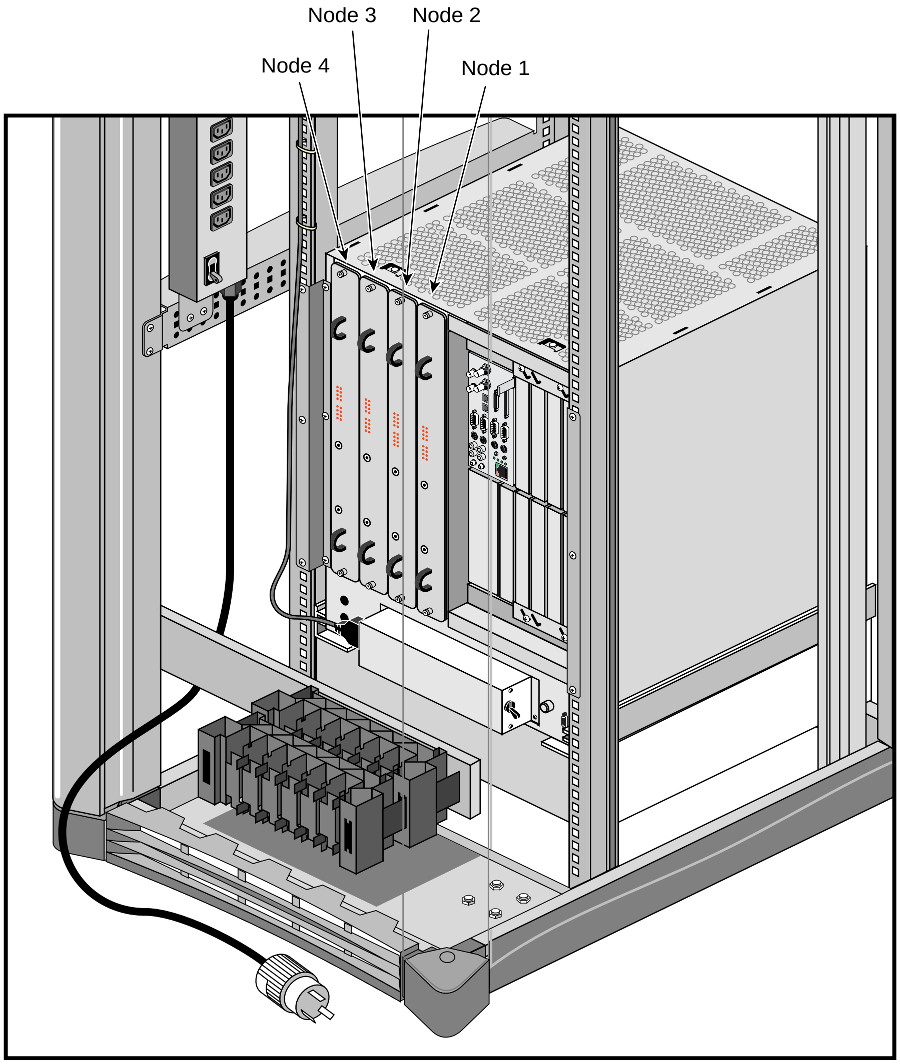

Figure 3-18 Node Board Positioning in a Module

Note: The slot designations are the same for a deskside or a rackmount system.

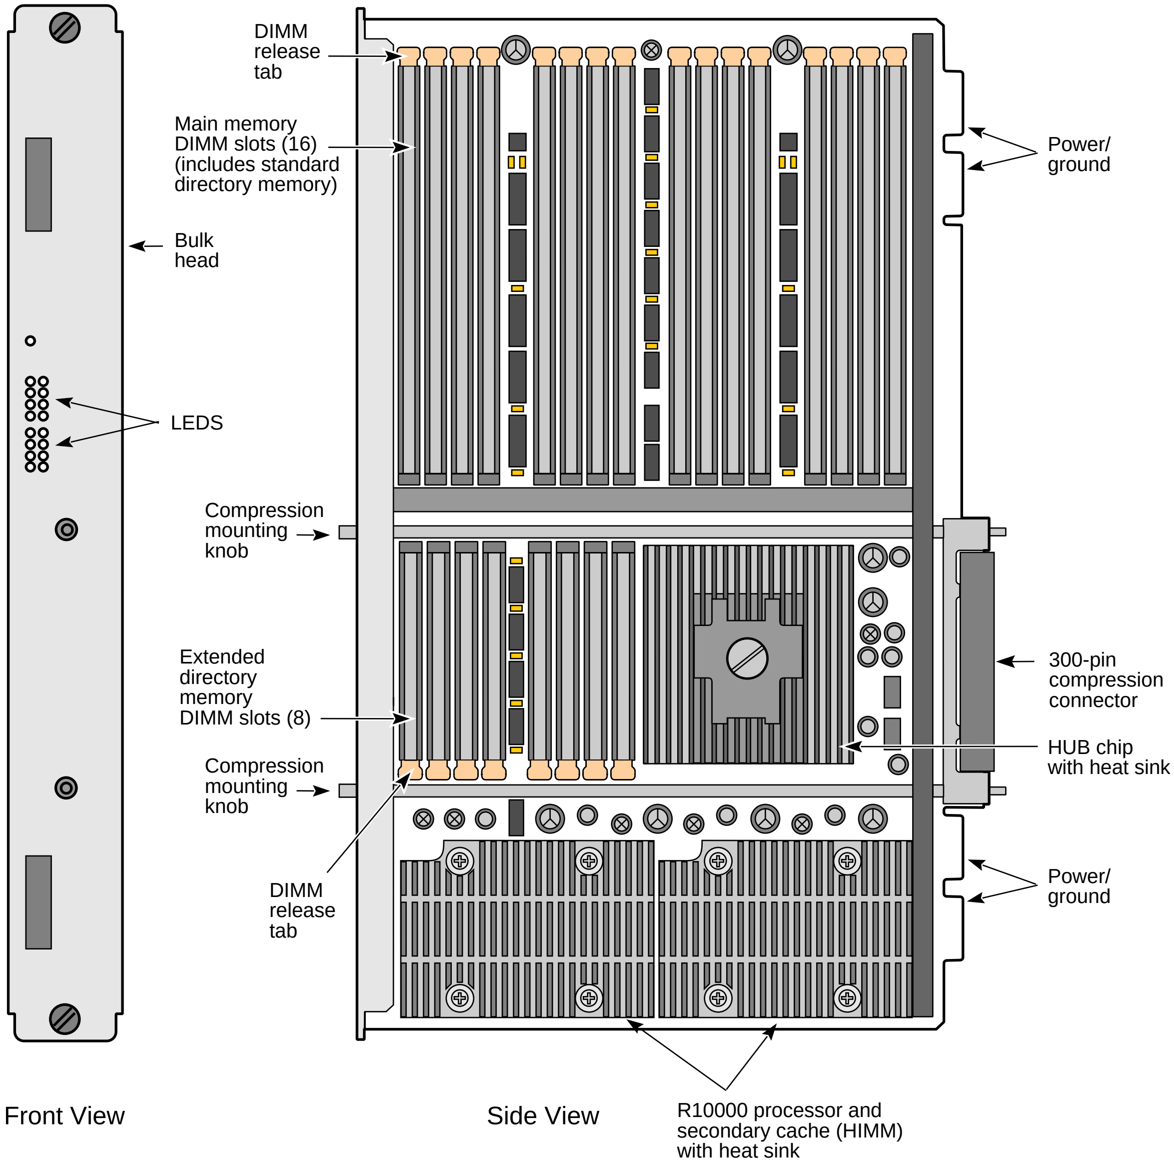

Figure 3-19 Node Board Components

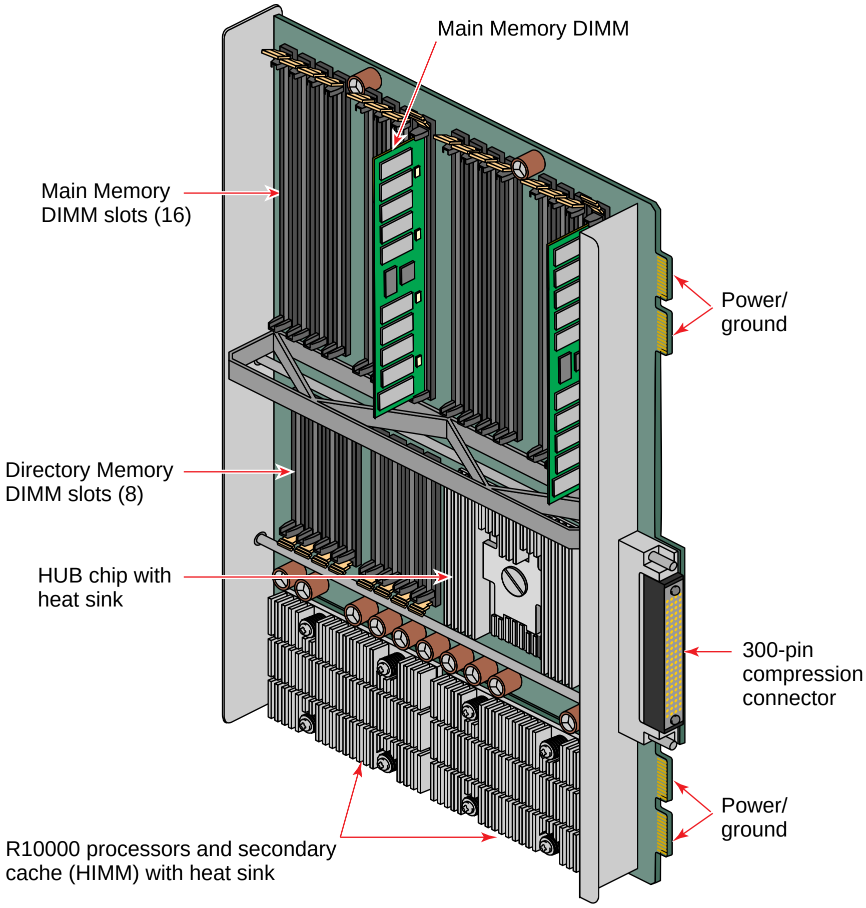

Figure 3-20 Node Board (Component Side View)

3.7.2 System Controller and Display

There can be two types of System Controllers for the Origin2000 and Onyx2 systems—the MSC, which is a single module System Controller, and an MMSC, which is a multimodule System Controller. The deskside system has the MSC, whereas the rackmount has both the MSC and MMSC and display. For detailed information on operating the System Controller, see the system *Owner's Guide*.

3.7.2.1 Module System Controller (MSC)

The MSC (see Figure 3-21) provides environmental and status monitoring for an individual system module.The module System Controller is also known as the ELSC—entry-level System Controller.

The MSC has a set of eight DIP switches for hardware debugging. The DIP switches are located in the back of the module controller panel (see Figure 3-22). These switches are described in Chapter 7, "Diagnostic Tools." The settings should normally be in the off position.

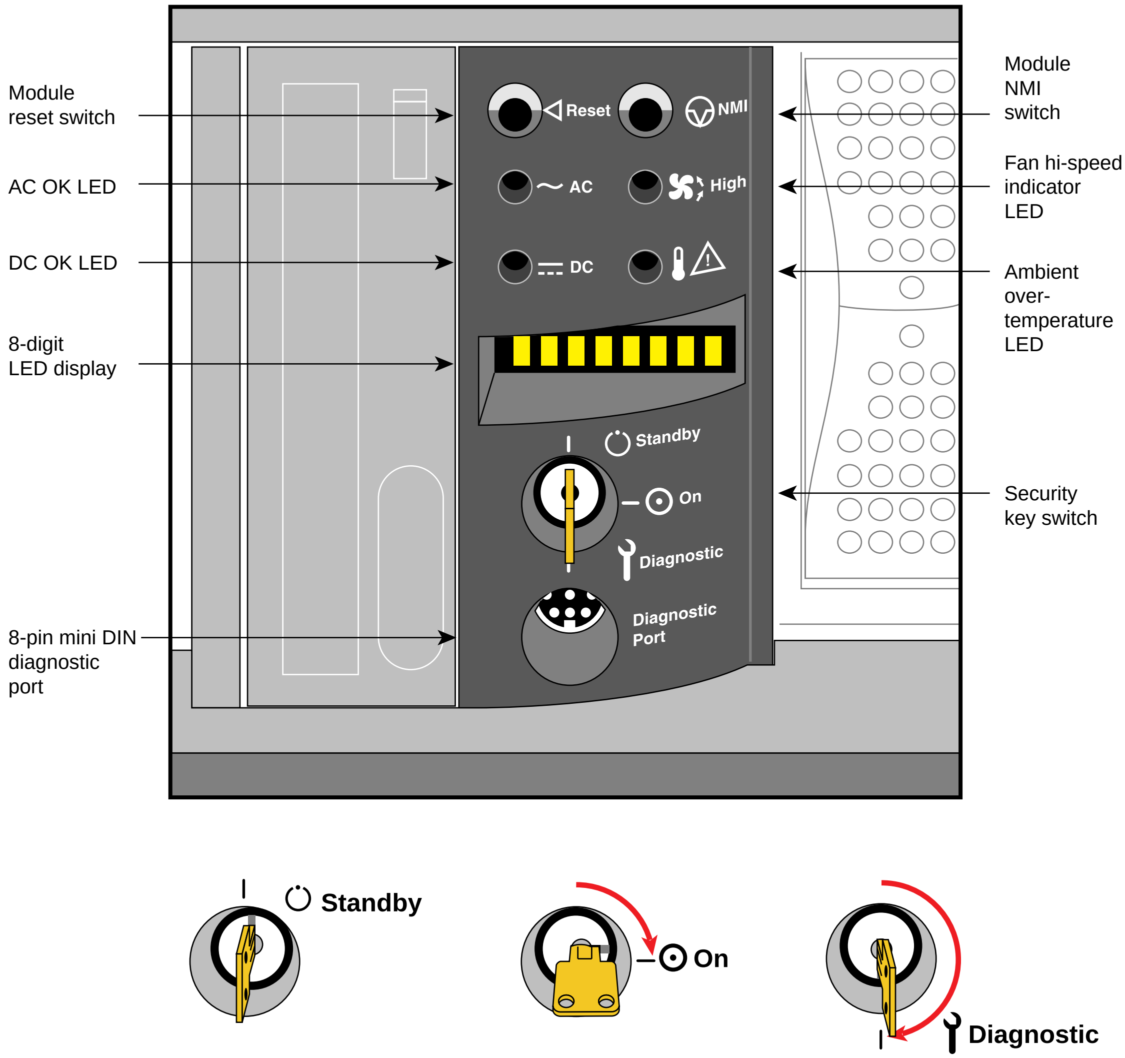

Figure 3-21 MSC (Module System Controller) and Display

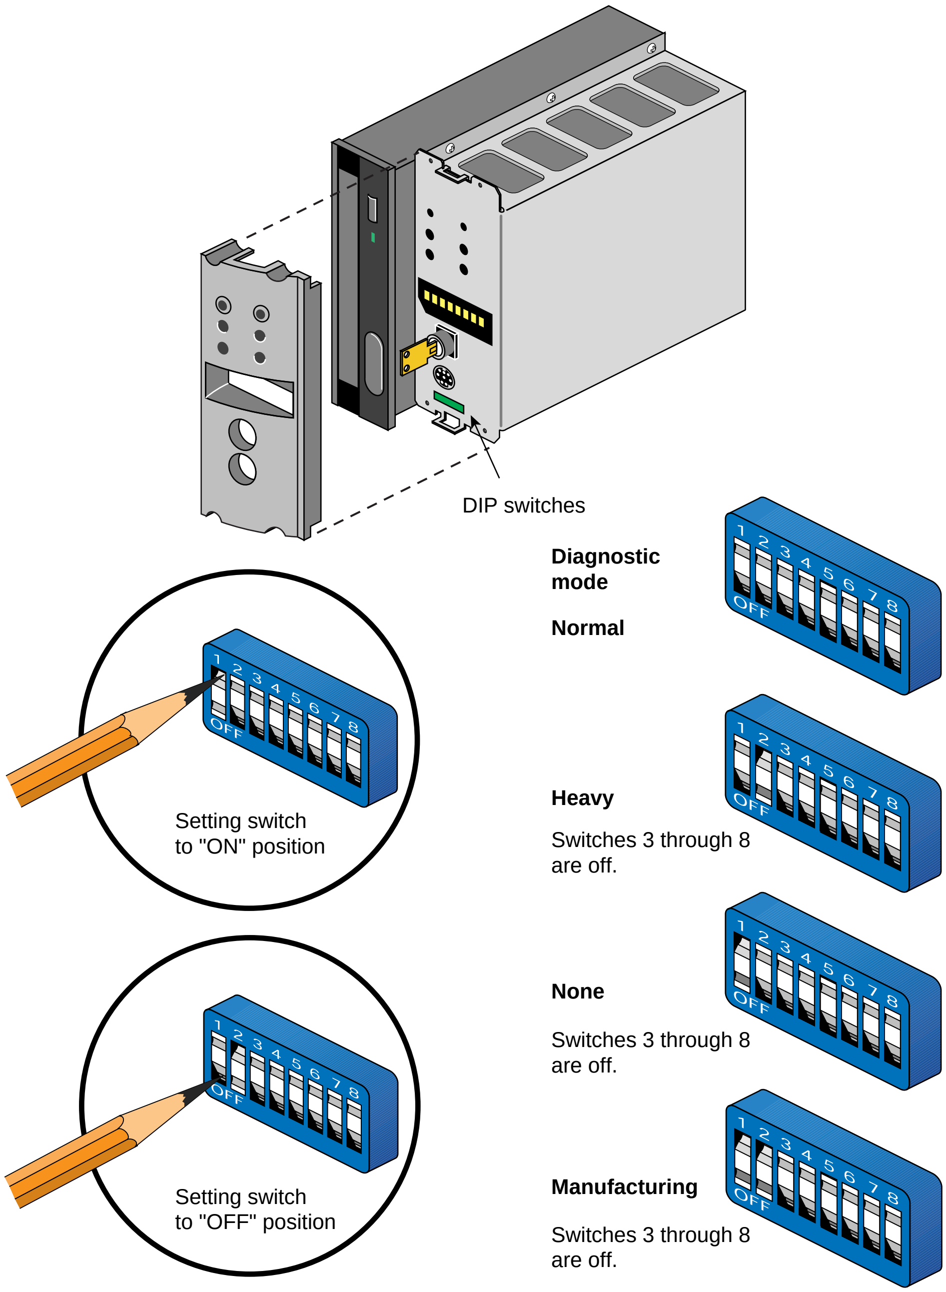

Figure 3-22 MSC Module Controller DIP Switches Location

3.7.2.2 Multimodule System Controller (MMSC)

The MMSC display (see Figure 3-23) is the single-point administration interface for the rackmount configuration. The individual module controllers are tied to the multimodule controller through an 8-pin serial connection from the back of the module chassis to the MMSC board in the rear of the rack (see Figure 3-24). The MMSC is also known as the full-feature System Controller—FFSC.

There should only be one multimodule MMSC display per configuration. The display is a 640 x 480-resolution true-color flat panel (TFP).

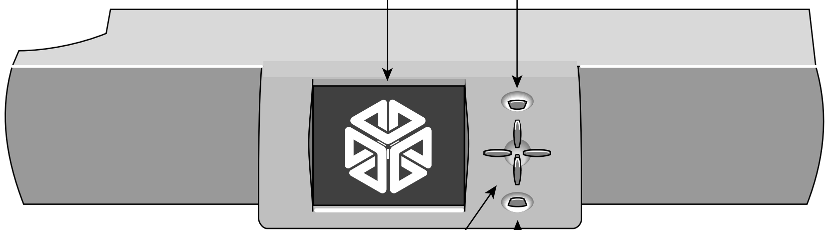

Figure 3-23 MMSC Multimodule Controller and Display

The MMSC has six serial ports, each capable of running at speeds of up to 115,200 bps. The ports are not interchangeable; each has a specific function, as summarized in Table 3-6.

Table 3-6 MMSC Serial Ports

Port	Function	Cabling
COM1	Terminal device or console connects to this port. In a multirack system, this port is typically used only with the system that has the master BaseIO.	Normally requires an 8-pin DIN to 9-pin serial converter, along with a 9-to-25-pin PC-style serial cable.
COM2	Connection to the MSC in the upper bay of the Origin2000 Rack. In a multirack system, this port is used on every MMSC.	Cable provided with system, or standard 9-pin serial to 8-pin mini-DIN null modem cable
COM3	Connection to the MSC in the lower bay of the Origin2000 Rack. In a multirack system, this port is used on every MMSC.	Cable provided with system, or standard 9-pin serial to 8-pin mini-DIN null modem cable.
COM4	Connection to the *master* BaseIO board tty_1 (console port). In a rack configuration with two BaseIO boards (an upper and a lower board) the master BaseIO is typically the one in the lower module. In a multirack system, this port is typically used only with the system that has the master BaseIO board.	Cable provided with system, or standard 9-pin serial to 8-pin mini-DIN null modem cable.
COM5	Remote service modem: not a general-purpose modem port. Also used for direct firmware downloads in emergencies.	Normal 8-pin mini-DIN to 8-pin mini-DIN null modem cable.
COM6	MMSC debugging port; not used in customer systems.	

The standard female 8-pin mini-DIN jack for each port should be labeled with its port number.

The system console (tty1) on the Origin2000 or Onyx2 rack is attached to the BaseIO through the MMSC. Thus the MMSC can intercept certain command strings and can obtain data (such as new firmware images) from either a running UNIX system or from the system console itself.

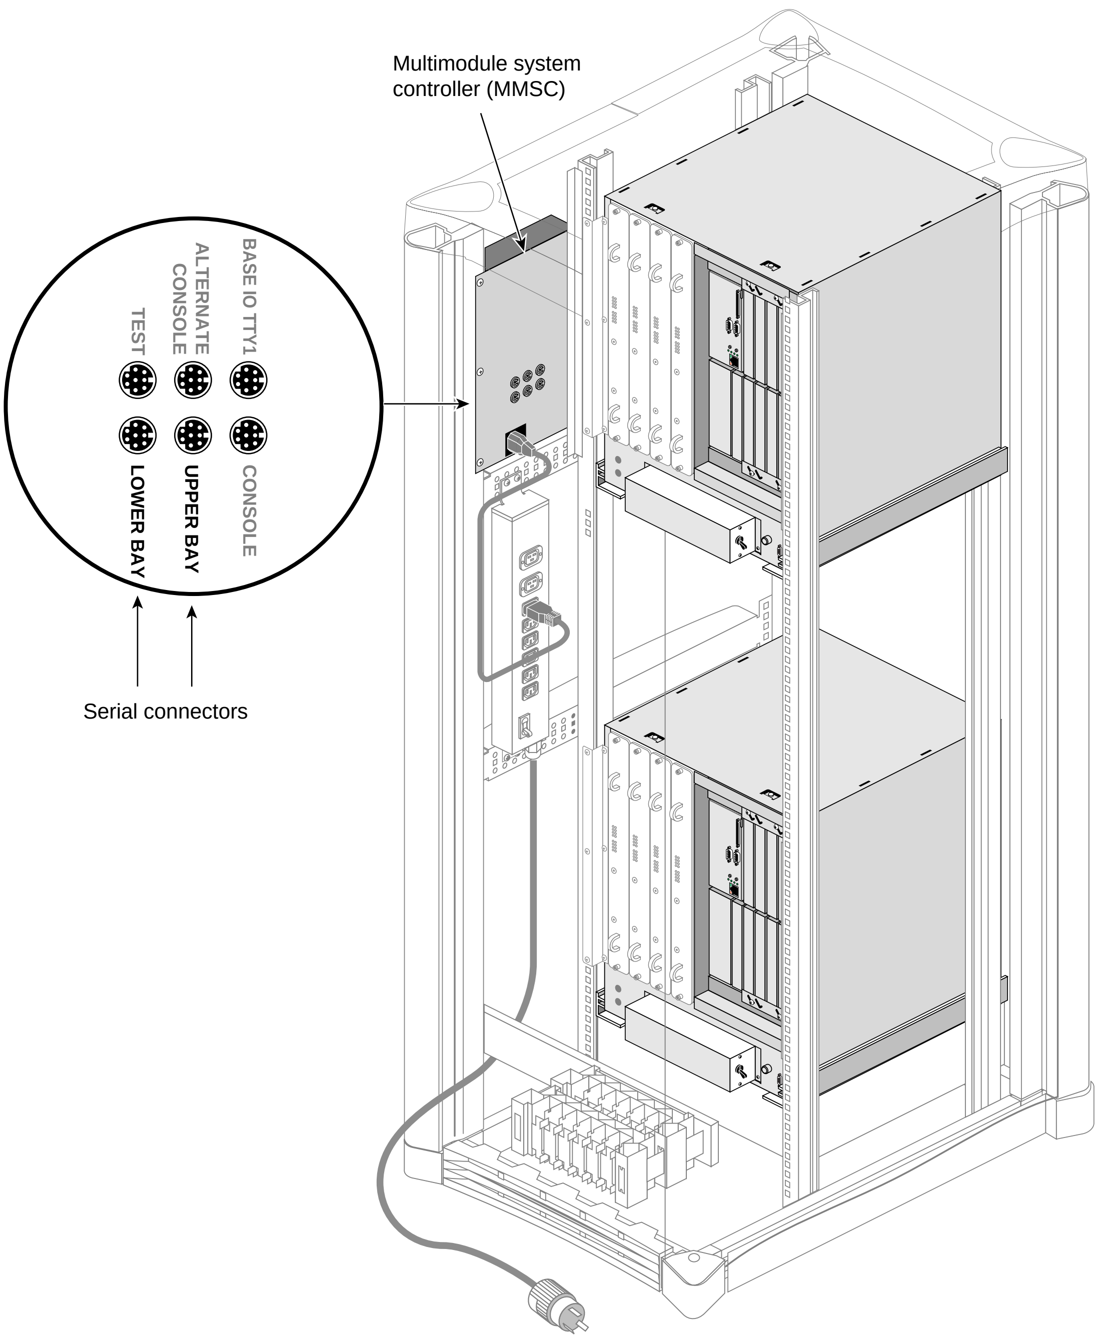

Figure 3-24 MMSC Serial Connectors

3.7.3 BaseIO Board (Server Version)

The BaseIO board is standard in all deskside and rackmount configurations. A deskside or system module cannot have more than one BaseIO board installed at a time. However, a rackmount configuration with two system modules can house two BaseIO boards. Figure 3-25 and Figure 3-26 show different views of the BaseIO board. Table 3-7 describes the connector functions.

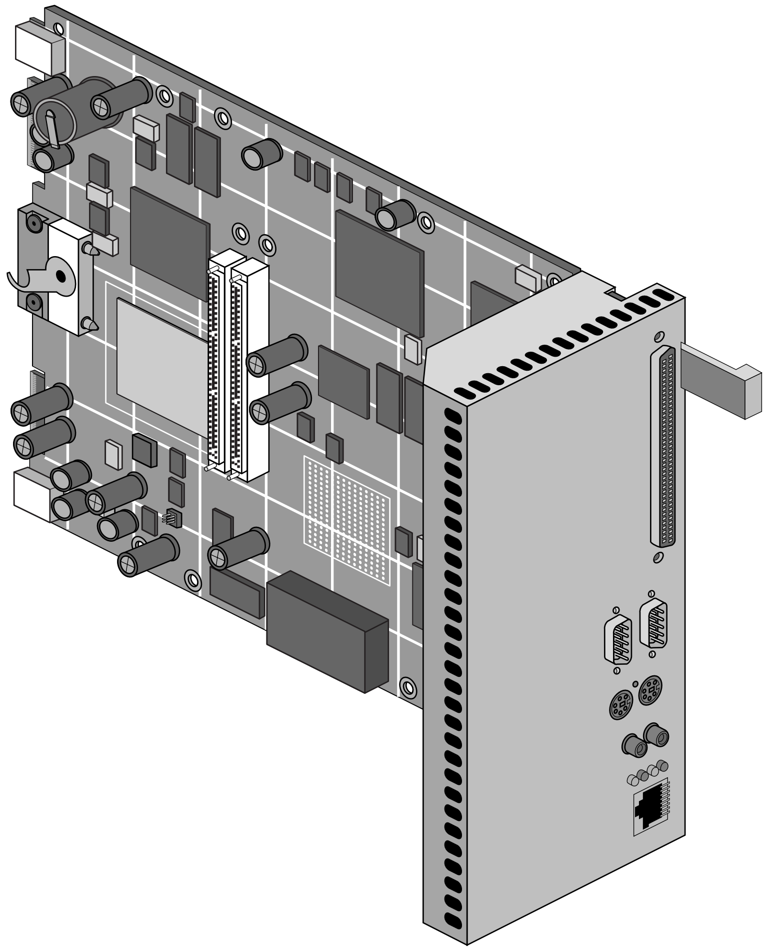

Figure 3-25 Server BaseIO Board

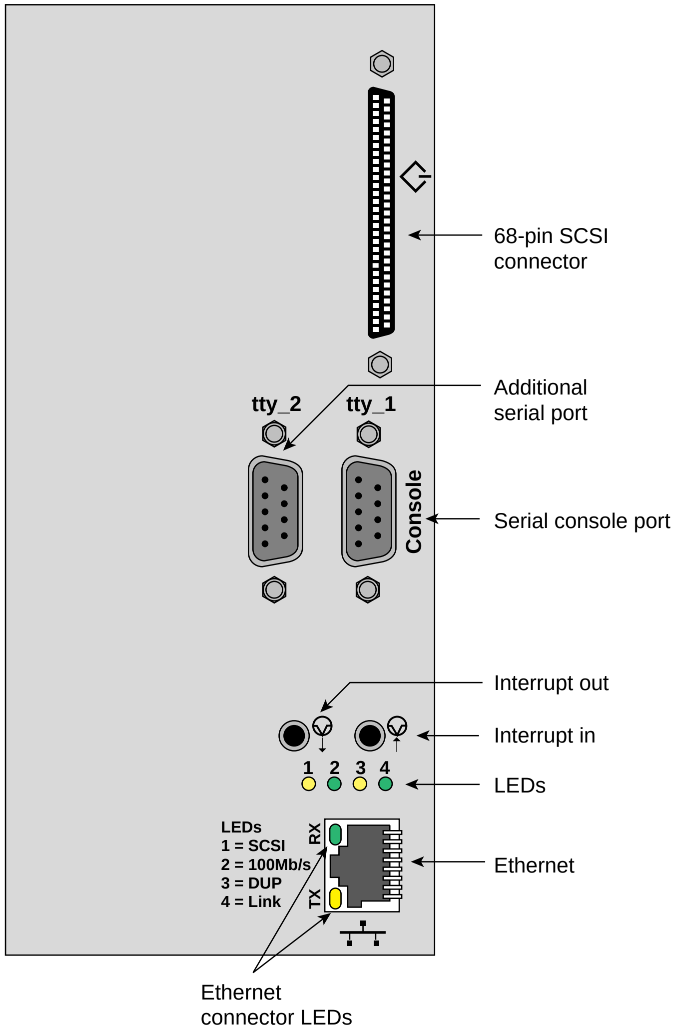

Figure 3-26 Server BaseIO Board Panel

Table 3-7 Server BaseIO Connectors

Connector Description	Connector Type	Function
100-Base T	8-pin jack	100-Mb per second Ethernet
Serial	9-pin DIN	RS-232 and 422 Serial
SCSI	68-pin	Ultra SCSI (Fast-20) single-ended)

3.7.4 BaseIO Board (Graphics Version)

The graphics BaseIO board is standard in all Onyx2 configurations. A deskside or system module cannot have more than one BaseIO board installed at a time. However, a rackmount configuration with two system modules can house two BaseIO boards. Figure 3-27 and Figure 3-28 show different views of the BaseIO board.

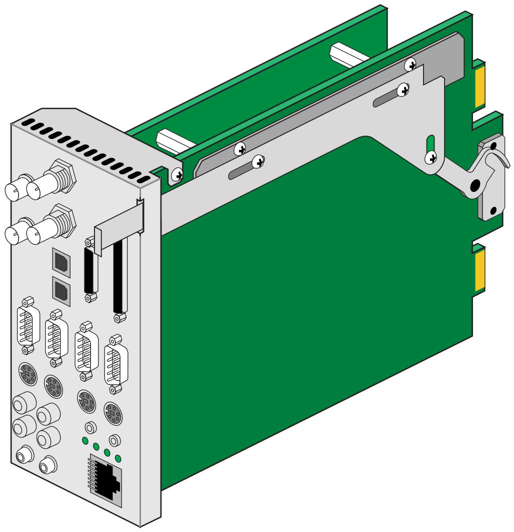

Figure 3-27 Graphics BaseIO Board

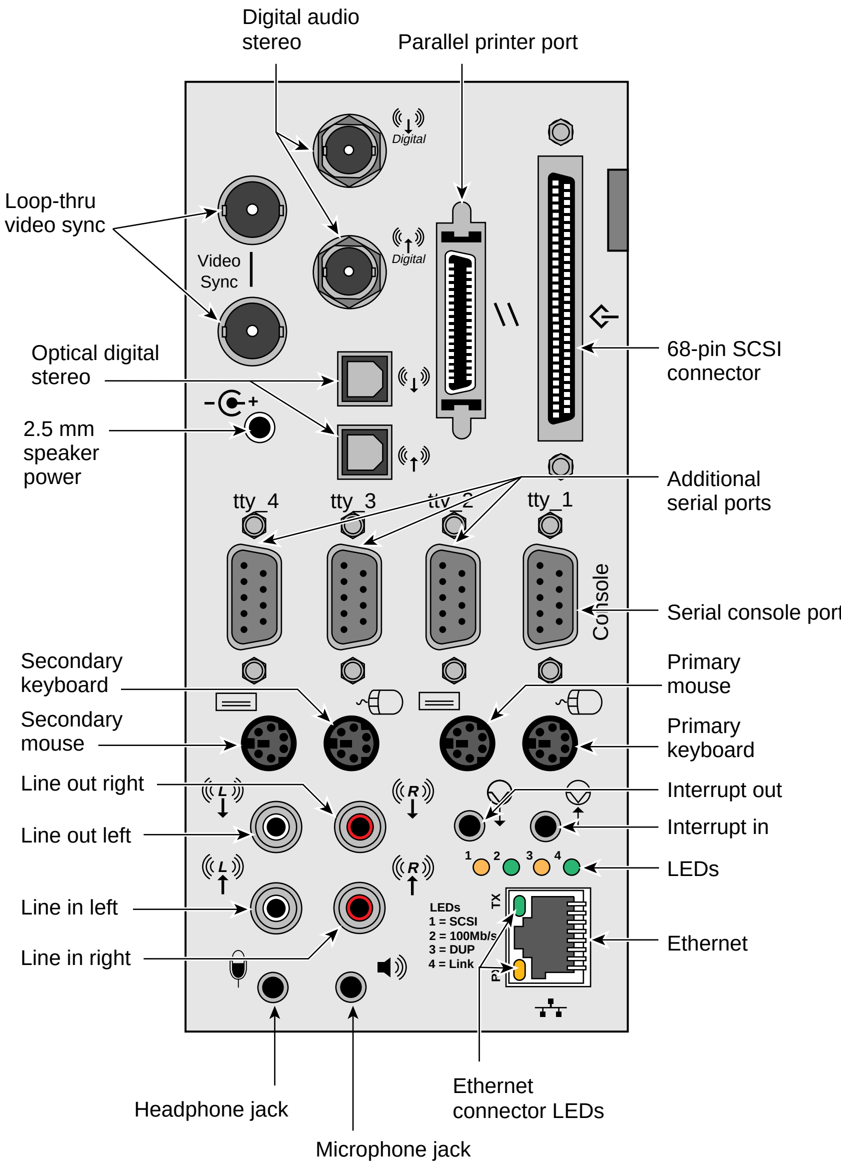

Figure 3-28 Graphics BaseIO Panel

3.7.5 Optional XIO Boards

XIO boards are optional products for Silicon Graphics platforms that are based on the scalable shared-memory multi-processing (SSMP™) architecture. XIO boards are installed into the XIO slots of Origin2000 and Onyx2 systems. Each active XIO slot provides up to 1600 megabytes per second of bidirectional bandwidth (that is, 800 megabytes in each direction) through a nonblocking crossbar switch on the system's midplane or frontplane. Specific XIO products may use a portion or all of this available bandwidth. All XIO slots in a system can be active simultaneously.

Table 3-8 summarizes components and features common to all XIO boards. Figure 3-29 illustrates these basic components.

Table 3-8 XIO Board Common Components

Component	Explanation
Compression connector	Provides communication between the board and the system via the midplane or frontplane.
Hooks on connector	Hold compression connector securely to midplane or frontplane. There is one hook on each side of the compression connector. The hook actuator (next entry in this list) pushes / pulls the hooks into / out-of the locked position.
Hook actuator	Device for moving hooks into and out of their locked position on the midplane.
Screw holes	For attaching a hook actuator to the board.
Panel plate	Provides cutouts for external cables and light-emitting diodes (LEDs).
Thick side of board	The surface of the board that has the compression connector and the tallest components.
Thin side of board	The side of the board with low-profile components.

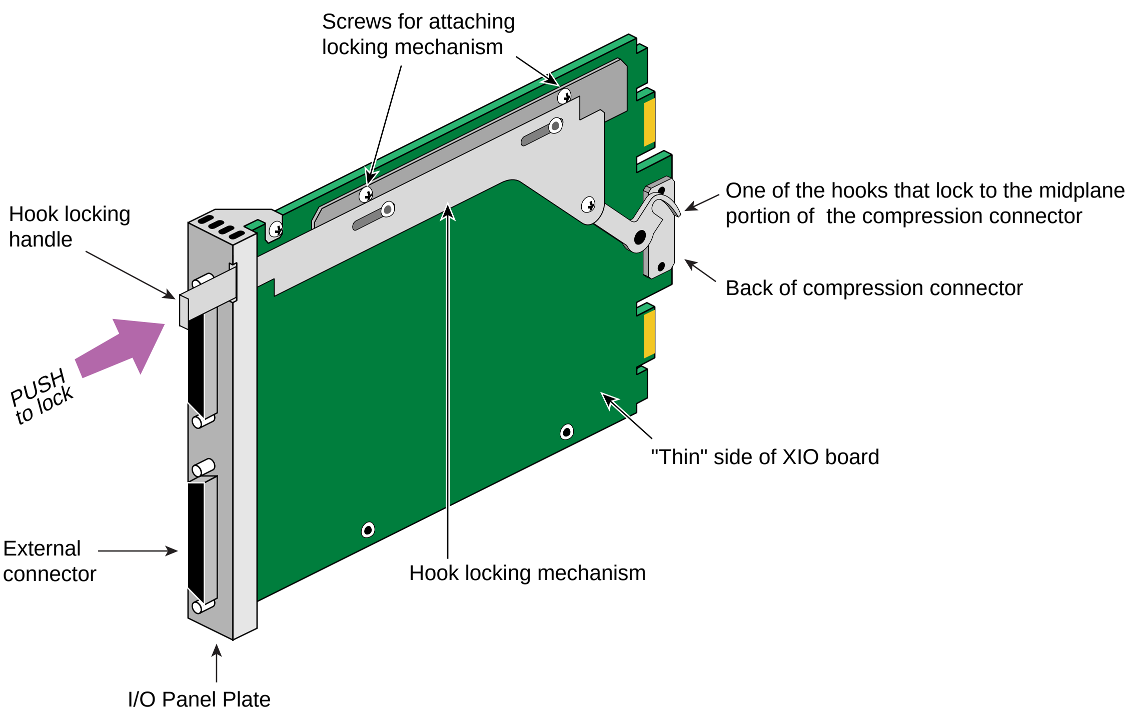

Figure 3-29 Generic XIO Board

3.7.5.1 Hook Actuator

For an XIO board to function, its compression connector must be locked tightly to a mate (other half) on the midplane or frontplane, inside the chassis. The hook actuator is designed to do this.

Each XIO board has two hooks (one on each side of the compression connector). A hook actuator presses against one of the hooks, thus moving the hooks into and out of their locked position. For the Origin2000 and Onyx2 platforms, the hook actuator consists of a horizontally sliding lever and a handle; each type of XIO board may have a unique design for its lever and handle. Figure 3-29 shows one design. These hook actuators are screwed onto the XIO board and attached to one of the hooks, as illustrated in Figure 3-29.

The method of operation is the same for all level and handle designs:

- Pushing the handle locks the hooks and seats the compression connector to the midplane.
- Pulling the handle releases the hooks, in preparation for removing the board.

3.7.5.2 Caring for the Compression Connector

The compression (CPOP) connector used for XIO boards (and router and node boards) has 96 pads that enable passage of signals between the system (via the midplane) and the board. This compression connector has two halves: one half is physically located on the board (illustrated in Figure 3-30); the other half is on the midplane of the chassis. Each pad on a midplane connector is a flat, gold-plated surface. Each pad on a board connector is composed of hundreds of tiny bristles (dendrites). When a bristled pad is pressed into a gold-plated pad, a connection is created for one signal.

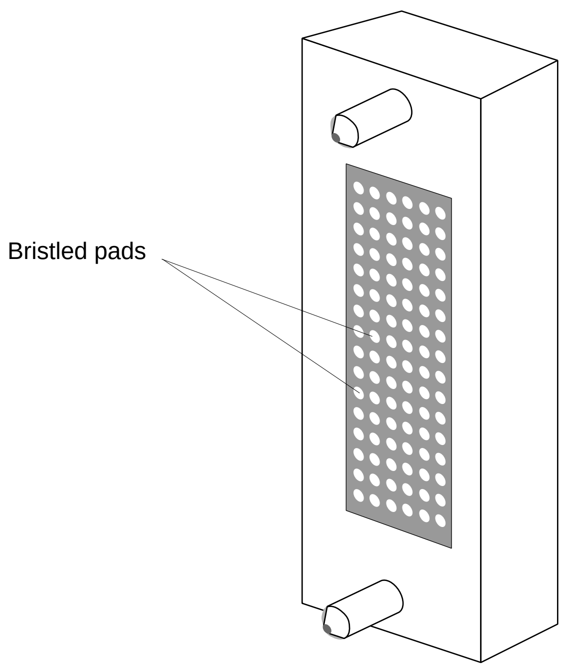

Figure 3-30 The Compression Connector Used on XIO, Node, and Router Boards

The bristled pads can attract and hold dust, lint, grease, powder, and dirt. These substances clog or damage the bristles and prevent them from making proper contact with the gold-plated pads on the system's midplane. It is important to prevent this from occurring. Section 3.7.5.3 explains how to keep the compression connector bristles clean; Section 3.7.5.4 explains how to clean them, in the event that they become dirty.

3.7.5.3 Storing and Handling the Compression Connectors on Boards

To avoid damaging a board's compression connector and to keep it in optimal working condition, follow these guidelines whenever the board is not installed:

Caution: Failure to follow these instructions can result in irreparable damage to the surface of the connector's pads which may result in intermittent or complete failure of the product.

- Do not wipe or touch the pads of the compression connector with anything (no human fingers, no brushes, no cloth, no probes), except as specified in the cleaning instructions. The bristles might be damaged.
- Whenever the board is not in a slot, put the protective cap over the compression connector and store the board in an antistatic bag. Make sure to close (fold over) the open end of the bag in order to minimize exposure to dust and atmospheric gases.
- Do not put anything (not even water) on the pads, except as specified in the cleaning instructions.
- Before laying the board down onto a surface, make sure that the surface is free of dust, lint, powder, metal filings, oil, water, etc.
- Do not blow dust, dirt, or powder anywhere near the board when it is not inside its protective bag.

3.7.5.4 Cleaning the Compression Connector on a Board

A compression connector should never need to be cleaned, if you keep the protective cover on whenever the board is not installed. However, if the connector becomes dirty, follow the instructions below for removing pollutants.

Note: Some pollutants can irreversibly damage (corrode or chemically alter) the pad surfaces. Although cleaning may remove the pollutant, it will not repair damage incurred by this contact.

To remove pollutants, follow these instructions:

1. Obtain a can of dry, compressed inert gas (100% pure nitrogen is recommended). The Envi-ro-tech™ Duster 1671 product manufactured by TECHSPRAY™ (telephone 806-372-8523) works extremely well for this application.

   Caution: Do not use a cleaning product that contains any of the following ingredients: halogenated hydrocarbons, aromatic hydrocarbons, ethers, sulphur, ketones, or solvents of any kind. These substances will cause irreparable damage to the connector's surface.

2. Prepare the can for use, as instructed on the can. For example, if provided, attach the tube to the can's dispensing mechanism.
3. Hold the can in a vertical position.
4. Place or hold the board so that the rounded edge of the compression connector faces up. Note that the rounded edge is completely closed, so that air cannot flow into the connector, whereas the squared edge has an opening.

   Caution: Spraying into the squared (open) edge of the connector can destroy it.

5. Position the board at an angle to the can, so that the tip of the can's applicator is 1 to 2 inches away from the first (topmost) row of pads (as illustrated in Figure 3-31). When you spray, the air will hit each pad and flow downward. Do not allow the applicator to touch the pads.

Figure 3-31 Position for Dry Gaseous Nitrogen Can When Cleaning Compression Connector

6. Start spraying. As you spray, move the spray along the side of the connector until the entire length has been sprayed. Move down a few rows and again spray along the entire length.

   Note: Do not shake the can. Stop spraying if any visible material (for example, foam) appears. This foam will blow away once you resume spraying just air.

7. Repeat until all the pads have been sprayed.
8. When you finish, cover the compression connector with its cap or immediately install the board in an open slot.

3.7.6 Router Boards

The Router board is a multiported, bidirectional data packet controller that can transport up to 800 MBs/sec per port per direction. A Router board is required only when two or more Node boards are present; each system module can have one, two, or no Router boards (depending on the number of Node boards present). Among other functions, Router boards enable two or more Node boards to communicate by providing an interface between the nodes through the midplane. The Router interfaces with the Hubs on Node boards and allows the R10000 processors on one Node board to directly access the main memory located on another Node board. The router boards connect to the Node boards through the midplane as shown in Figure 3-32.

There are three Router board types are described in Table 3-9.

Table 3-9 Origin2000 Router Boards

Router Board Type	Where Used
Null Router (see Figure 3-33)	Used in an Onyx2 or Origin2000 Deskside to directly connect two Node boards. (See Chapter 4, "System Configurations" for additional information.)
Star Router (see Figure 3-34)	Used in a 3- to 4- node Onyx2000 deskside system to enable direct connection to each of the nodes for improved connectivity and throughput. The Star Router is always used with a Rack Router. (See Chapter 4, "System Configurations" for additional information).
Rack Router (see Figure 3-35)	The Rack Router has three external ports for interconnection fabric connectivity.

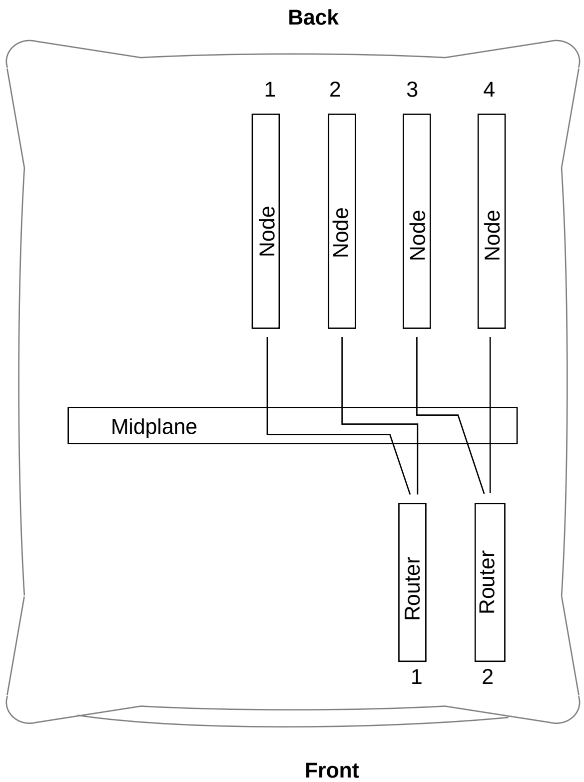

Figure 3-32 Router Board Slot Positions and Connection to Node Boards (Top View)

If a module has only one Node board, no Router is required. A Router board is required only when there are two or more Node boards.

3.7.6.1 Null Router Board

The Null Router board is used in modules with only two Node boards (see Figure 3-33). It provides a low-cost method to connect two Node boards. The Null Router board *cannot* be used for CrayLink Interconnect linking, and it does not have any external router connectors.

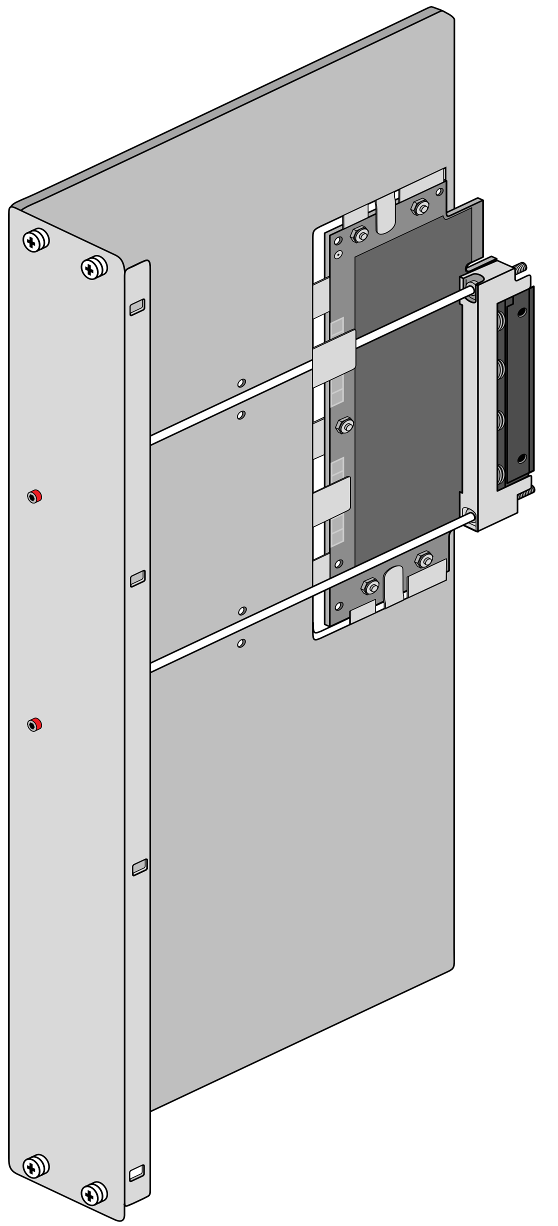

Figure 3-33 Null Router Board

Note: There are no external routing ports on the Null Router board.

3.7.6.2 Star Router Board

The Star Router board is used in a module with three or four Node boards (see Figure 3-34). The Star Router is always paired with a Rack Router board for proper operation. This cost-effective router board provides connections with all the Node boards in a module but cannot be used for chassis external CrayLink Interconnections. The Star Router has one external connector, which connects to a port on the companion Rack Router board through a CrayLink jumper.

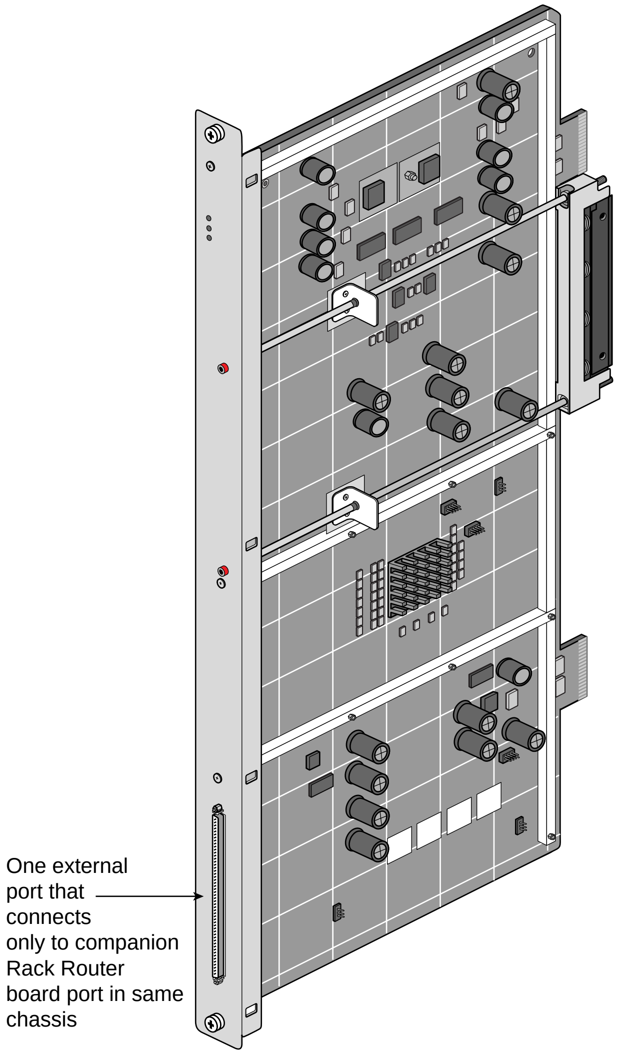

Figure 3-34 Star Router Board

Note: The Star Router board can be used only with a Rack Router board. It cannot be used as a standalone board and is not used in an Onyx2 deskside system.

3.7.6.3 Rack Router

For CrayLink Interconnect cabling, there are normally two Rack Router boards in a system module. The Rack Router (see Figure 3-35) has six ports that route data at up to 800 MBs / sec per port in each direction. Three of the ports connect internally. The fourth, fifth, and sixth ports connect to external router ports on different modules.

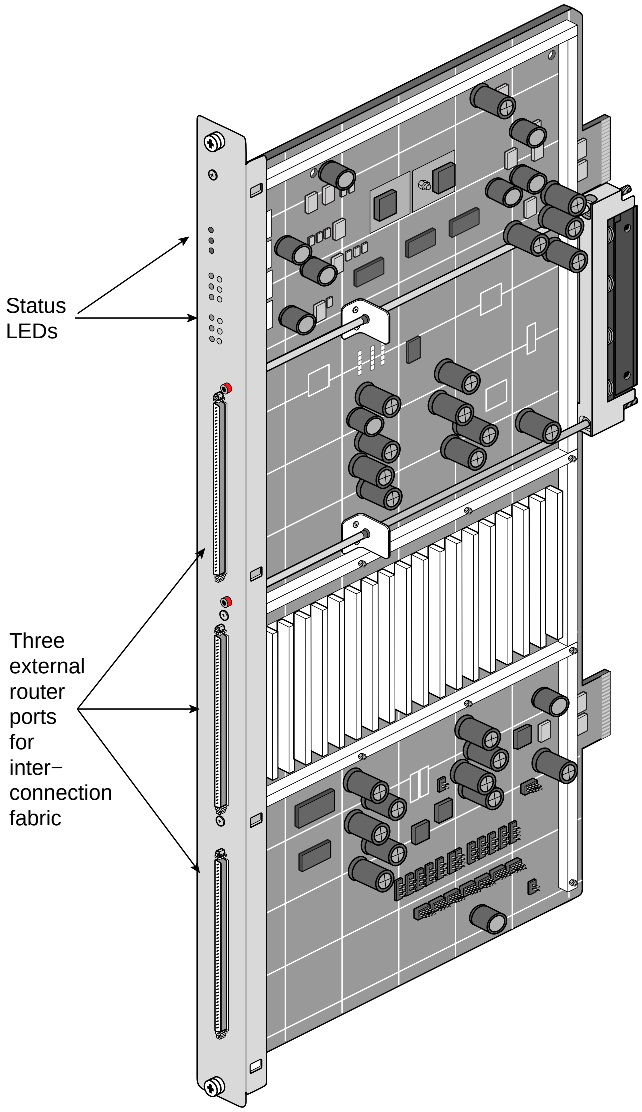

Figure 3-35 Rack Router Board

3.7.6.4 Cray Router

The Cray Router chassis (not shown) is used for CrayLink connectivity for configurations with up to 128 processors. The Cray Router comes in two modules that house additional router boards to provide interconnection fabric support for systems with more than 64 processors.

3.7.6.5 Router LEDs

The Rack Router boards have a set of LEDs on the bulkhead (see Figure 3-35). For an interpretation of the LED states, see Chapter 7, “Diagnostic Tools.”

3.7.7 Midplanes

This section discusses the different midplanes in the Origin2000 and Onyx2 systems.

Note: Check the location of the system NIC (number in a can). This removable battery like part contains the system serial number that uniquely identifies the system, licensing information and other important system information. If you replace the midplane, you must a remove the system NIC from the old midplane and insert it into the replacement midplane.

3.7.7.1 Origin2000 Midplane

Figure 3-36 shows the major parts of the Origin2000 midplane.

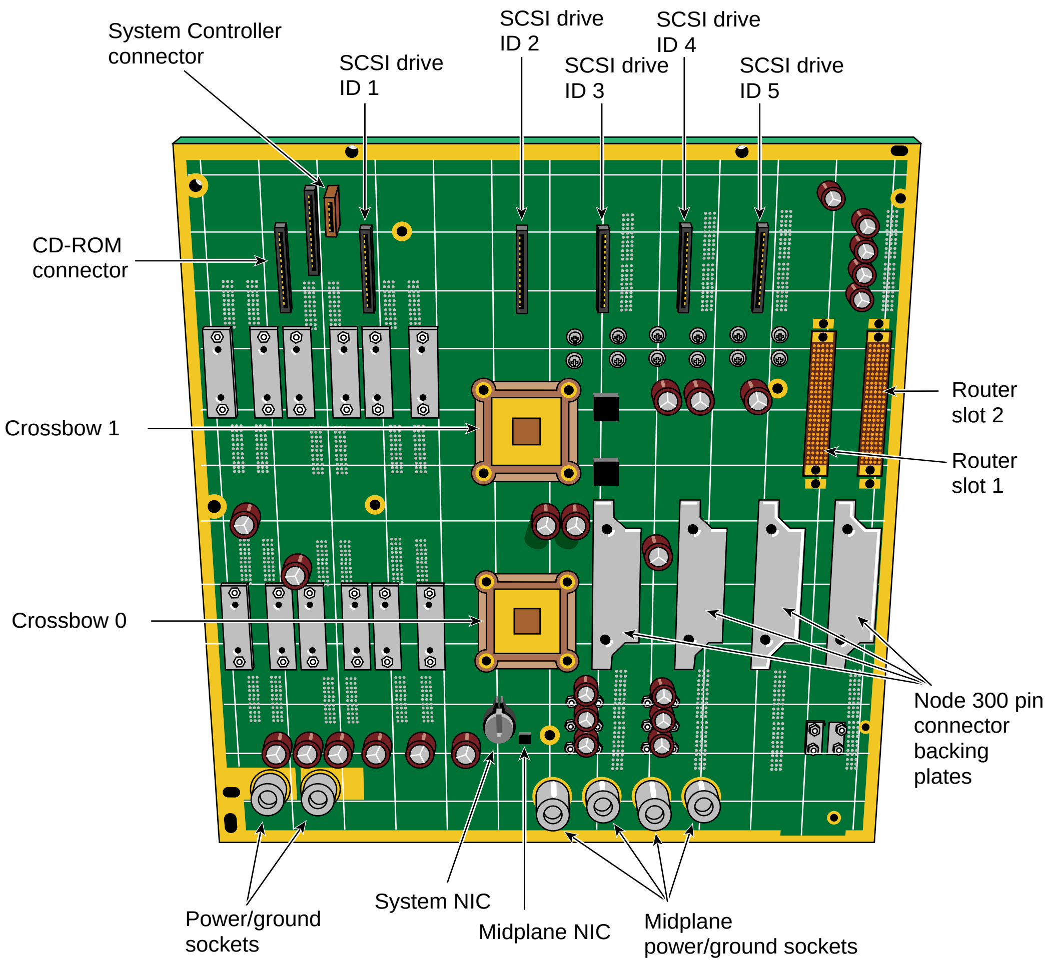

Figure 3-36 Origin2000 Midplane Major Parts

Note: This midplane is also used in the server module of the Onyx2 rack system.

Figure 3-37 shows the major parts of the Onyx2 deskside graphics midplane.

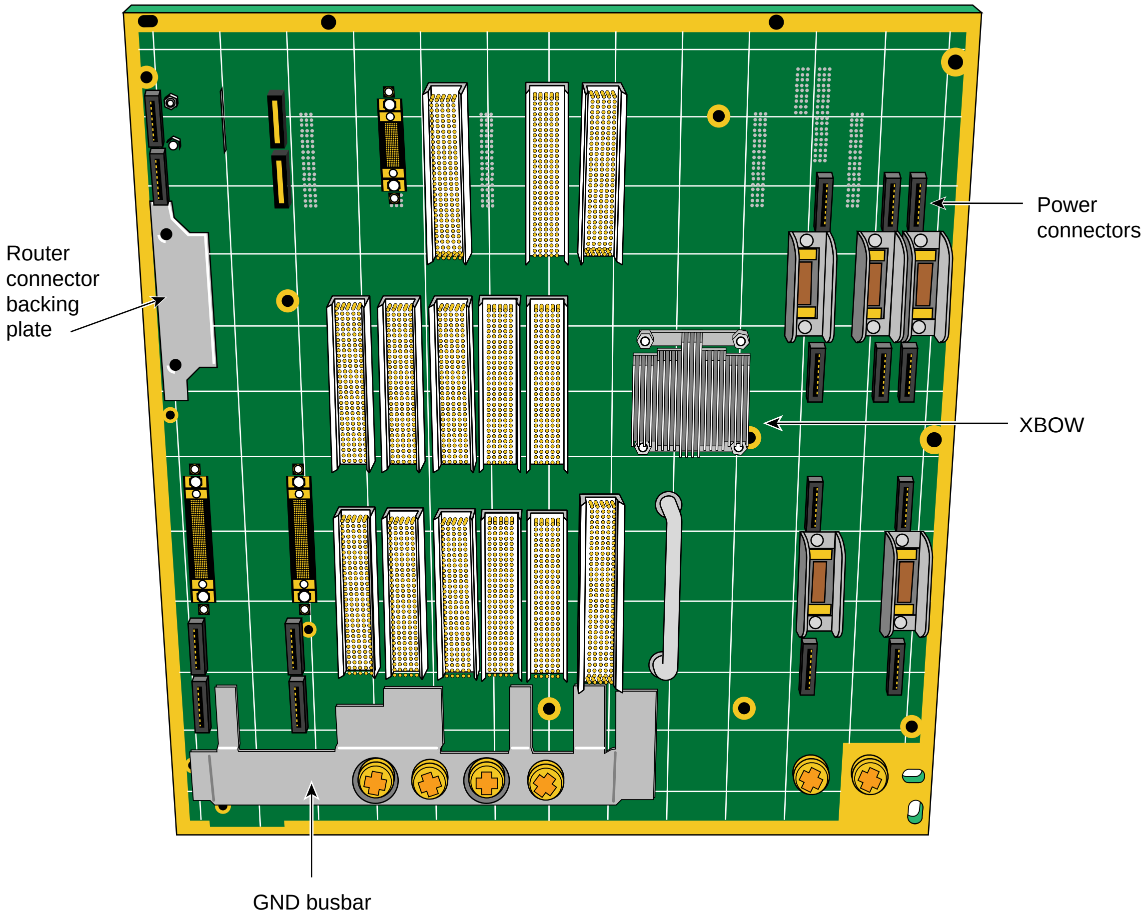

Figure 3-37 Onyx2 Deskside Midplane (Rear of Board)

3.7.7.2 Onyx2 Graphics Module Midplane

Figure 3-38 shows the major parts of the Onyx2 rack graphics midplane.

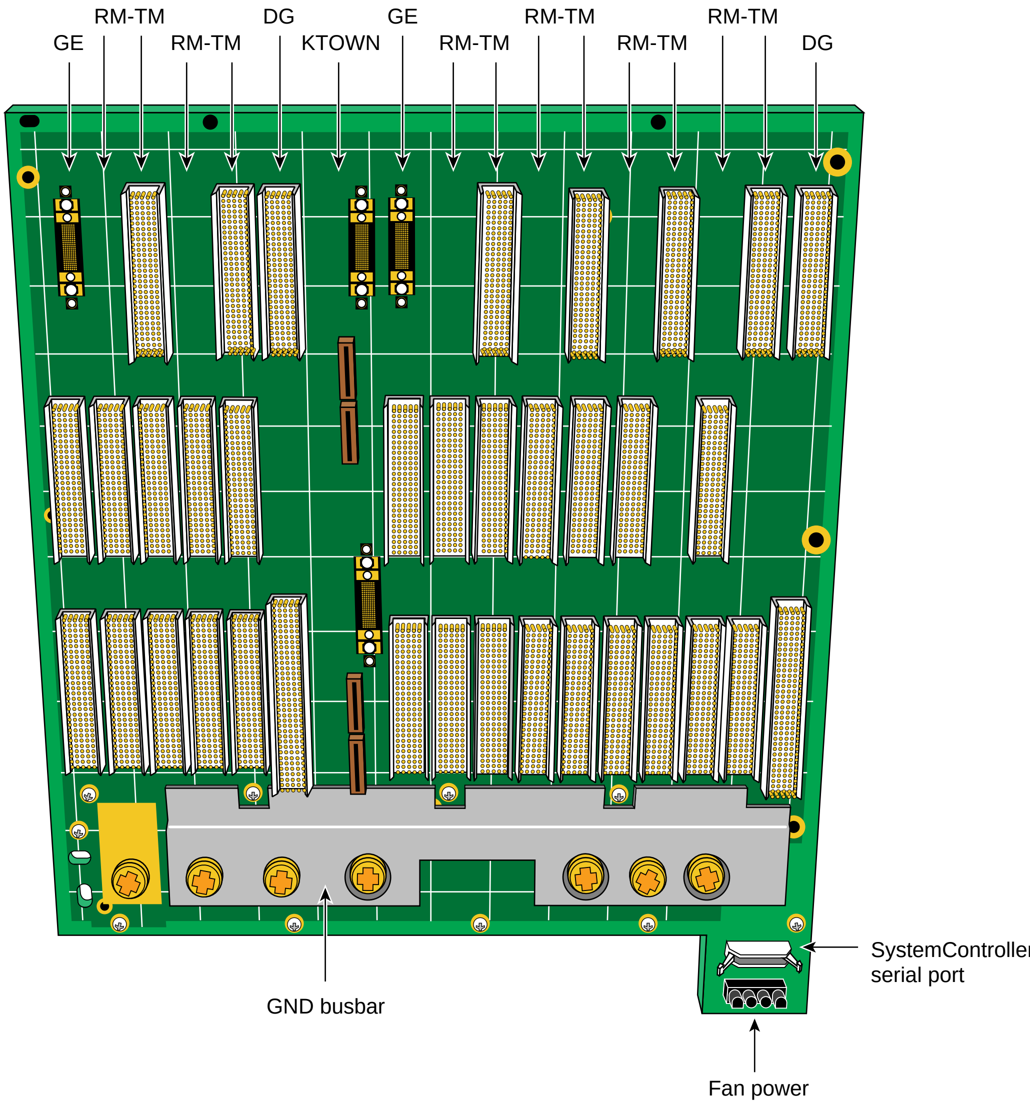

Figure 3-38 Onyx2 Rack Graphics Module Midplane

3.7.8 Power Distribution Unit (PDU)

The PDU supplies 220 volts for the rack and provides a common AC connection source for all the modules and other peripheral devices.

Note: The power supply in the system module is auto-ranging from 110 to 220 volts.

3.7.9 CrayLink Interconnect, Xpress Link, and Crosstown Cables

The CrayLink Interconnect and the Xpress links are the cabled interface that runs between Rack Routers boards (see Figure 3-39). These cables provide a high-speed (1600 MBs / sec), scalable interconnection between different modules. These cables also supply physical link redundancy so that if a link fails, another link can take its place.

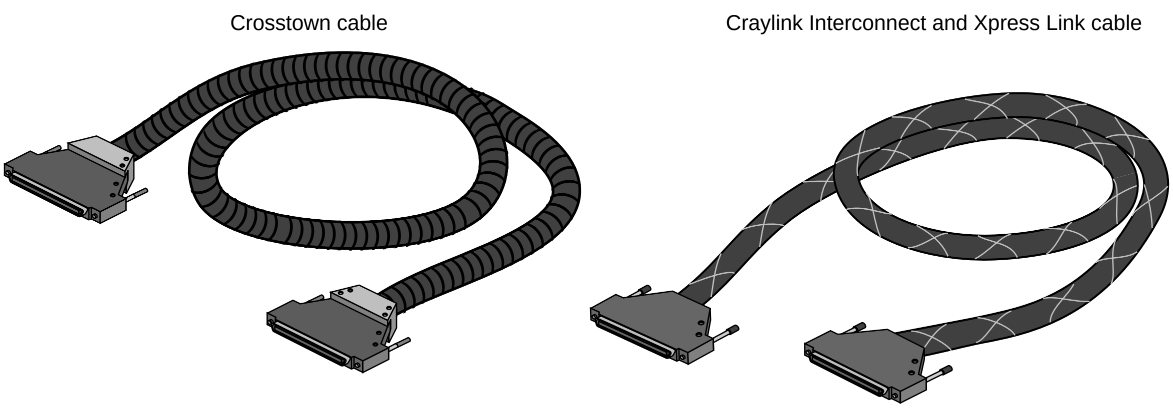

Figure 3-39 Crosstown, CrayLink Interconnect, and Xpress Link Cables

3.7.9.1 Cautionary Guidelines

CrayLink Interconnect and Xpress link cables are delicate. Observe the following cable handling guidelines:

- Avoid bending the cables more than a 1.25-inch radius.
- Avoid stepping on the cables.
- While both Origin2000 modules are up and running, avoid “hot plugging” in or removing cables. Either can *hang* or *crash* the entire system configuration.

3.7.9.2 Xpress Links

The Xpress links provide greater system bandwidth, as they supply additional 800-MBs/sec routing paths. The Xpress links connect to Router board ports that are not being used by the CrayLink Interconnect cables. For additional information on the use of Xpress links, see Chapter 4, “System Configurations.”

3.7.9.3 Crosstown Cable

The Crosstown cables are similar to the CrayLink Interconnect and Xpress cables, so the same cautionary guidelines apply. The cables are used to connect to external peripherals and provide the same I/O performance as the CrayLink Interconnect and Xpress link cables.

3.7.10 PCI Adapter

The peripheral connector interface (PCI) adapter can house up to three third-party PCI boards (see Figure 3-40). Two of the PCI boards can be full-size boards. The third PCI slot accepts only half-size boards. The PCI adapter provides up to 75 watts of power, approximately 25 watts per slot.

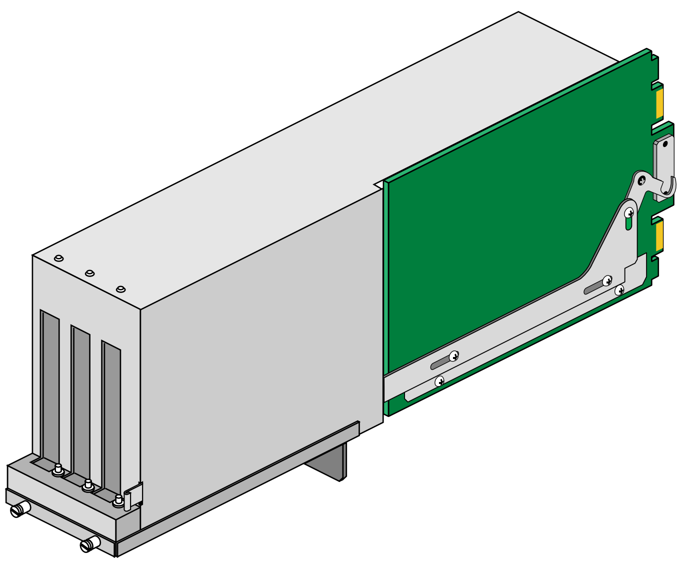

Figure 3-40 Optional PCI Adapter

3.7.11 Remote Serial Connector

This connector (see Figure 3-41) can serve two major functions. It can provide a modem connection linkup for a remote monitoring purposes. It can also connect to a multimodule System Controller in a rack system so that an MMSC can monitor and control various system-level activities such as fan speed, CPU activity, and operational status.

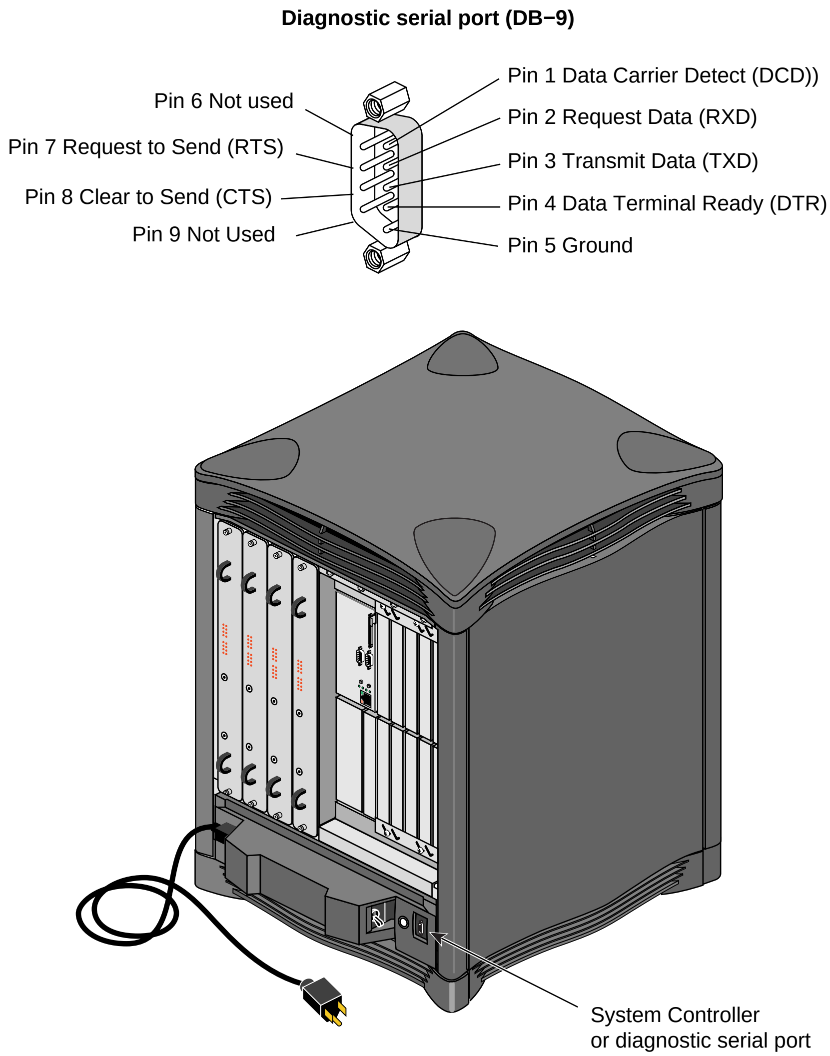

Figure 3-41 Remote Serial Connector

3.7.12 DG5 Display Generator Board

There are two versions of the DG5 board: one with two channel outputs (see Figure 3-42) and a second with eight channel outputs (see Figure 3-43).

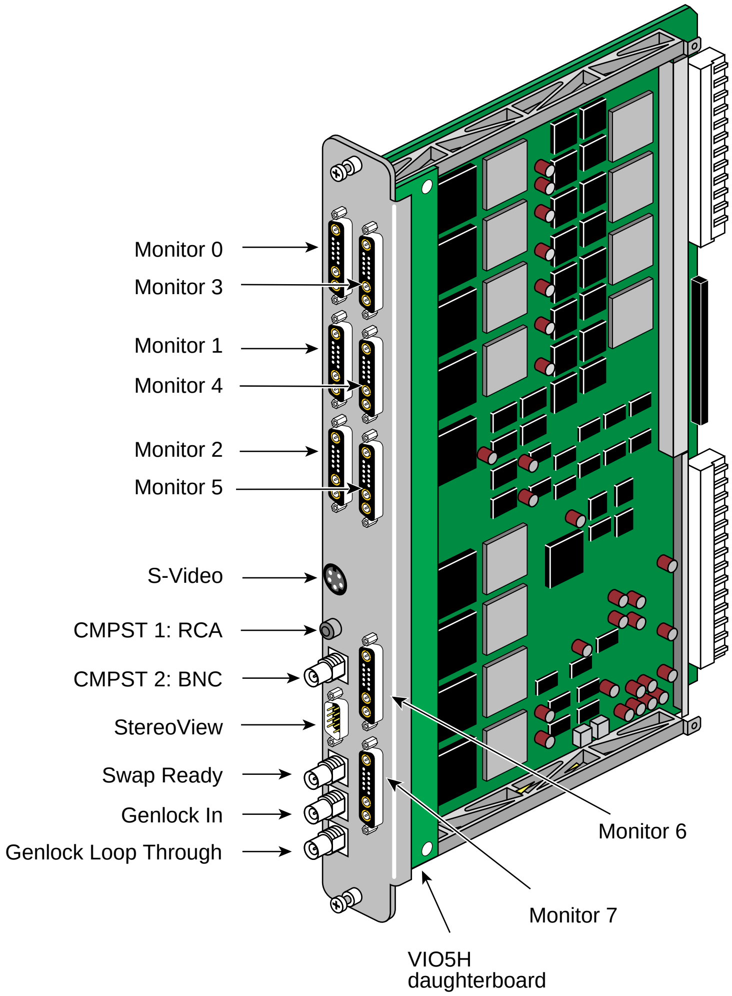

Figure 3-42 DG5-8 / GVO Display Generator Board

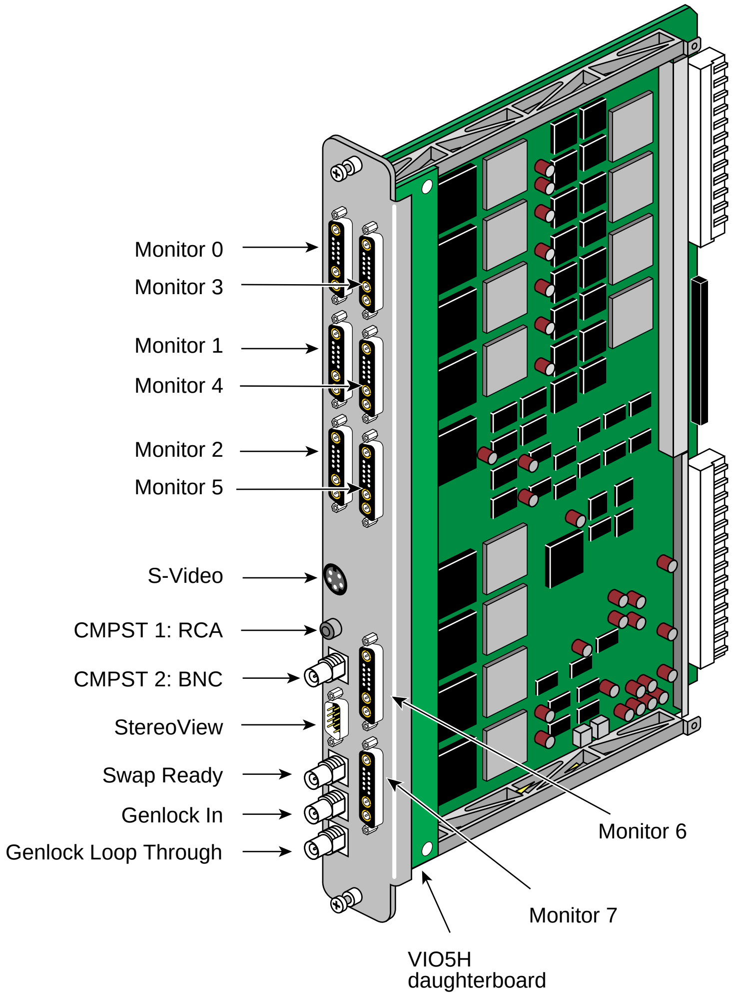

Figure 3-43 DG5-8 Display Generator Board

3.7.13 Other Reality and InfiniteReality Graphics Boards

Figure 3-44 shows the other boards (GE and RM / TM) that make up the Reality and InfiniteReality board set. The GE-14 board and Ktown board (which is used only in Onyx2 rack systems) install using compression connectors. The external connector on the RM / TM board is for future development. The TM is a daughterboard for the RM board.

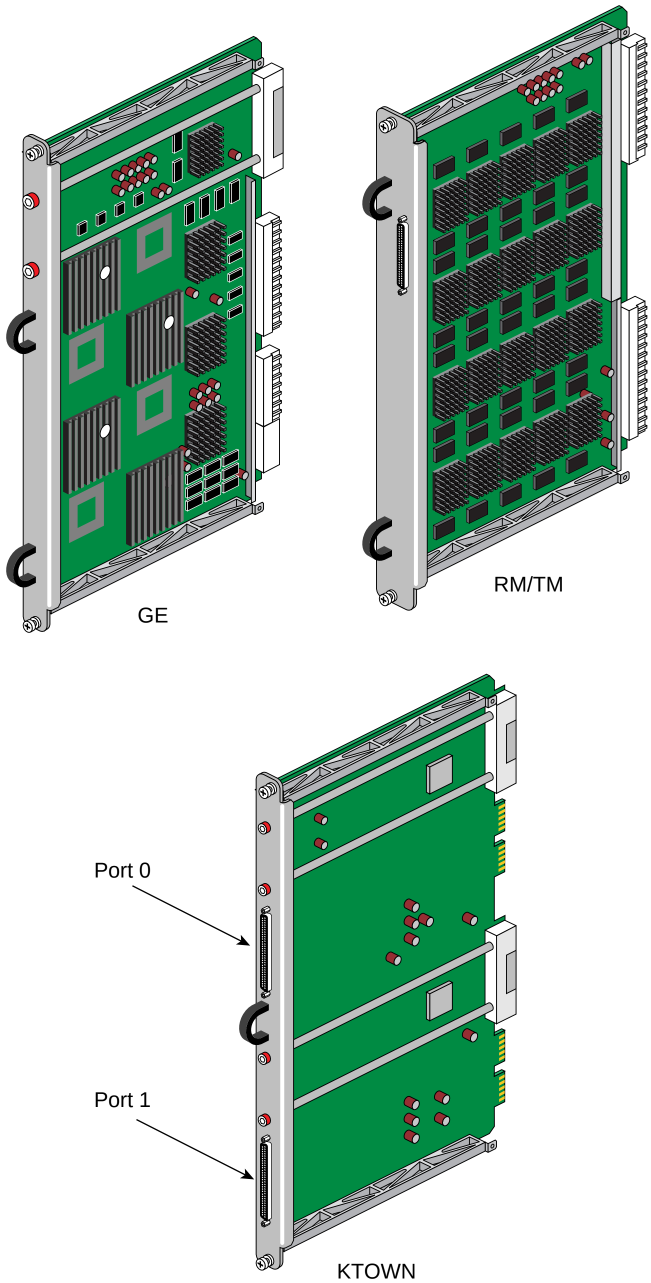

Figure 3-44 Other Reality and InfiniteReality Graphics Boards

Note: The Ktown board is present only in the Onyx2 rackmount systems.

3.7.14 24-Inch Monitor

Figure 3-45 shows the superwide, 1,920 x 1,600 monitor that has been developed with Sony® to provide superior real-time images for the Onyx2 InfiniteReality graphics set. When used with the Reality graphics set, this monitor has a limited set of video formats.

See the *InfiniteReality Video Format Combiner User's Guide* (P / N 007-3279-xxx) for additional information on available formats.

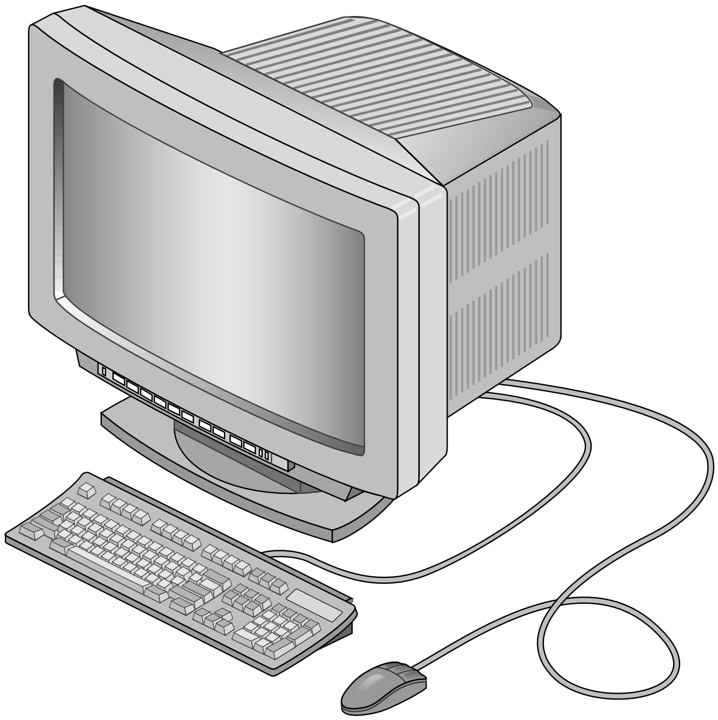

Figure 3-45 24-Inch Monitor for InfiniteReality Graphics

3.7.15 Serial Cables

The Origin2000 and Onyx2 systems use standard PC-style serial cables for connection to an ASCII monitor and for cabling the MMSC and MSC. You cannot use earlier Silicon Graphics serial cables. The DIN-8 is identical to the Indy DIN-8. The DB9 is PC standard; standard PC serial cables work with Origin2000 and Onyx2 systems. See Appendix B, "Connector Pinouts."

To connect the BaseIO to a laptop, dumb terminal, or Indy, obtain

- female-female null modem cable
- female-to-male gender changer (male-male)
- 9-pin to DIN converter (for Indy)

Chapter 4

System Configurations

There are three major Origin2000 and Onyx2 configurations: the deskside, rackmount, and multirack system. Most of these configurations contain the following major common hardware:

- IP27 Node board with one or two R10000 processors
- 64 MBs to 4 GBs of main memory (per Node board)
- a Router board
- a BaseIO board (that provides the system I/O ports such as serial, Ethernet, and SCSI interfaces)
- CD-ROM and System Controller display assembly
- up to five SCSI drives per module

Additional system modules and individual hardware components can be added to suit growing computational and I/O requirements.

Note: Configurations such as a 4-processor or a 16-processor system are sometimes referred to as a “4P” or “16P” system. The *P* stands for the R10000 *p*rocessor. For example, an 8P system has eight R10000 processors.

4.1 Basic System Configurations

This section provides a brief overview of the many different configurations that are available for the Origin2000 and Onyx2 systems.

4.1.1 Origin2000 Server Configurations

Figure 4-1 and Figure 4-2 show various system configurations based on the number of processors in the system. These figures provide information on the type of Router board(s) that are required to support that configuration. The diagrams also show how the Origin2000 system expands from a basic building block module to a multimodule, multirack configuration.

As Figure 4-1 shows, a base module, known as the Origin2000 *deskside chassis* supports one to eight processors. To go beyond eight processors, a rackmount system is needed.

Figure 4-2 shows that configurations with more than 16 processors or 8 Node boards require at least two rack systems. Larger configurations with 32 processors or 16 Node boards require four racks. The largest configurations, with greater than 64 processors (up to 128), not only require even more racks but also a different type of Router hardware called a *Cray Router* (see Figure 4-2).

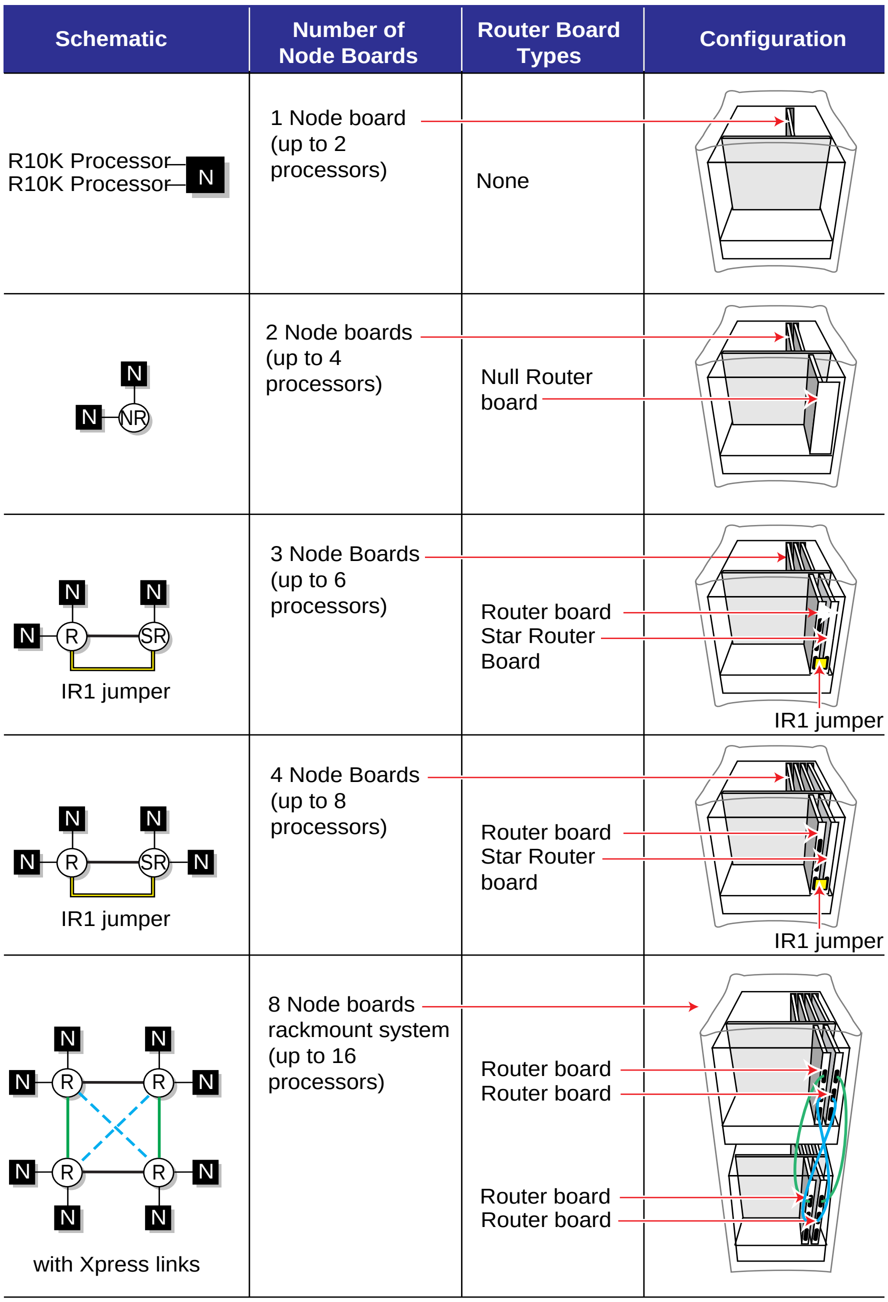

N = Node board
NR = Null Router board
SR = Star Router board
R = Router board

Figure 4-1 Origin2000 System Configurations and Router Board Type Used

Schematic	Number of Node Boards	Configuration
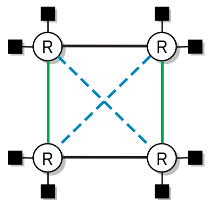 With Express Links	8 Node Boards (up to 16 processors) Rackmount System	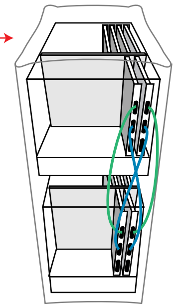
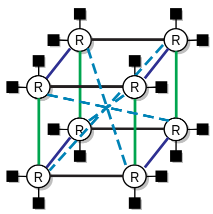 With Express Links	16 Node Boards (up to 32 processors) Multi Rackmount System	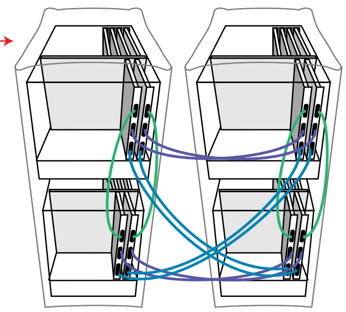
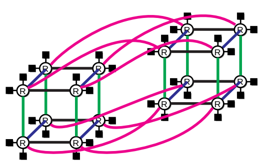	32 Node Boards (up to 64 processors) Multi Rackmount System	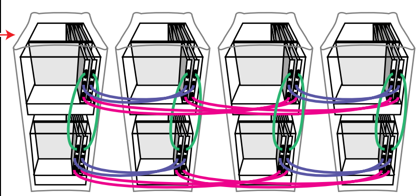
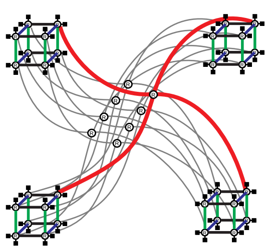	64 Node Boards (up to 128 processors) Multi Rackmount System with a Meta Router	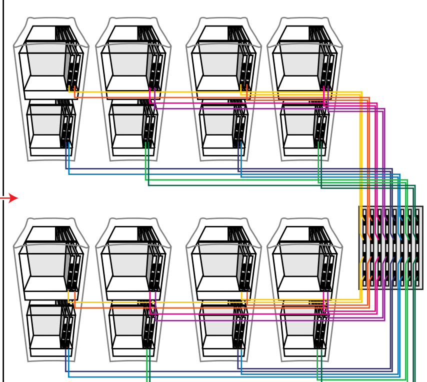 Cray Router Only Meta Router connections shown

Figure 4-2 Different Origin2000 Rackmount System Configurations

4.1.2 Onyx2 Graphics Configurations

Figure 4-3 and Figure 4-4 show various graphics system configurations based on the number of processors in the system and the type of graphics that are installed. These figures also provide information on the type of Router board(s) that is required to support that configuration. In addition, the diagram shows the Onyx2 system expanding from a deskside system to a multimodule, multirack configuration.

As Figure 4-3 shows, the Onyx2 deskside supports one to four processors. To go beyond eight processors, a rackmount system is needed. The bottom of Figure 4-3 illustrates an eight-processor system with a graphics subsystem—the Onyx2 rackmount system. Figure 4-4 illustrates various Onyx2 multirack configurations.

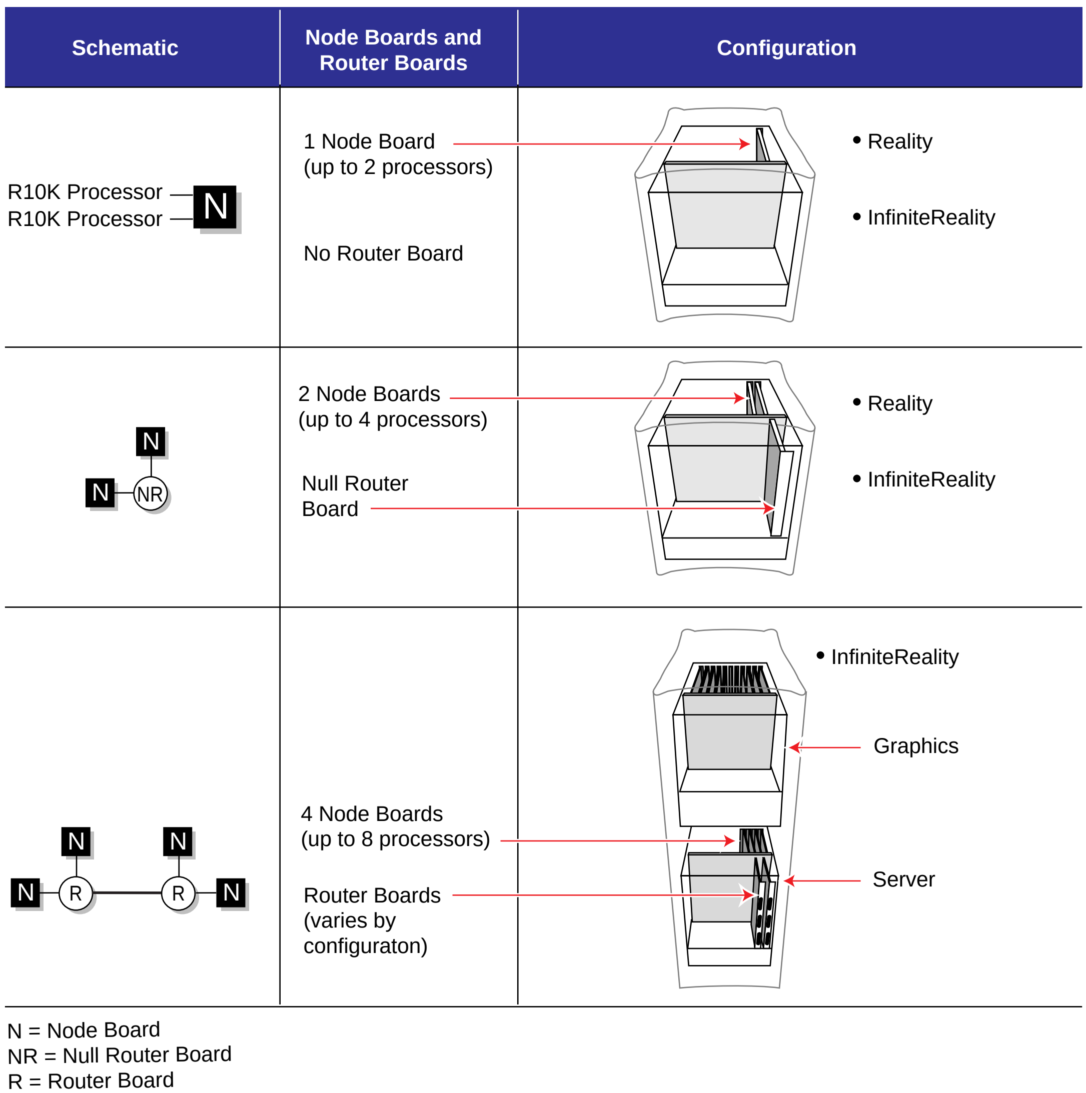

Figure 4-3 Onyx2 System Configurations (Part 1)

Schematic	Node Boards and Router Boards	Configuration
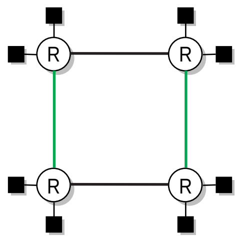	8 Node Boards (up to 16 processors) Router Boards Router Boards	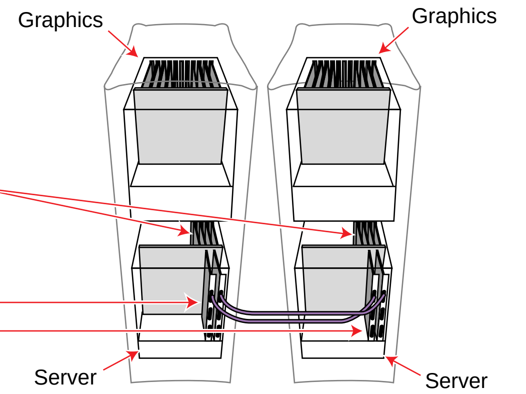
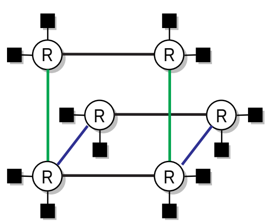	12 Node Boards (up to 24 processors) Router Boards Router Boards Router Boards	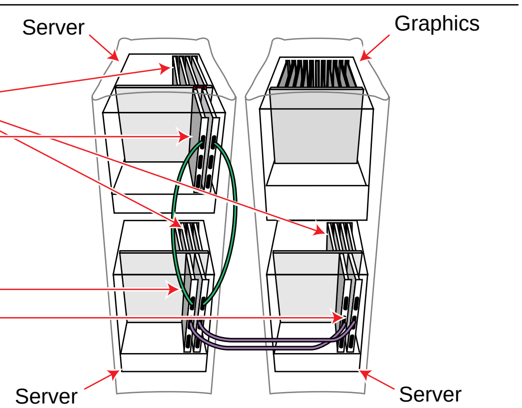

Figure 4-4 Onyx2 System Configurations (Part 2)

4.2 How the Origin2000 and Onyx2 Systems Expand

The CrayLink Interconnect building block is shown schematically in Figure 4-5. The Router board and Node board connection is represented by a Router "ball" attached to two Node boards, each with two R10000 processors (four total R10000s). This represents a 4*P* deskside configuration.

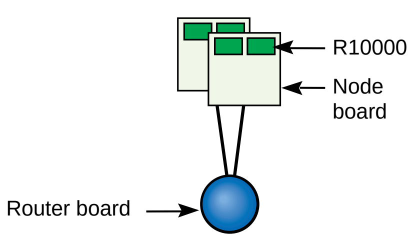

Figure 4-5 CrayIInk Interconnect Building Block (4P Configuration)

Figure 4-6 doubles this configuration with four additional processors and a second router board. This is the largest possible deskside configuration (four Node boards and two router boards). To simplify the schematic representation, a single ball represents a 4P configuration.

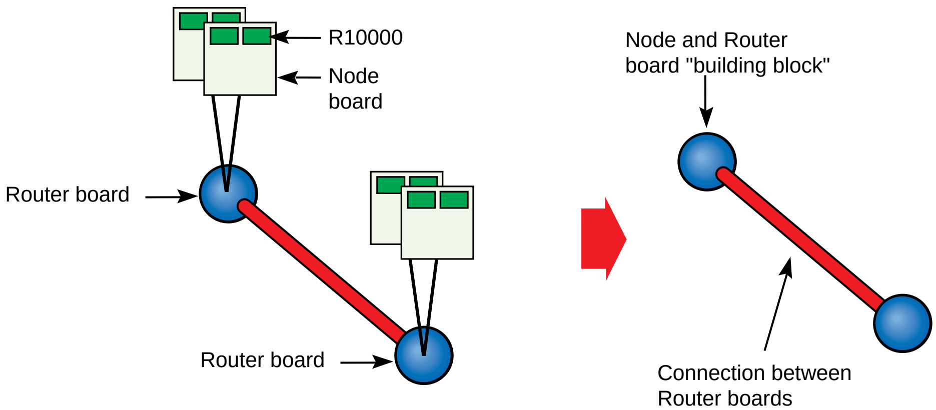

Figure 4-6 Router-to-Router Connection (8P Configuration)

Figure 4-7 represents a 16P rackmount system. R1 through R4 represent the four Router boards and are shown schematically as a square. CrayLink Interconnect cabling has been added to show how the modules connect. Note how the top ports of each router board (port 1) connect to one another. These are *corresponding* ports. To help guarantee proper communication lines, you need to install CrayLink Interconnect cables between the proper corresponding ports. Additional configurations are described later in this chapter.

Schematic Diagram	Configuration

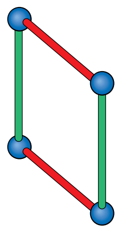

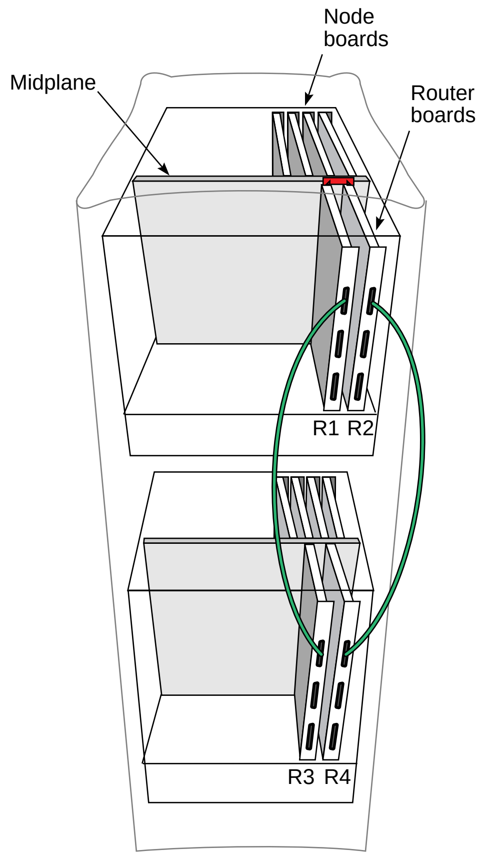

Figure 4-7 16P Rackmount System Configuration

4.3 Origin2000 Deskside Configurations

The Origin2000 deskside (see Figure 4-8 through Figure 4-10) uses a modular, 19-inch chassis. Table 4-1 summarizes the range of hardware components that can be supported by the deskside system.

The deskside can be used as a standalone chassis or be modified for mounting in a Silicon Graphics Origin2000 rackmount system or in any standard 19-inch rack unit. In addition, the Origin2000deskside has from one to eight processors and also contains up to 12 I XIO slots to house additional Silicon Graphics network or peripheral support boards. The Origin2000 deskside includes an entry-level System Controller module for basic system control and administration functions and five SCA, single-ended Ultra-SCSI disk drive slots.

Table 4-1 Single Module Chassis Configurations

Hardware Component	Range
IP27 Node boards	1 to 4
R10000 CPUs	1 to 8
Main memory per Node board	64 MBs to 4 GBs (using 256-MB memory DIMMs)
Main memory per chassis	64 MBs to 16 GBs (using four Node boards)
Number of internal SCA drives	1 to 5
Number of XIO slots	6 or 12 for a single-module chassis
I/O interfaces	- One (single-ended) SCSI connector - Two 9-pin D-sub RS-232 or RS-422 serial connectors - Two hardware interrupt lines - One 10/100 Base-T Ethernet connector

Figure 4-8, Figure 4-9, and Figure 4-10 show 1-2P, 2-4P, and 4-8P configurations, respectively.

The 1P-2P system configuration shown in Figure 4-8 has one Node board with one or two R10000 processors and does not require any type of Router board.

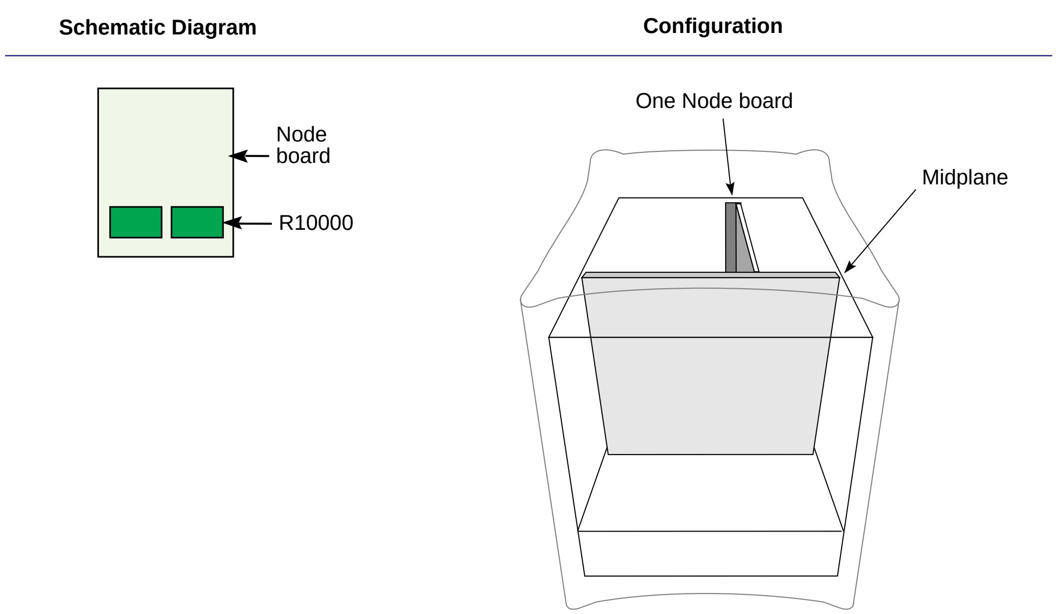

Figure 4-8 One or Two Processor Configuration (No Router Board Required)

The 2P-4P system deskside configuration uses two Node boards (each with one or two R10000 processors). Figure 4-9 shows a 4P configured with one Null Router board. In this configuration all 12 XIO board slots are enabled because the Node boards are positioned in slots 1 and 2. For more information, see “Board Configuration and Layout” in Chapter 3.

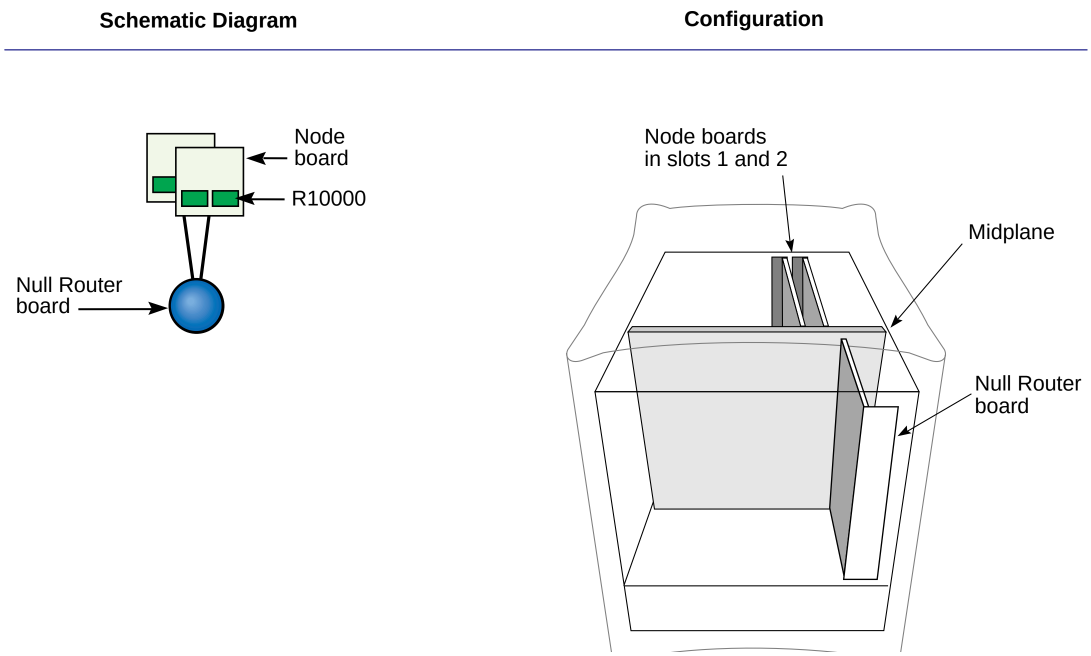

Figure 4-9 4P Configuration (Deskside Only)

The 4P-8P system4 to 8P system deskside configuration in Figure 4-10 shows a four-Node-board system with four to eight R10000 CPUs that uses a Star Router board along with a Rack Router board. The Star Router board has one external port that is used exclusively to connect with a port on a companion Rack Router board. Because of power grounding issues, it is not possible to use the remaining unused ports on the Rack Router to connect with other modules.

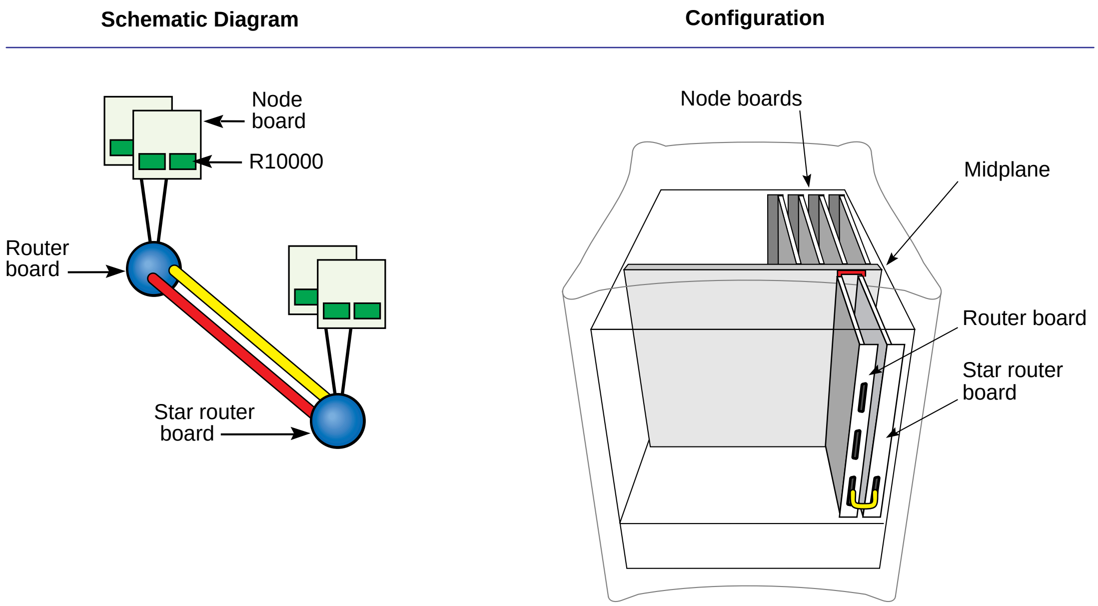

Figure 4-10 Origin2000 Deskside 8P with Star Router

4.4 Sample Origin2000 Rackmount Configurations

Figure 4-11 and Figure 4-12 show various single rackmount Origin2000 configurations. The single rackmount has either one or two module systems. Table 4-2 and Table 4-3 summarize the range of hardware components that can be supported by either a one-or two-module configuration.

The rackmount configurations also house Origin Vault drive boxes to provide additional storage capacity. One Origin Vault SCSI drive box comes standard with the rackmount configuration. A rackmount system can support up to four Origin Vault drive boxes, if a second module is not present.

The Origin2000 chassis is a standard 19-inch rack system that comes with a multmodule System Controller display and provides convenient cable management hardware for CrayLink Interconnect cables.

Table 4-2 Single-Module Rackmount Configurations

Hardware Component	Number Range
IP27 Node boards	1 to 4

Table 4-2 Single-Module Rackmount Configurations

Hardware Component	Number Range
R10000 CPUs	1 to 8
Main memory per Node board	64 MBs to 4 GBs (using 256-MB memory DIMMs)
Main memory per chassis	64 MBs to 16 GBs (using 4 Node boards)
Number of internal SCA drives	1 to 5
Number of XIO slots	6 or 12 for a single-module chassis
Number of Origin Vaults	1 to 4 (one Vault is standard with the rack)
I/O interfaces	- One (single-ended) SCSI connector - Two 9-pin D-sub RS-232 or RS-422 serial connectors - Two hardware interrupt lines - One 10/100-Base-T Ethernet connector

Table 4-3 Dual-Module Chassis Configurations

Hardware Component	Number Range
IP27 Node boards	2 to 8
R10000 CPUs	2 to 16
Main memory per Node board	64 MBs to 4 GBs (using 256-MB memory DIMMs)
Main memory per chassis	128 MBs to 32 GBs (using 8 Node boards)
Number of internal SCA drives	2 to 10
Number of XIO slots	18 or 24 for a dual-module chassis
Number of Origin Vaults	1 only
I/O interfaces	- One to two (single-ended) SCSI connectors - Two or four 9-pin D-sub RS-232 or RS-422 serial connectors - Two or four hardware interrupt lines - One to two 10/100-Base-T Ethernet connector

Figure 4-11 shows a single-module Origin2000 rackmount system. This configuration can house up to four Origin Vault drive boxes and one to eight CPUs. If desired, this rack can be cabled to other racks using the CrayLink Interconnect. This allows the other Origin2000 systems in the configuration to have high-speed access to those drives.

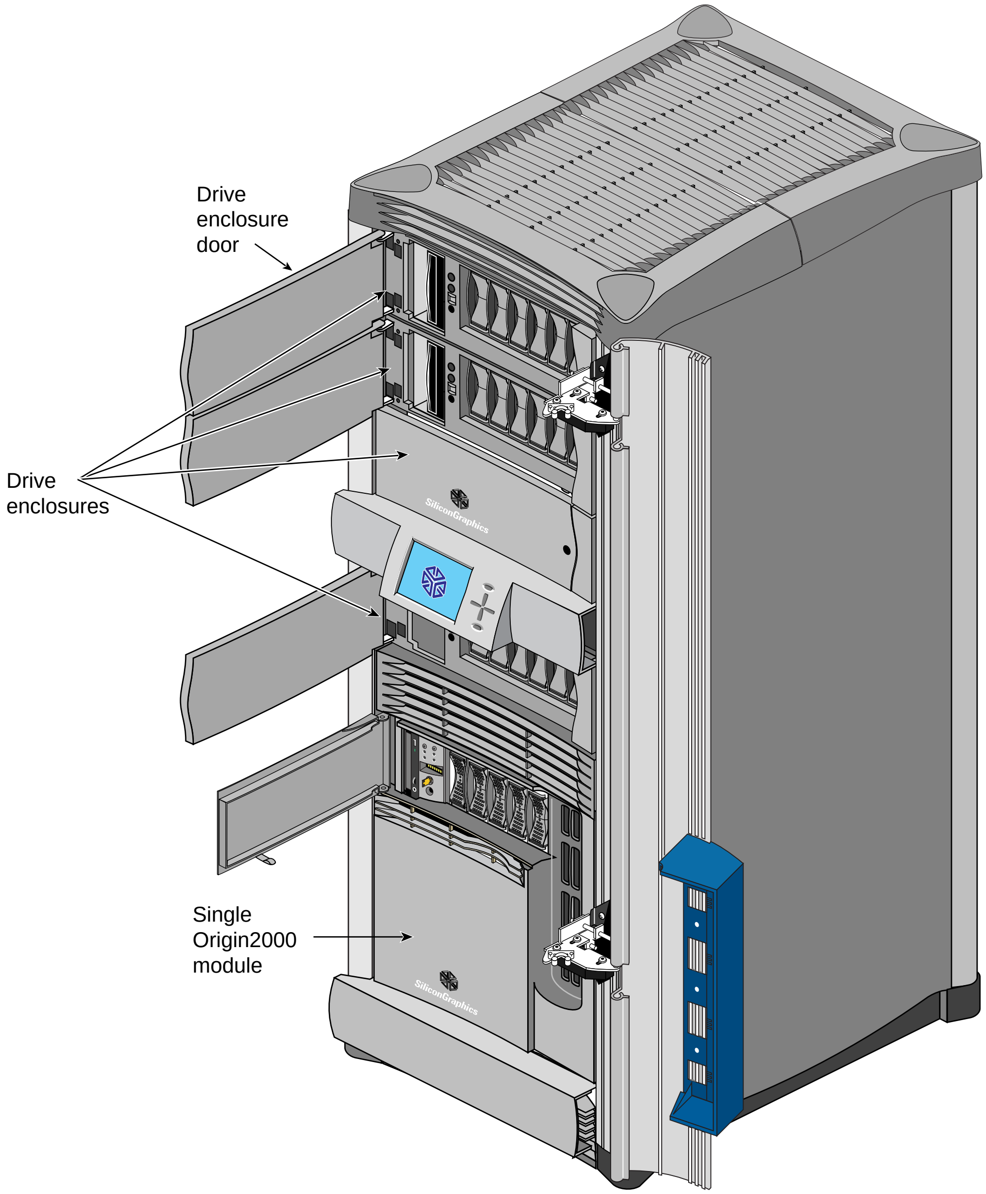

Figure 4-11 Single Origin2000 Module in a Rack With Four Origin Vault Drive Boxes

Figure 4-12 shows a 16P configuration using a combination of CrayLink Interconnect cabling and Xpress Links. The Xpress links can effectively double the interconnect bandwidth within a rack by providing additional 800-MBs / sec routing paths. The schematic in Figure 4-12 shows the criss-cross connection across Router boards.

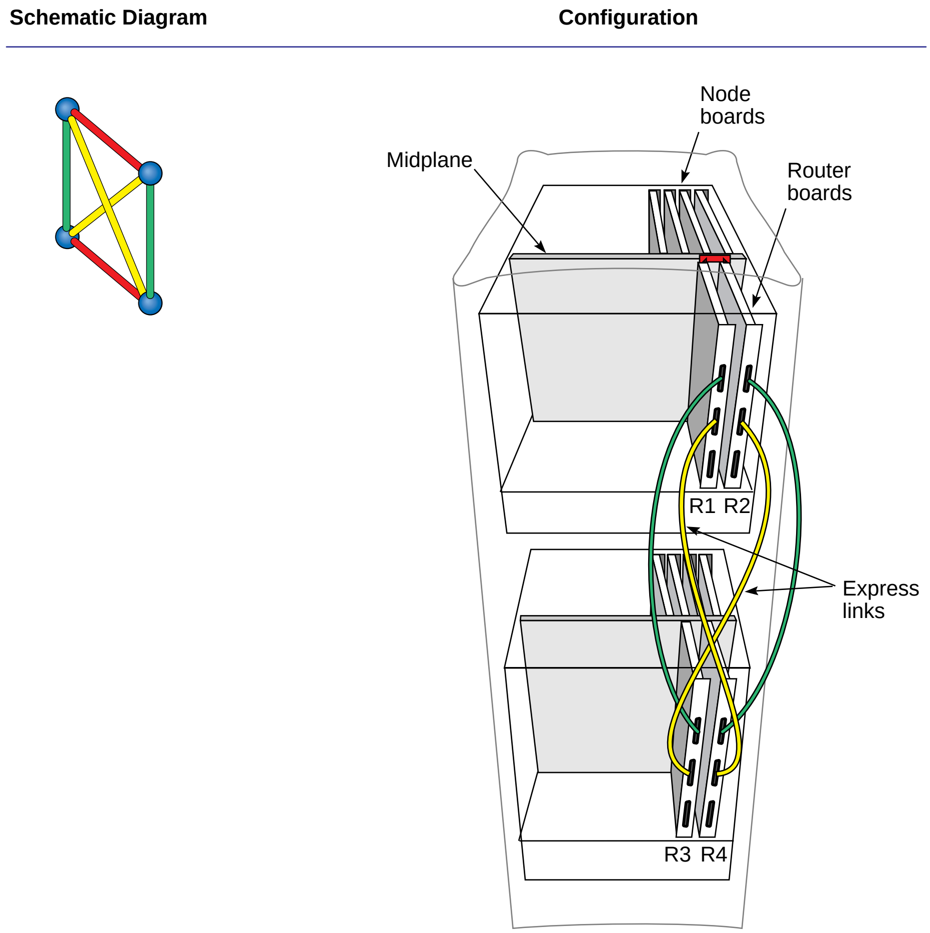

Figure 4-12 Origin200016P Rackmount With Xpress LInks

Figure 4-13 shows the 16P configuration without Xpress Links.

Schematic Diagram **Configuration**

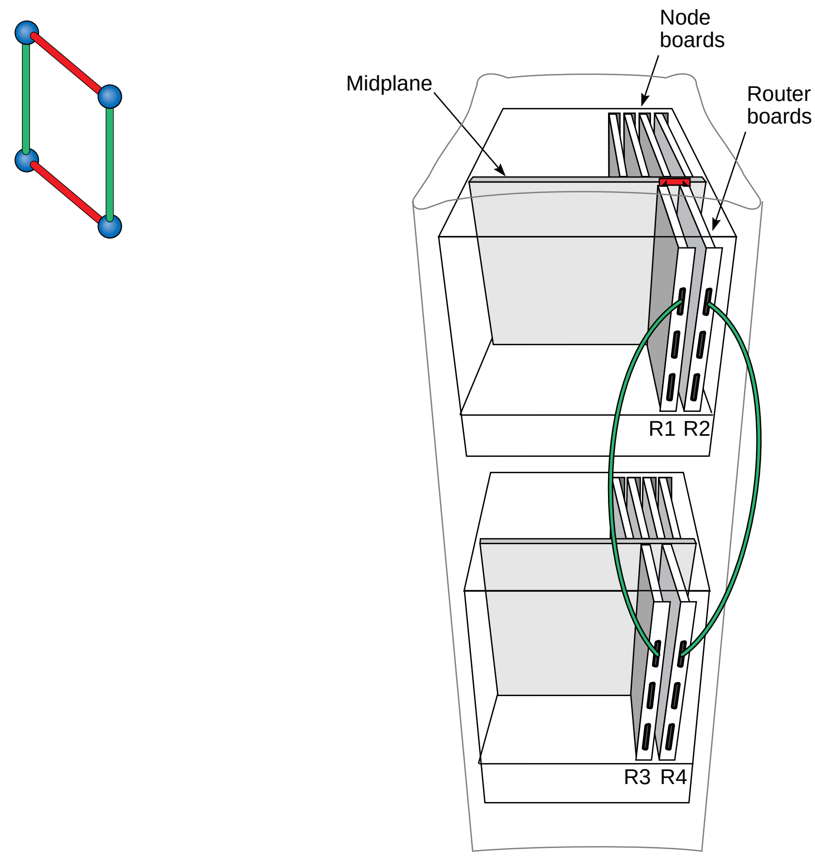

Figure 4-13 Origin200016P Rackmount

4.5 Origin2000 Multirack Configurations

Sample Origin2000 multirack configurations are shown in Figure 4-14 through Figure 4-17. These can be very large system configurations requiring a significant amount of floor space, along with special power and cooling considerations. The multirack configuration can have between 4 and 128 processors, up to 256 GBs of main memory, and up to 192 XIO slots through 16 separate system modules. If customers wish to expand their configuration, they should also consult the *Site Preparation* for *Origin Family and the Onyx2* manual (P / N 007-3452-xxx), a facilities engineer, and a reputable electrician.

Table 4-1 lists the various hardware components that can be supported through a multirack configuration.

Table 4-4 Multiple-Rack Configurations

Hardware Component	Number Range
Racks	2 to 8
IP27 Node boards	2 to 64
R10000 CPUs	2 to 128
Main memory per Node board	64 MBs to 4 GBs (using 256 MB-memory DIMMs)
Main memory per chassis	128 MBs to 32 GBs (using 8 Node boards)
Main memory (entire configuration)	128 MBs to 256 GBs)
Internal SCA drives	5 to 80
XIO slots	24 to 192

With router board connectors, the usual practice of connecting from top to bottom does not always apply. Table 4-5 summarizes their use.

Table 4-5 Router Board Connector Use

Position	Use
Top (A)	Vertical connection to the top connector on the Router immediately above or below
Middle (B)	16P or greater: Horizontal connection to adjacent cabinet in the pair; used only to connect the bottom pair of modules in a configuration with an odd number of modules
Bottom (C)	9-16P and 25-32P: Xpress Links 17P-24P: Diagonal link (pseudo Xpress Link) 33P-64P: Horizontal link across multiple cabinet 65P-128P: Link to Cray router

For example, for a 24P configuration, only the bottom modules would use middle connectors; the top module in rack 1 uses its bottom connector for the diagonal connection; its middle connector would be unused.

Figure 4-14 and Figure 4-15 show two examples of 32P (or processor) configurations. These systems require two racks along with the necessary CrayLink Interconnect cabling. Each rack has two modules, and each module has four Node boards and two Router boards. This configuration simply doubles the 16P configuration, by adding a second (16P) rack. The 32P configuration is represented schematically by a cube or a three-dimensional *hypercube* (see also “CrayLink Interconnect” in Chapter 1). The four ball points, or vertices, of the schematic square shown in Figure 4-7 and Figure 4-12 are duplicated and connected to form a cube. The eight points of the cube represent the eight Rack Router boards required for this configuration.

Note: A configuration can have any number of processors between 16P and 32P. The 16P and 32P configurations are used because they represent the maximum number of processors for a one- or two-rack system.

Schematic Diagram	Configuration

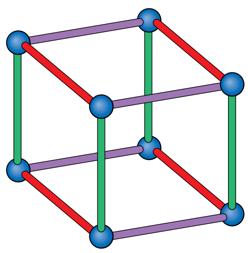

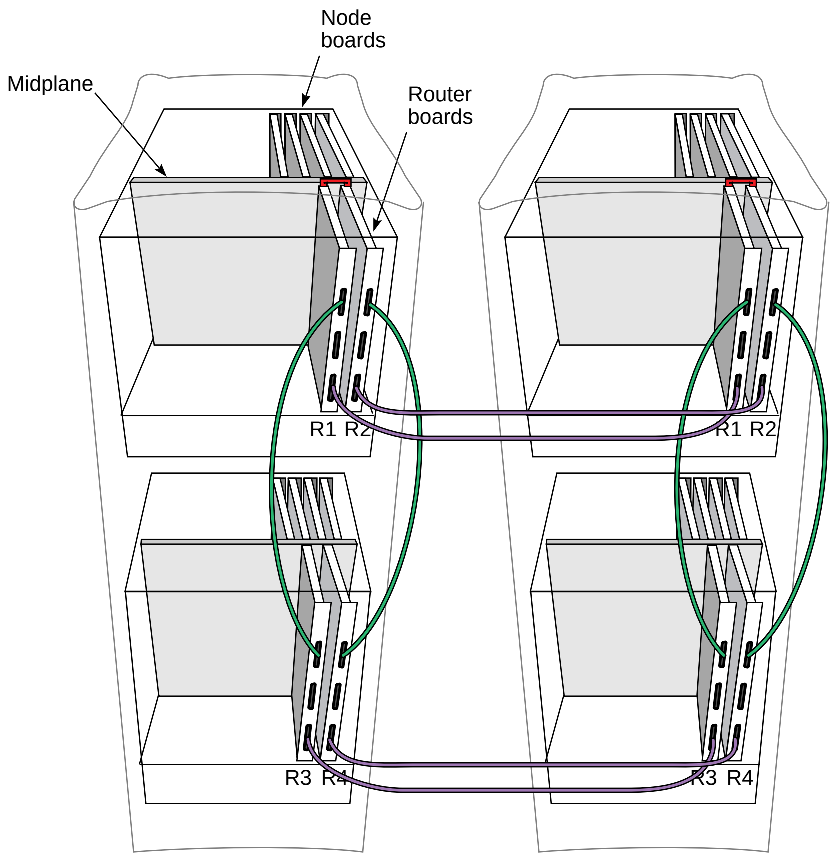

Figure 4-14 32P Configuration

The configuration shown in Figure 4-15 uses Xpress links to increase the transmission throughput between modules and racks. The Xpress links provide additional data routing which reduces potential transmission latency. In Figure 4-15, all the available router ports are used. To go beyond a 32-processor configuration, the Xpress links must be removed and more CrayLink Interconnect cables must be installed.

Schematic Diagram **Configuration**

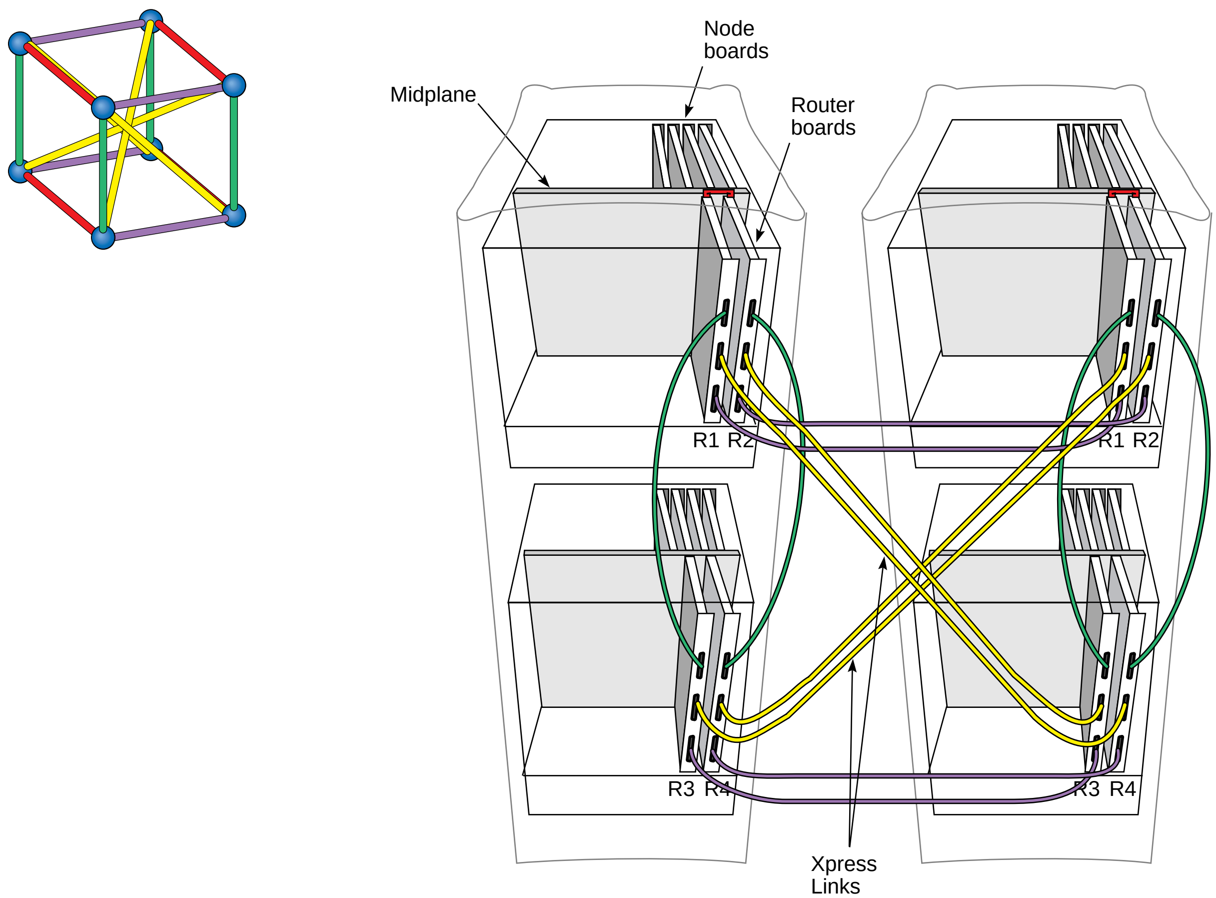

Figure 4-15 32P with Xpress Links

Figure 4-16 represents a 64P configuration which is simply a double of the previous 32P system. This configuration consists of 4 racks, 8 system modules, 16 Rack Router boards, and 32 Node boards. Schematically, a 64P configuration consists of two cubes that are joined at common points or vertices.

As the Figure 4-16 shows, all of the available router ports are now in use. The 64P configuration is the largest one that can be supported using the Rack Router board. To go beyond 64P requires the use of a Cray Router chassis. With a Cray Router chassis, configurations of up to 128P are possible (see Figure 4-17).

Schematic Diagram

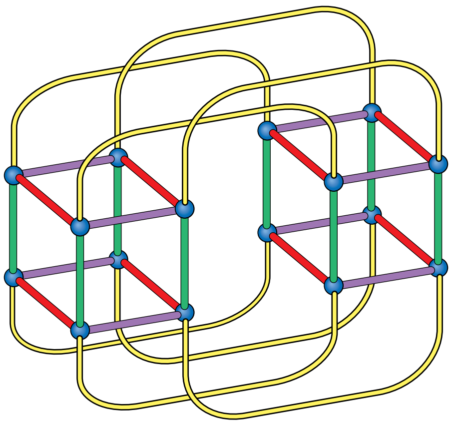

Configuration

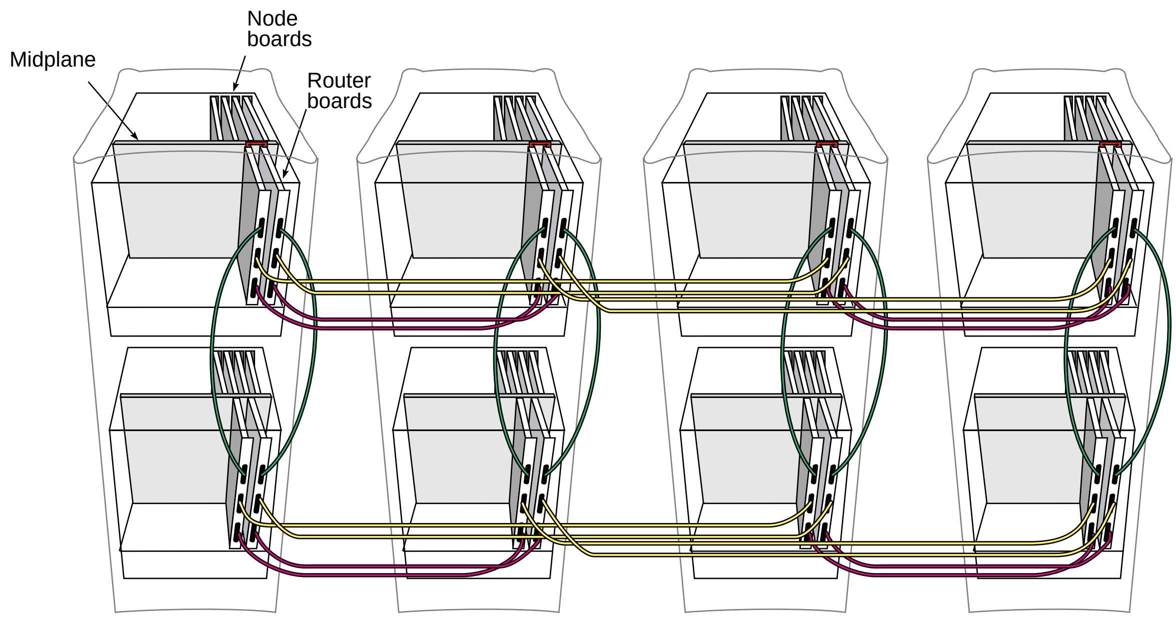

Figure 4-16 64P Configuration

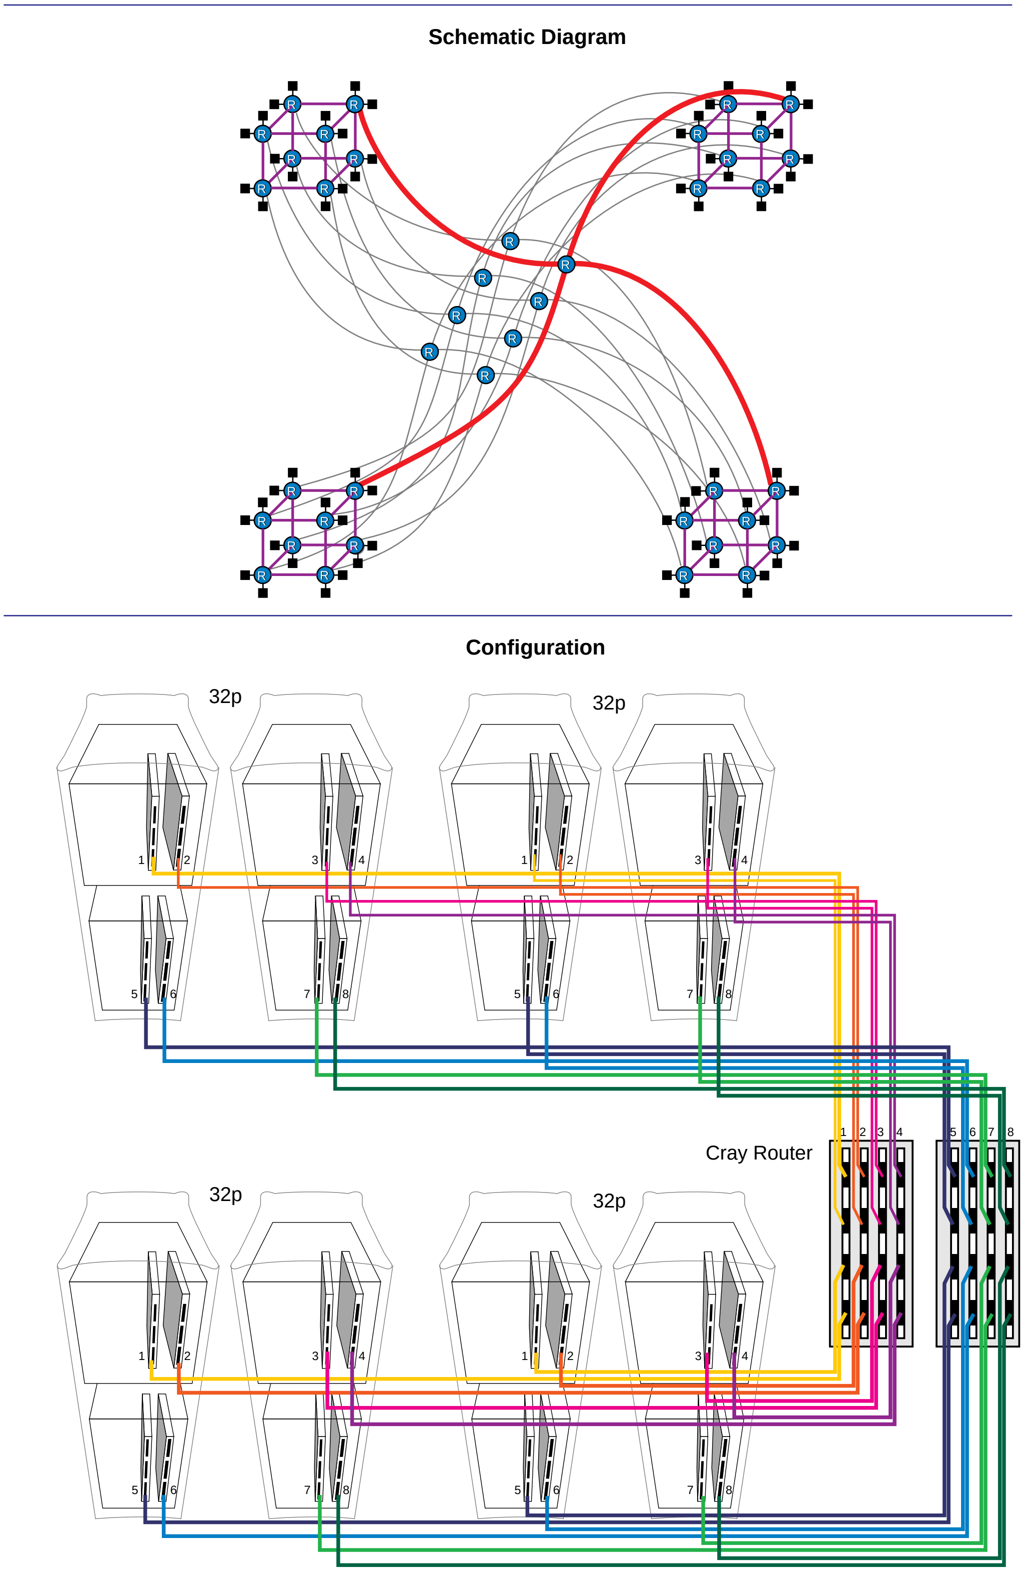

Figure 4-17 128P Configuration

4.6 Onyx2 Graphics Configurations

This section describes the Onyx2 Reality and InfiniteReality graphics configurations. In many instances, the Onyx2 system and Origin2000 server systems use much of the same hardware. As a result, many components are swappable between server and graphics systems including:

- Node boards
- Router boards
- drives
- System Controllers
- power supplies
- XIO boards and other parts

The Onyx2 also comes in deskside, rackmount, and multirack configurations; however, the dual rack (with 24P or six graphics pipes) is the largest configuration available at the initial release of this product.

4.6.1 Onyx2 Deskside

There are two graphics deskside versions:

- Reality
- InfiniteReality

4.6.1.1 Onyx2 Reality Deskside System

The entry-level Reality graphics system ships *only* as a deskside configuration. The following lists the major hardware components:

- one to four R10000 (1 to 2 Node boards)
- 64 MBs to 8 GBs of main memory
- GE14-4 with four Geometry Engines (GEs)
- five XIO slots
- 20-inch, 1280 x 1024 monitor
- one to two RM8 boards, each with a 40-MB framebuffer and either 16 or 64 MBs of texture memory

4.6.2 Onyx2 InfiniteReality Deskside System

The Reality and InfiniteReality deskside systems differ in the three following aspects.

- The InfiniteReality deskside uses an RM7/TM7 board; the Reality deskside uses an RM8.
- The InfiniteReality deskside system can fully support a super wide, 24-inch monitor; the Reality deskside can best support the smaller 20-inch monitor.
- two to four R10000 (one to two Node boards)

  Note: The InfiniteReality deskside configuration requires at least two R10000 processors.
- 64-MB to 8-GB RAM
- five I/O slots
- superwide (24-inch) 1,920 x 1,200 monitor
- GE14-4 (with four Geometry Engines)
- RM7 board with 80-MB framebuffer
- TM7 daughterboard with either 16 or 64 MBs of texture memory

Note: You cannot mix 16- and 64-MB TM7s. Both TM7 boards must have the same texture memory size.

4.7 Onxy2 Rackmount Configurations

All Onyx2 rack systems use the InfiniteReality graphics and are housed in a standard 19-inch chassis (see Figure 4-18). Like the Origin2000 server racks, the Onyx2 systems come with cable management hardware (both in the front and in the back of the chassis), multimodule System Controller, and CrayLink Interconnection capability (except between graphics modules).

Note: Only server modules can use the CrayLink Interconnection.

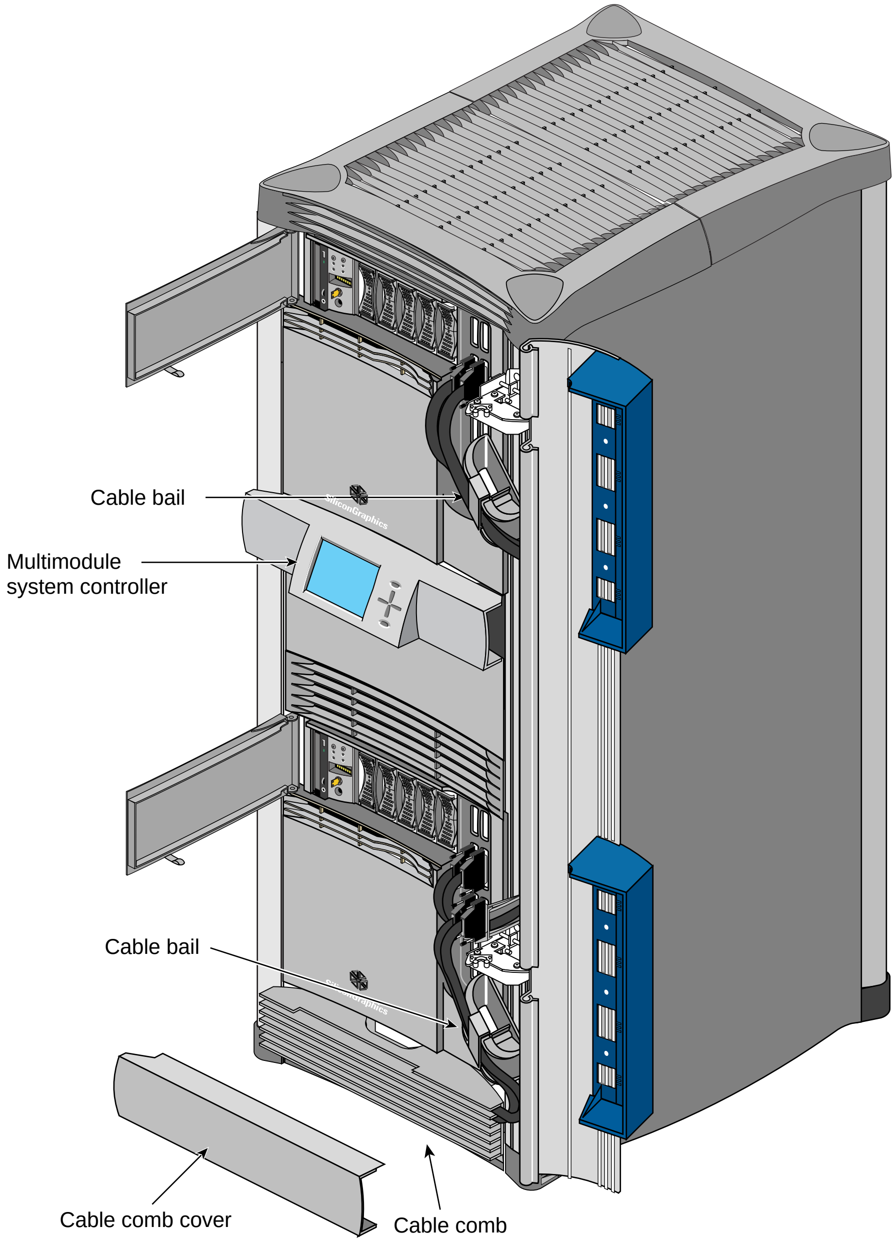

Figure 4-18 Onyx2 Single-Rack Configuration

4.7.1 Onyx2 Single Rack System

The following lists the basic hardware components of the Onxy2 rackmount (see also Figure 4-18 and Figure 4-19):

- up to two graphics pipes (one with two RMs and the other with four RMs)

  Note: Each pipe requires at least one Node board for proper processing support. For example, if an Onyx2 rack has two pipes, then at least two Node boards must be installed.

- two to eight processors (one to four Node boards)
- 128 MBs to 16 GBs RAM
- CD-ROM drive
- up to five internal drives, including the system hard drive
- module System Controllers (MSCs) for both the graphics and server modules
- up to 12 (half-height) XIO slots or up to six (full-height) XIO slots
- Superwide (24-inch) 1,920 x 1,200 monitor
- GE14-4 (with four Geometry Engines)
- RM7 board with 80-MB framebuffer
- TM7 daughterboard with either 16 or 64 MBs of texture memory

Note: You cannot mix 16- and 64-MB TM7s. Both TM7 boards must have the same texture memory size.

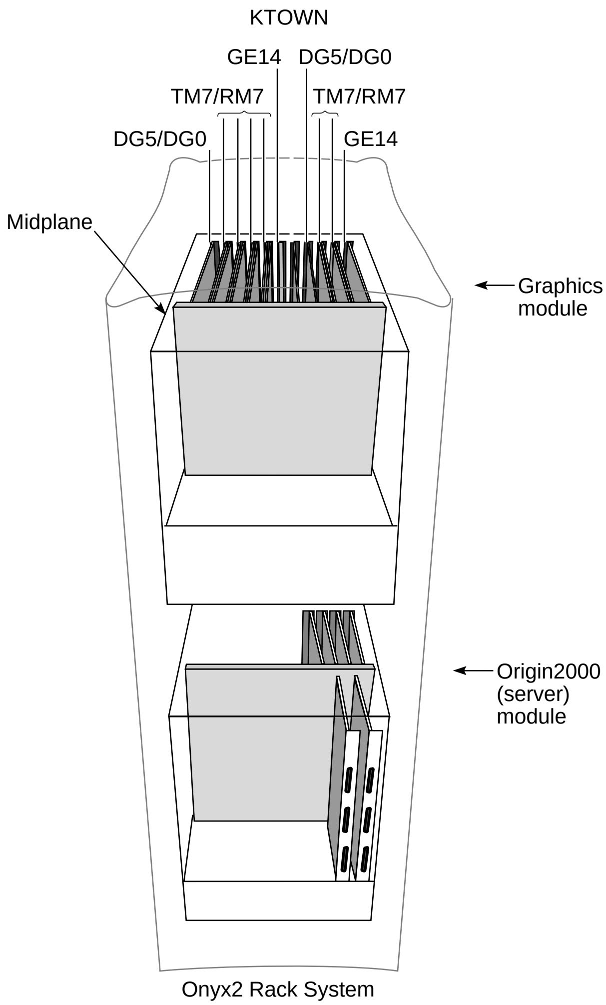

Figure 4-19 Onyx2 Rackmount (8P)

4.7.2 Onyx2 Multirack Systems

The Onyx2 multirack configurations can have between 2 and 24 processors, from one to four pipes, up to 48 GBs of main memory, and up to 36 XIO slots through four separate system modules. If customers wish to expand their configuration to larger graphics configuration, they should also consult the *Site Preparation* for *Origin Family and the Onyx2* manual (P / N 007-3452-001), a facilities engineer, and a reputable electrician.

Each pipe requires at least one Node board for proper processing support. For example, if an Onyx2 rack has two pipes, then at least two Node boards must be installed.

Note: For information on how the pipes are numbered, see “Pipe Numbering in the Onyx2 Rack System” in Chapter 10.

Figure 4-20 shows a 16P dual-rack Onyx2 configuration with four graphics heads. Each rack is a fully-independent system and uses CrayLink Interconnection to route data between the server modules.

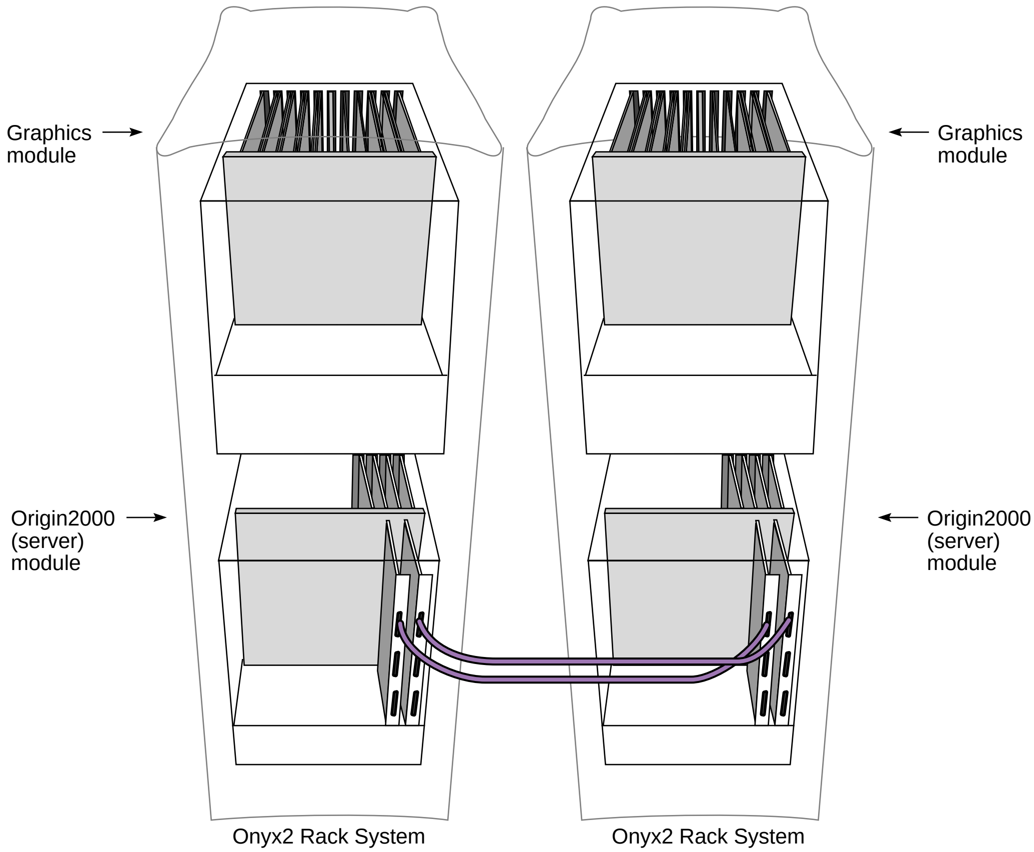

Figure 4-20 Two Onyx2 Racks With 16 Processors and Four Graphics Pipes

Figure 4-21 is a 24P configuration with two graphics pipes. Only one of the racks has a graphics module. The other rack is an Origin2000 server system (see Figure 4-22).

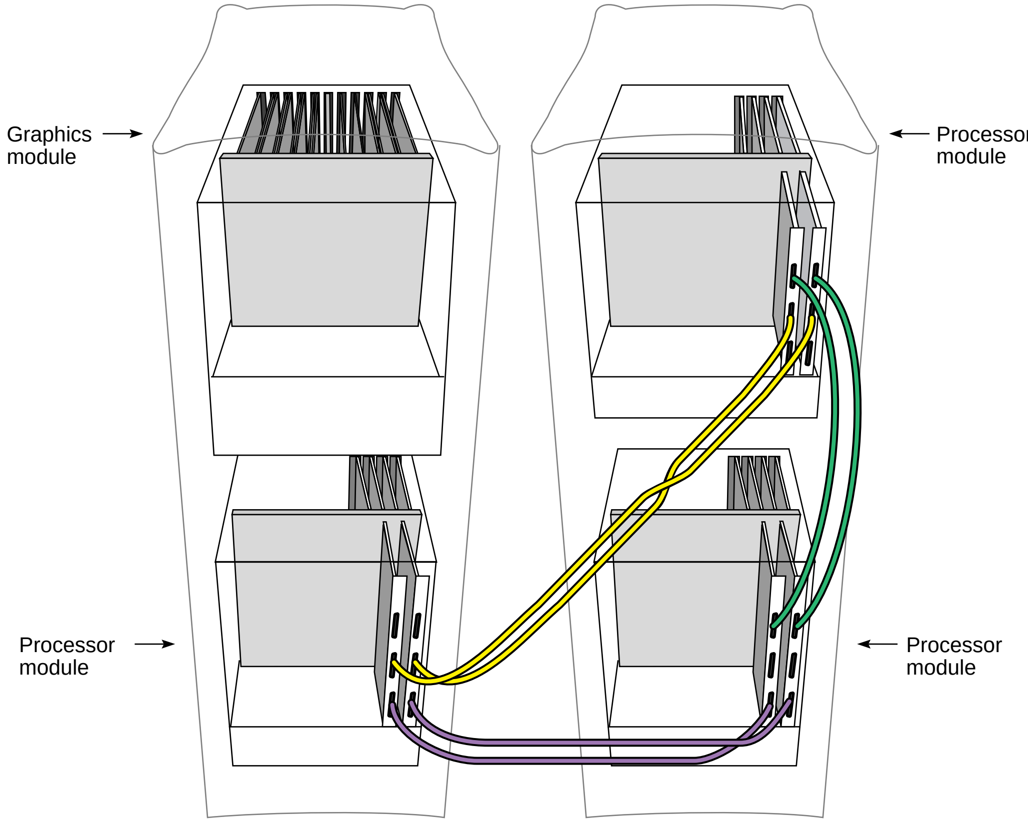

Figure 4-21 24P Onyx2 Configuration with Two Graphics Pipes

Figure 4-22 Actual 24P Configuration

4.8 Configuration Guidelines for the Origin2000 and Onyx2

This section provides guidelines to follow when adding Node and XIO boards to a Origin2000 deskside and rackmount system and to an Onyx2 rackmount system. This information helps optimize system performance and maximize compute and I/O resources. *This section does not apply to Onyx2 deskside systems.*

Configuration guidelines in this section are given for the following:

- single module systems
- systems with more than one module (that is, rackmounted systems containing more than one system and / or graphics module)

These rules and guidelines assume that a multimodule system has two 8-processor (8P) router cards in each module — an assumption that is not a strict configuration requirement.

Figure 4-23 shows the locations of Node and XIO boards as viewed from the rear of the Origin2000 and Onyx2 chassis.

- The odd Node board slots (1,3) are connected to Crossbow 1, which in turn is connected to the lowest six XIO boards (slots 1 through 6).
- The even Node board slots (2,4) are connected to Crossbow 2, which in turn is connected to the highest six XIO boards (slots 7 through 12)

Figure 4-23 also shows how the XIO boards are oriented in their slots.

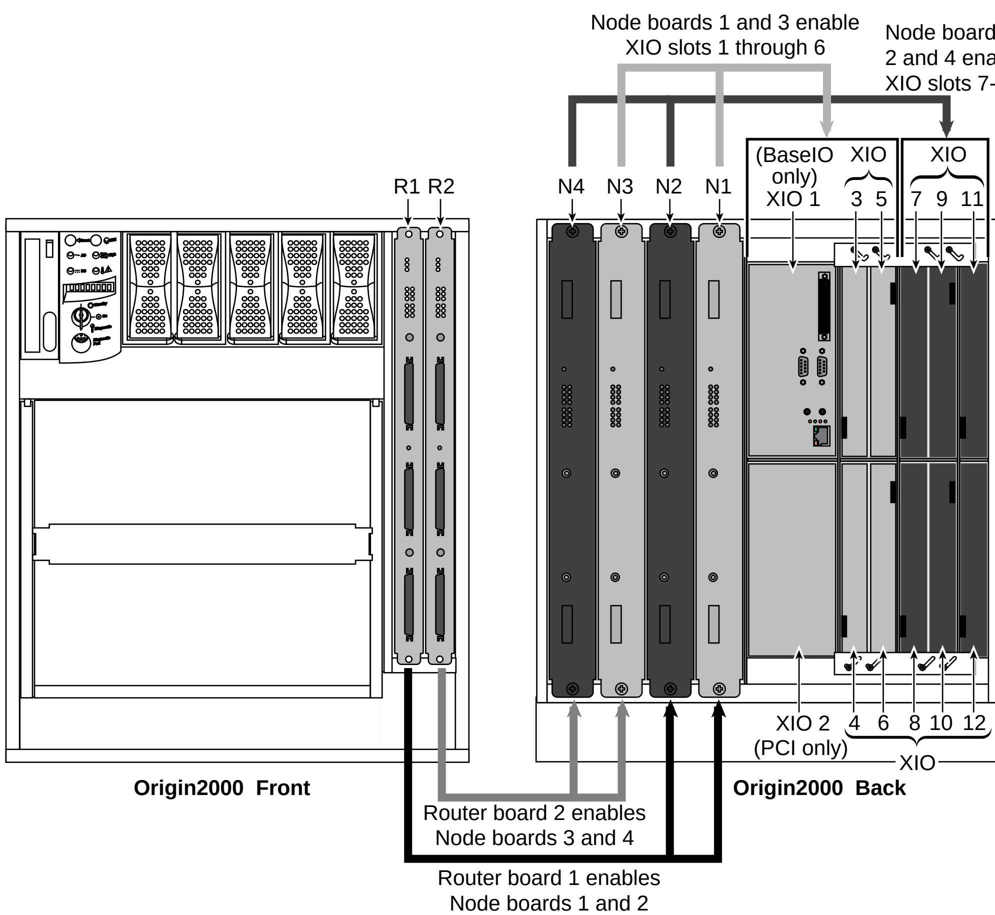

Figure 4-23 Board Positions in Origin2000 Chassis

There are six rules, for both single-module and multiple-module systems, that *must be followed* when configuring the Origin2000:

- each Origin2000 module must have a Node board
- there must be at least one BaseIO board in the system (whether single module or multiple module system)
- the BaseIO board must go in slot IO1
- the PCI expansion must go in slot IO2 (beneath the BaseIO board)
- neither a BaseIO board nor an MSCSI board can ever go into slot IO2
- to use all twelve of the XIO slots in a module, there must be two Node boards installed: one Node board in an odd slot (N1 or N3) and one Node board in an even slot (N2 or N4).

4.8.1 Single Module Guidelines

This section provides guidelines for adding Node and XIO boards to a single module system.

4.8.1.1 Adding Node Boards To a Single Module System

When filling a single module, install the Node boards from center outward, starting at slot N1, then N2, N3, and finally N4.

4.8.1.2 Adding XIO Boards To a Single Module System

Install the BaseIO board in slot IO1 and install the PCI expansion module in slot IO2.

If there is only one Node board in the module, it is placed in slot N1 and just the leftmost IO slots (slots 1 through 6) are usable; in this case fill the slots sequentially, 3 through 6. As described above, slot IO1 holds the BaseIO board, and slot IO2 contains the PCI expansion module.

If there is more than one Node board, fill the odd XIO slots first (1,3,5,7,9,11) and then fill the even XIO slots (2,4,6,8,10,12), so that XIO cards are spread fairly evenly across the system.

4.8.2 Multiple Module Guidelines

This section provides guidelines for adding Node and XIO boards to a system with more than one module. The configuration guidelines are based on the following constraints:

- maximizing the number of Node boards in the system (in which case, see the section titled, Section 4.8.2.1, “Maximizing the Number of Node Boards in More Than One Module”)

- maximizing the number of XIO boards in the system (in which case, see the section titled, Section 4.8.2.2, “Maximizing the Number of XIO Boards in More Than One Module”)

4.8.2.1 Maximizing the Number of Node Boards in More Than One Module

If the configuration of a multiple-module system is determined by the number of Node boards you want to install, the following guidelines should be followed:

- Add Node boards within a module until the module is full, starting with slot N1, then N2, N3, and finally N4. Only then should you add Node boards to the next module. In other words, do not have more than one module unfilled with Node boards. For instance, in a five Node board system, four of the Node boards should go in one module, with the remaining Node board going into the second module (4N+1N).

  Note: The step above differs importantly from XIO installation guideline presented in the next section.

- Install all available BaseIO boards in IO1 slots. Remember, there can only be one BaseIO board in a module, and it can only go in the IO1 slot.
- Fill any remaining unpopulated IO1 slots with MSCSI boards, if available.
- Install any PCI expansion boxes in IO2 slots. Remember, there can only be one PCI expansion box in a module, and it can only go in the IO2 slot.
- Fill any remaining unpopulated IO2 slots with XIO boards. Remember, MSCSI boards can never go in the IO2 slot.
- Take the remaining XIO boards and divide them evenly among the modules. Fill the odd XIO slots first (3, 5, 7, 9, 11) and then fill the even XIO slots (4, 6, 8, 10, 12), so that XIO cards are spread fairly evenly across the system.

4.8.2.2 Maximizing the Number of XIO Boards in More Than One Module

If the configuration of a multiple-module system is constrained by the number of XIO boards you want to install, use the following guidelines:

- Distribute the Node boards among slots N1 and N2 as evenly as possible throughout all modules. That is, fill slots N1 and N2 in all modules before adding any Node boards to slot(s) N3, and then N4. This allows maximum distribution of XIO cards in each module: to use all twelve XIO slots, there must be a Node board installed in one of the odd slots (N1, N3) and a Node board in one of the even slots (N2, N4). For instance, in a five Node board system, three of the Node boards should go in one module, with the remaining two Node board going into the second module (3N+2N).

  Note: The step above differs importantly from Node board installation guidelines presented in the previous section.

- Install all available BaseIO boards in IO1 slots. Remember, there can only be one BaseIO board in a module, and it can only go in the IO1 slot.
- Fill any remaining unpopulated IO1 slots with MSCSI boards, if available.
- Install any PCI expansion boxes in IO2 slots. Remember, there can only be one PCI expansion box in a module, and it can only go in the IO2 slot.

- Fill any remaining unpopulated IO2 slots with XIO boards. Remember, MSCSI boards can never go in the IO2 slot.
- Take the remaining XIO boards and divide them evenly among the modules. Fill the odd XIO slots first (3,5, 7, 9, 11) and then fill the even XIO slots (4, 6, 8, 10, 12), so that XIO cards are spread fairly evenly across the system.

Chapter 5

Preinstallation Checklist

This chapter provides the steps that must be performed before you begin the actual installation. It is extremely important to read and perform these steps.

Note: It is highly recommended that the SSEs be familiar with configuring a network.

5.1 Before You Begin Installation

The installation will go more smoothly if a set of preinstallation information is obtained in advance of hardware bring-up, as described in Table 5-1. Much of this information can be obtained days or weeks in advance of customer hardware rollout.

The site should have already been selected and prepared as described in the *Site Preparation for Origin Family and Onyx2* manual (P/N 007-3452-xxx). In addition, be sure to perform and check off the following technical, administrative, and bureaucratic tasks before installation to help ensure a successful setup.

Table 5-1 Preinstallation Checklist

Activity	Date Completed	Notes
Read supporting manuals, web pages, release notes and reference (man) pages.	__/__/__	
Gather the supporting documentation.	__/__/__	

Lookup the World Wide Web (WWW) sources for the latest firewall-protected information:
NSD Tech Pubs Next Generation Doc Archive -
IP27 PROM Technical Reference Manual
ASD'S Web Document Depot -
`http://b7.asd.sgi.com/doc/index.html`
Origin2000 Software
`http://babylon.engr.sgi.com/lego/`
	__/__/__	

Table 5-1 (continued) Preinstallation Checklist

Activity	Date Completed	Notes
Read the related release notes and reference pages. These documents contain important information not covered in the manuals such as diagnostics testing and command definition.		
Establish clear lines of communication with the customer.	__/__/__	
Ensure that you have strategic points of contact with the field and with the customer (see Section 5.2). The customer contacts should represent the customer's needs and have the authority to make decisions.	__/__/__	
Verify that all the required hardware and software have been received and inventoried (see Section 5.3).	__/__/__	
Obtain the inventory list and the system serial numbers from the customer.	__/__/__	
Verify that floor space, local area networking, and environmental preparations have been properly completed (see Section 5.4).	__/__/__	
Determine the local area network (typically Ethernet) IP addresses and hostnames and be sure they've been properly allocated and registered at the site (see Section 5.5).	__/__/__	
Label the networking and CrayLink Interconnect cables (see Section 5.6).	__/__/__	

5.2 Establish Customer Contact

Because Origin2000 and Onyx2 systems are commonly placed in large organizations, it is important to configure the system with respect to the administrative structure in place at the site. Be sure that your contacts have the ability and authority to make the required decisions.

5.3 Verify and Inventory the Components

The Origin2000 and Onyx2 system may consist of many layered hardware and software components, which if absent prevent the complete and proper operation of the system. Carefully ensure that all inventory has arrived at the site and has cleared the customer's local inventory control procedures. Collect and record the system serial numbers (see Chapter 7, "Diagnostic Tools," for information on obtaining the serial numbers through the *hinv* command).

5.4 Complete Floor Space, Network, and Environmental Planning

Be sure the proper physical, electrical, and environmental preparations have been completed in advance of the installation.

Be sure proper floor space has been allocated to permit connectivity of the Origin2000 and Onyx2 system configuration. Be sure networking plans have been constructed to place the Origin2000 or Onyx2 system on a local Ethernet. Be sure proper power and cooling have been arranged.

5.5 Determine Local Network IP Addresses and Hostnames

You must have the following information to configure the LAN:

- a network number
- IP addresses and hostnames on this network for each Origin2000 and Onyx2 system

These addresses and hostnames must be allocated and registered at the customer's site. Be sure to have the network address, host IP addresses and hostnames before continuing.

At this point, the preparatory steps should have been completed. Verify that the preconfiguration checklist in Table 5-1 is complete before you begin the actual installation.

5.6 Label Networking and Craylink Interconnect Cables

Label both ends of the networking and Craylink Interconnect cables with numbers that can be matched up later.

Chapter 6

Installation

This chapter describes how to unpack, cable, and configure the Origin2000 and Onyx2 deskside and rackmount systems.

6.1 Safety

Read these safety statements carefully before you install or remove any part of the system.

Warning: Installation of this product requires specific training and technical knowledge. These instructions have been provided for use by Silicon Graphics system support engineers (SSEs) and Silicon Graphics–trained personnel only. This equipment uses electrical power internally that is hazardous if the equipment is improperly assembled or disassembled.

Warning: Wait two minutes after powering off a rackmount system before working on any part of the power supply or midplane / backplane. Because of the large amount of capacitance in the system, a significant fraction of the operating voltage remains on the midplane / backplane for up to two minutes after the system is powered off.

Caution: This equipment is extremely sensitive to damage from electrostatic discharge (ESD) caused by the build-up of electrical potential on clothing and other materials.

Follow these ESD preventive measures:

- Attach a ground strap to your wrist when connecting / disconnecting boards or peripherals.
- Ensure that you and all the electrical equipment you handle during this installation are at ground potential to avoid damage from ESD.
- Keep the boards in the antistatic bags until they are needed.
- Place the boards only on an antistatic mat. Do not place boards on top of an antistatic bag unless the outside of the bag also has antistatic protection.
- Remove a board from its antistatic bag only when you are properly grounded to the chassis ground with a ground strap.

- If you are servicing a system or installing a hardware upgrade, do not disconnect the power cord from the wall socket *and* the chassis. You will lose the system ground and could damage the equipment as a result.
- Do not use an ohmmeter on the boards.

6.2 Unpacking a Deskside System

Read this section for information on how to unpack the deskside system.

Warning: Be careful when unpacking and moving the Origin2000 or Onyx2 deskside chassis. Ensure that the chassis remains on a level surface and that the chassis weight remains evenly distributed across the four casters. If the casters must be lifted over an obstacle, such as a door threshold, use proper lifting techniques and employ a minimum of two people.

Caution: Do not subject the cabinet to any unnecessary shocks or vibration while unpacking and installing the system.

See Figure 6-1 and follow these procedures to unpack the box.

1. Lift up the top of the cardboard box to expose the top of the system, remove the protective foam, then remove the bottom cardboard skirt.
2. Locate the two polystyrene ramps supplied with the packing material.
3. Set the ramps in place in front of the chassis. Secure the ramps by installing the braces into the notches of the ramp and platform.
4. Fit the strap onto the sled and slowly pull it down the ramp.
5. Loosen the wing nuts that secure the chassis to the platform sled and release the anchoring plates. Roll the chassis off the sled.

   Warning: The deskside chassis weigh as much as 170 lbs (77 Kg). Take care when rolling it down the ramp.

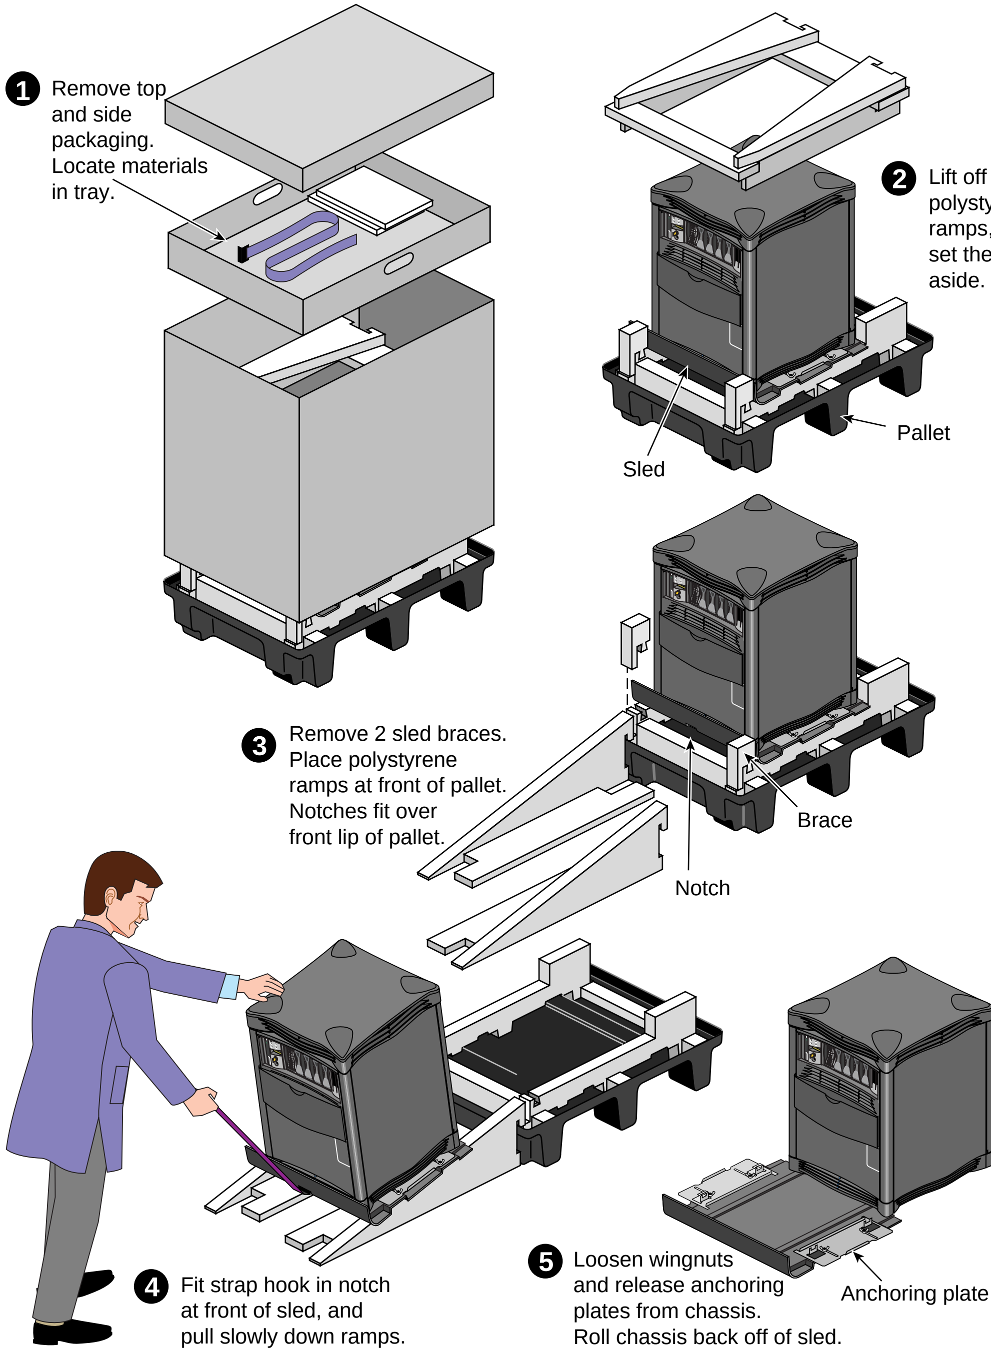

Figure 6-1 Unpacking an Origin2000 and Onyx2 Deskside System

6.3 Unpacking a Rackmount System

Read this section for information on how to unpack the rackmount system.

Warning: In its maximum configuration, the rackmounted system weighs approximately 900 pounds (408 kg). Use caution when unpacking and moving the system chassis. Ensure that the chassis remains on a level surface and that the chassis weight remains evenly distributed across the four casters. If the casters must be lifted over an obstacle, such as a door threshold, use proper lifting techniques and employ a minimum of two people. Exercise the same caution when moving monitors or terminals.

The system arrives at the customer site in a wooden crate. Packed along with the system are the documentation carton and the accessories carton. The documentation carton contains the system manuals as well as warranty and licensing information. The accessories carton contains the system cables and any additional connectors or tools required for a specific configuration. The monitors and / or terminal are shipped in separate cartons.

The following steps describe how to unpack a rackmount system, illustrated in Figure 6-2 and Figure 6-3. Note that a minimum of two people are required to safely remove a rackmount system from the packing crate:

1. Release the latches securing the top side panels.
2. Remove the top of the crate and the side panels.
3. Remove the two wing nuts from the front and rear support brackets.
4. Remove the roped padding.
5. Remove the cable door and drive access door(s). See “General Information” in Chapter 10 for instructions.
6. Position the ramp on the slot immediately in front of the system platform.

   Note: When positioning the ramp, flip the ramp up, place the supports underneath, and make sure that the bottom edge contacts evenly with the floor.

7. Using two people, carefully roll the system out of the crate and down the front door.

   Warning: Use extreme caution when rolling the system chassis down the door. The weight and dimensions of the system make it difficult to control. Personal injury and system damage could result if the chassis becomes unbalanced or builds too much momentum coming down the ramp.

8. Replace the doors and install the top plenum panels. See “General Information” in Chapter 10 for instructions.
9. Once in place, extend the four foot bolts at the bottom of the chassis to prevent the system from rolling.

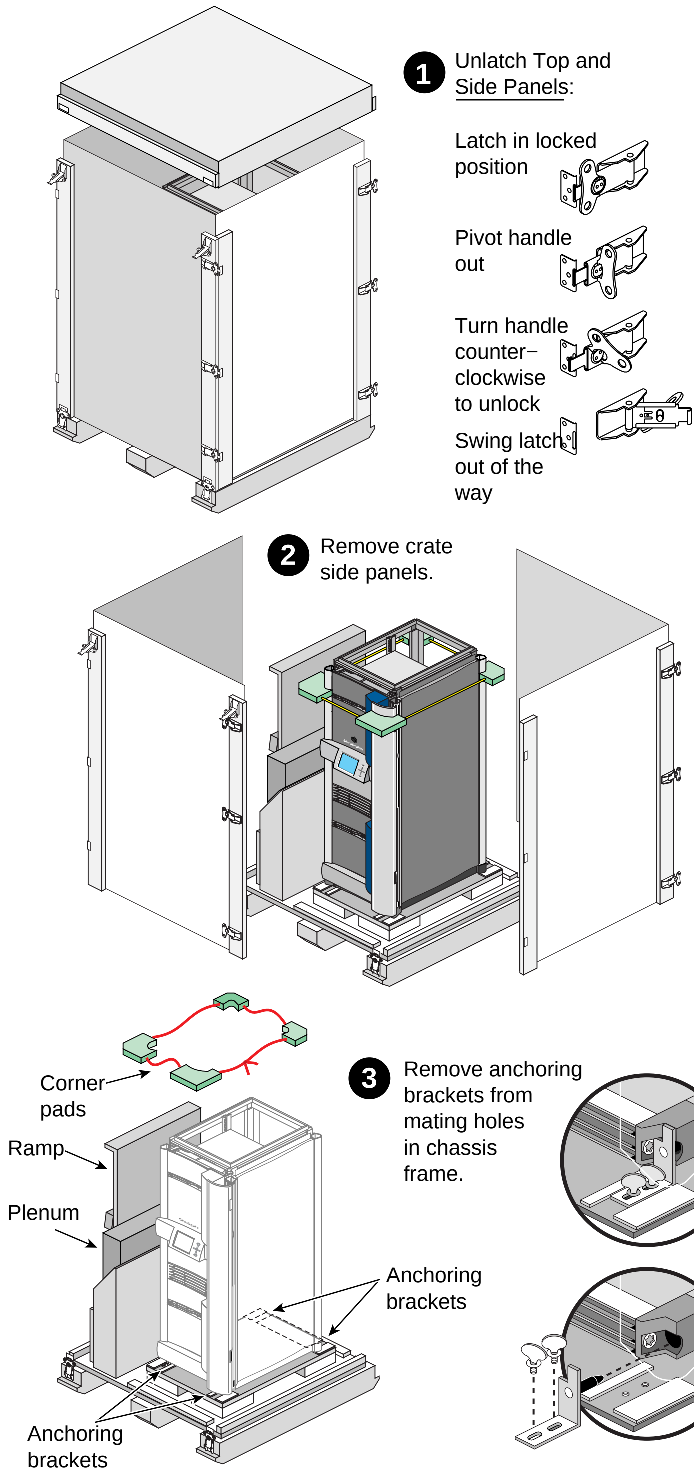

Figure 6-2 Unpacking the Rackmount System Shipping Container (Part1)

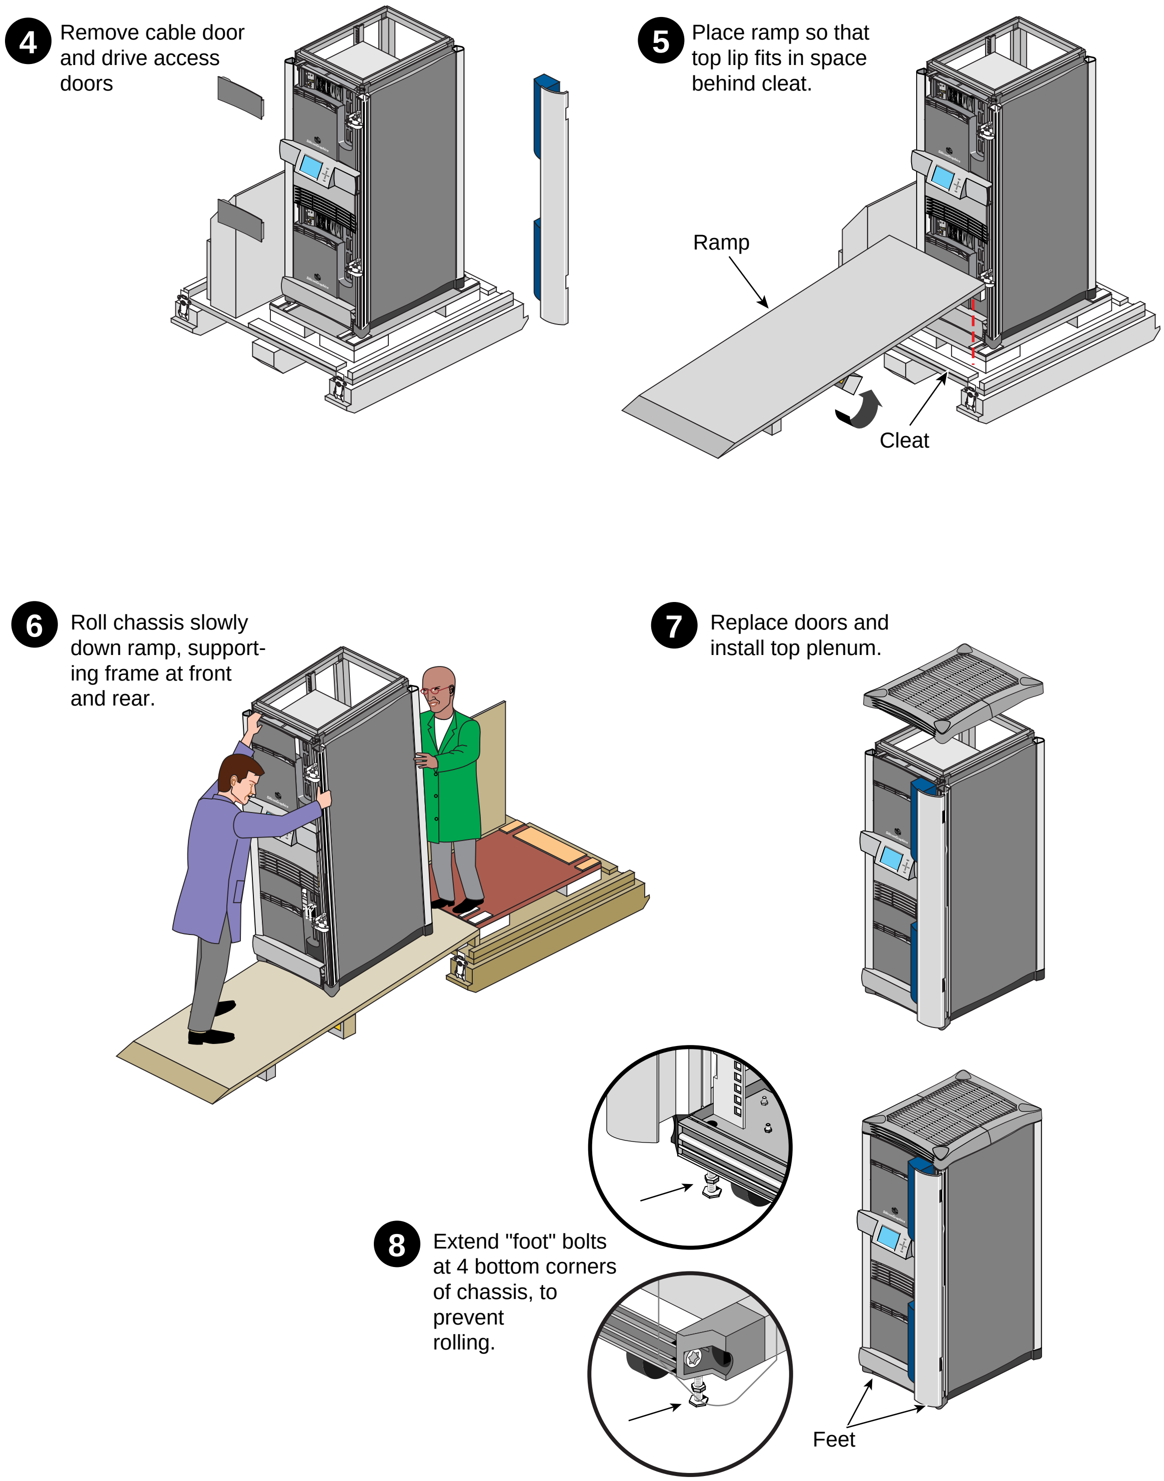

Figure 6-3 Unpacking the Rackmount System Shipping Container (Part II)

6.4 Key Points in This Chapter

This section highlights key points of information to help you during system installation.

- You cannot use CrayLink Interconnect cabling on standalone deskside systems.
- The Origin2000 and Onyx2 systems use PC-style serial cables. You cannot use earlier Silicon Graphics serial cables.
- In a multirack configuration, always position the rack with the multimodule controller and display in the leftmost spot and place the additional racks to the right of this rack.
- On Origin2000 systems and Onyx2 rackmount systems, the XIO boards have a different installation orientation from slot to slot. XIO boards IO1/ IO2, IO5/IO6, and IO9/IO10 are installed with the component side facing left. XIO boards IO3/IO4, IO7/IO8, IO11/IO12 are installed with the component side facing right.

6.5 Graphics System Installation

Before beginning the installation procedure, verify that the installation site meets the space, power, and environmental guidelines found in the *Site Preparation for Origin Family and Onyx2* manual (P/N 007-3452-xxx).

6.5.1 Graphics System Installation Summary

See Table 6-1 for a quick installation checklist for the Onyx2 deskside or rackmount system.

Table 6-1 Quick Onyx2 Deskside and Rackmount System Installation Checklist

Tasks	Location in Manual
1. Read the safety precautions.	See Section 6.1, “Safety.”
2. Unpack the system.	See Section 6.2, “Unpacking a Deskside System.”
3. Connect the monitor.	See Section 6.5.2, “Monitor Cabling.”
4. Connect the keyboard and mouse.	See Section 6.5.3, “Keyboard and Mouse Cabling.”
5. If you have a rack system, set up the System Controller connections.	See Section 6.7, “Setting Up the System Controller Communications (Rack Systems Only).”
6. Power on the system and verify normal operation.	See Section 6.8, “Powering On and Verifying Operation.”
7. Load the front-loading drives.	See Section 6.9, “Loading the Single Connector Assembly (SCA) Drives Into the System.”
8. Install optional XIO boards.	See Section 6.10, “Installing XIO Boards.”
9. Connect the printer.	See Section 6.11.1, “Printer Cabling.”

Table 6-1 (continued) Quick Onyx2 Deskside and Rackmount System Installation Checklist

Tasks	Location in Manual
10. Connect the modem (optional).	See Section 6.11.2, “Modem Cabling.”
11. Load the software.	See Section 6.12, “Software Installation.”
12. Configure the monitor.	See Section 6.13, “Configuring the Monitor.”
13. Install the CrayLink Interconnect cabling (if required)	See Section 6.14, “CrayLink Interconnect Cabling.”
14. Setup a multirack configuration (if required).	See Section 6.16, “Setting Up Multiple-Rack Configurations.”
15. Connect to Ethernet.	See Section 6.17, “Connecting the System to Ethernet LAN.”

6.5.2 Monitor Cabling

The Onyx2 Reality deskside system ships with a 20-inch, 1,280 x 1,024 monitor. The Onyx2 InfiniteReality deskside and rackmount configurations ship with a super-wide 24-inch, 1,920 x 1200 monitor. Both monitors use a 13W3-to13W3 cable (see Figure 6-4).

Two additional cables are also available for connection with other monitors (see Figure 6-4).

- 13W3-to-5 BNCs (red, green, blue, vertical sync, and horizontal / composite sync)
- 13W3-to-13W3 with two BNCs (vertical sync and horizontal / composite sync)

To connect a monitor, install the appropriate monitor cable to the DG5 board in the back of the Onyx2 system (see Figure 6-4). The default monitor port is Monitor 0. Connect the other end of the cable to the monitor. Figure 6-5 shows how to install monitors on a multipipe configuration.

For additional information regarding monitor operation, see the monitor’s user’s guide.

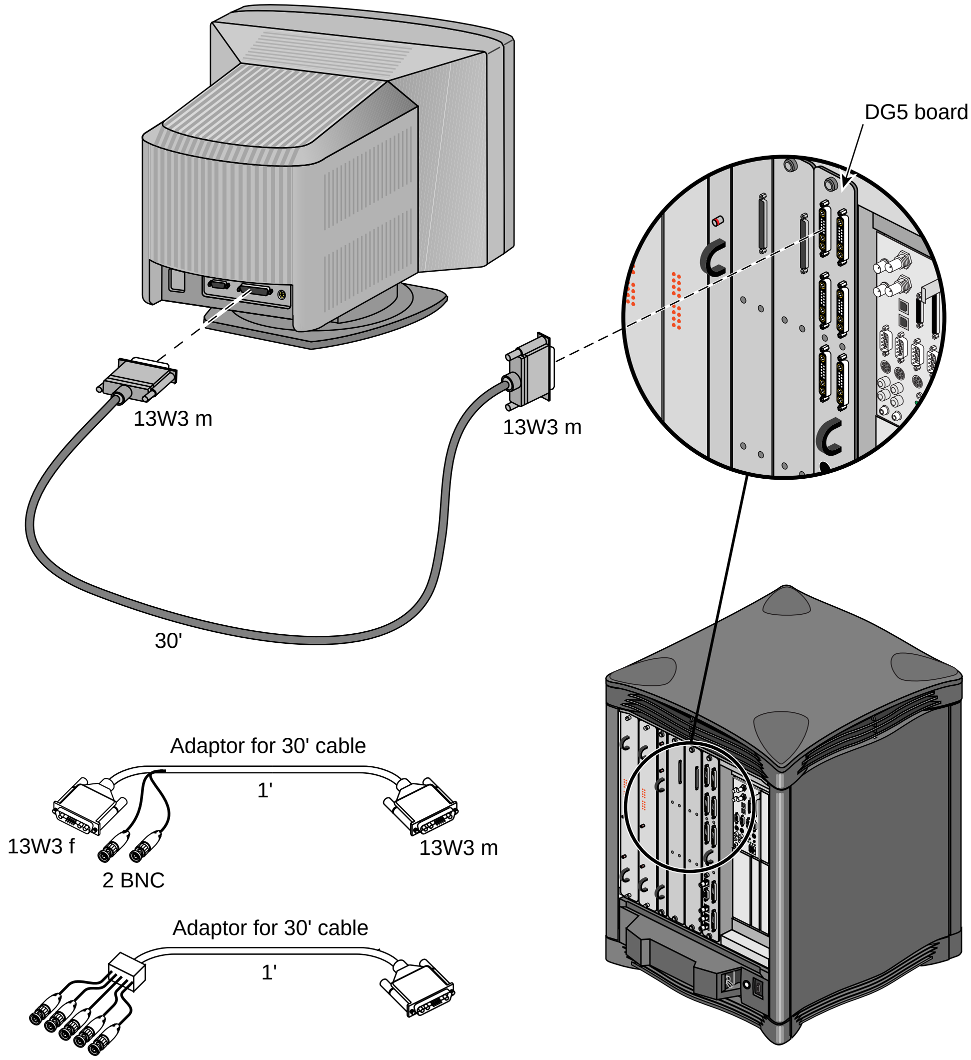

Figure 6-4 Connecting a Monitor

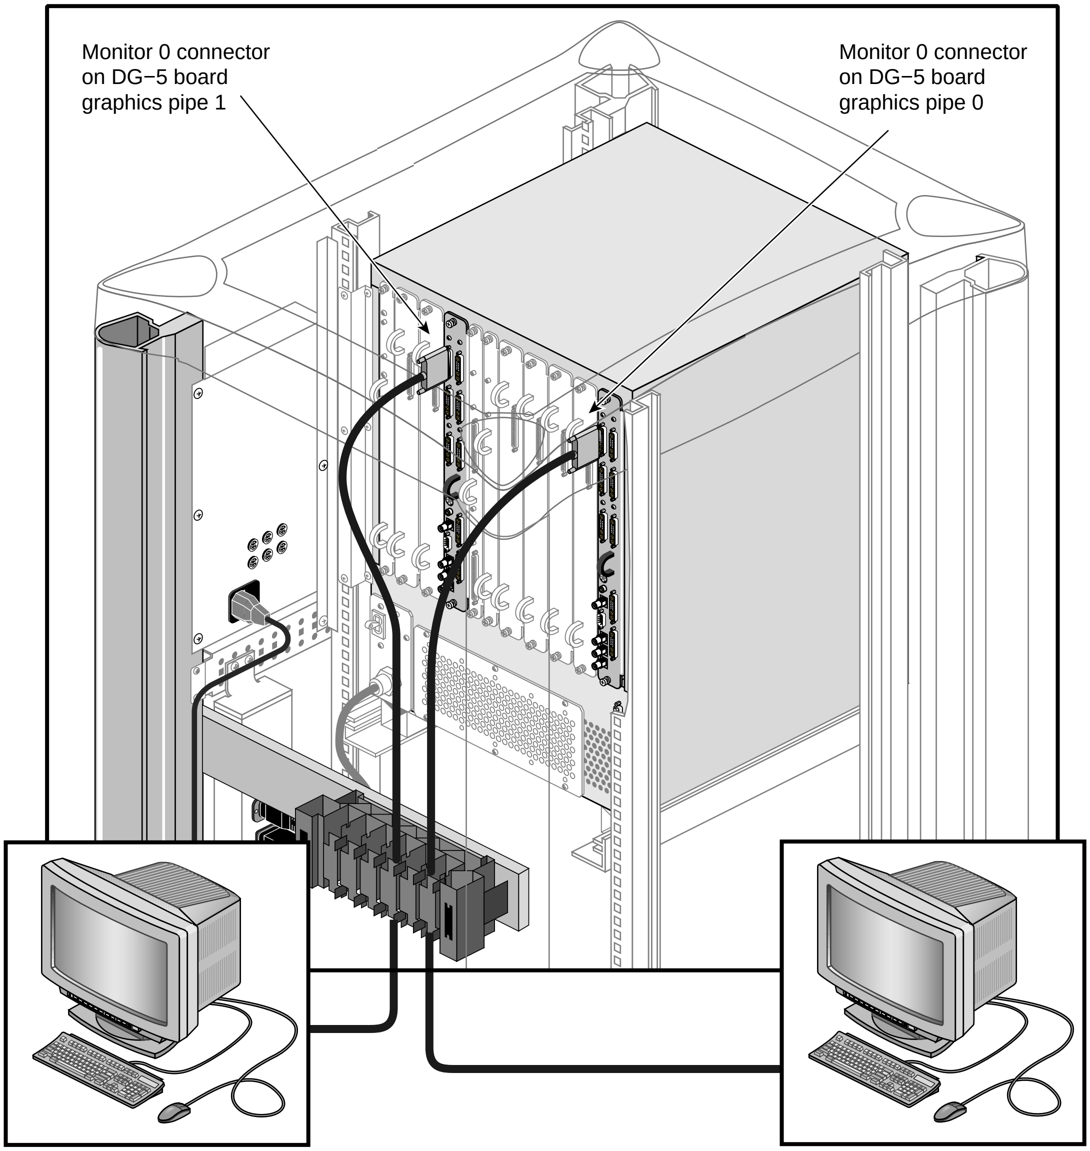

Figure 6-5 Connecting Monitors to a Multipipe Graphics Module

6.5.3 Keyboard and Mouse Cabling

The Onyx2 system ships with a standard 101-key international keyboard. The keyboard has two identical plug receptacles located in the upper right and left corners. These receptacles will accept the six-pin connectors from either the keyboard cable or the mouse. Attach the keyboard as shown in Figure 6-6.

Note: After you install the keyboard and mouse, go to Section 6.8, "Powering On and Verifying Operation."

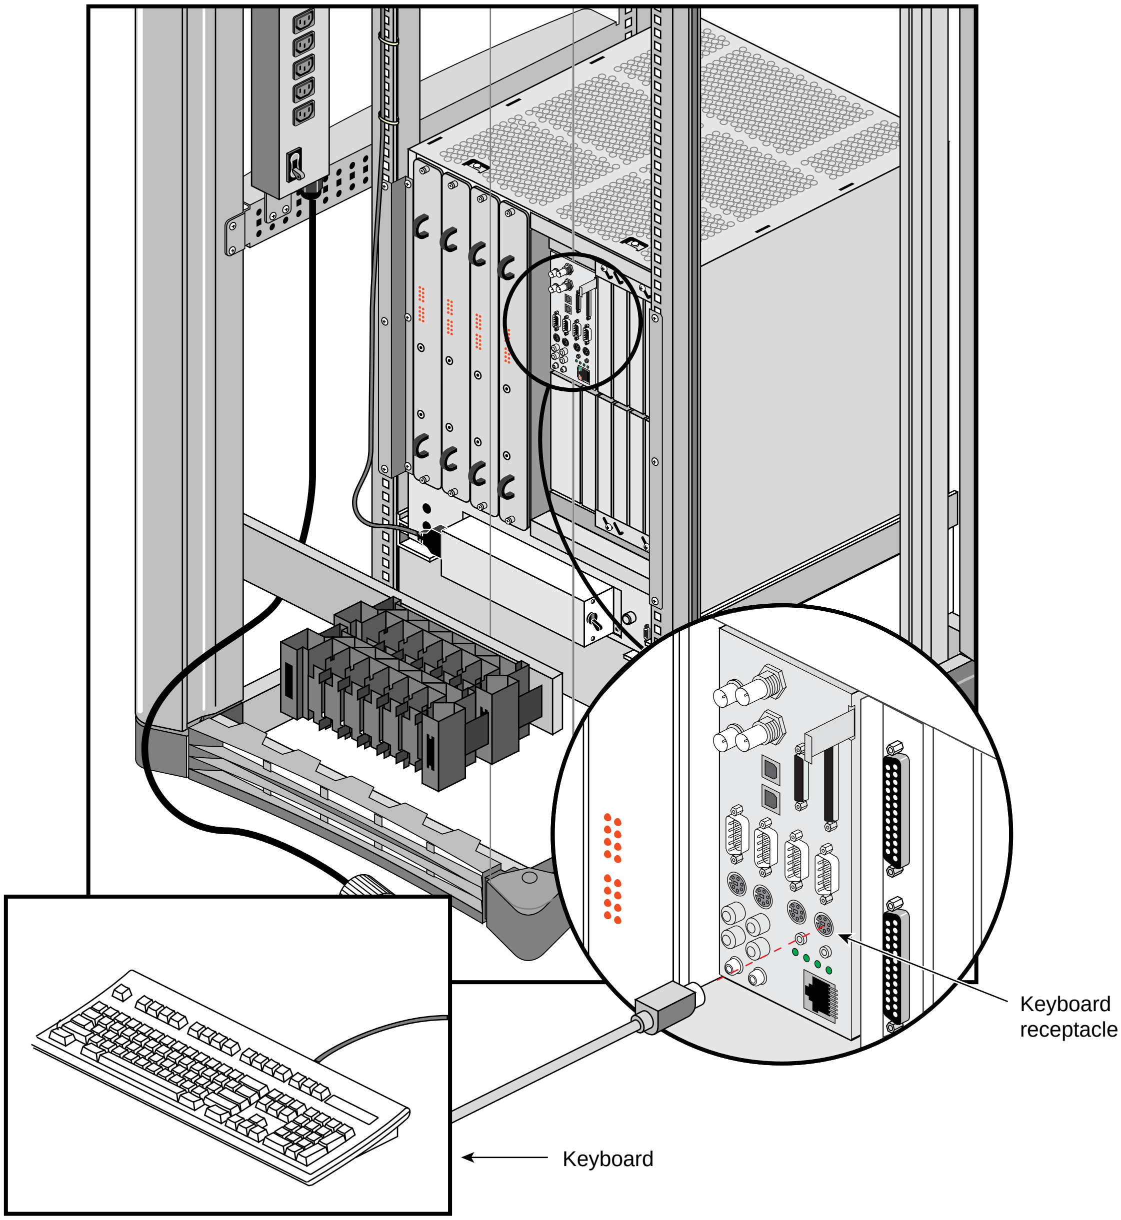

Figure 6-6 Connecting the Keyboard

6.6 Server Installation

Before beginning the installation procedure, verify that the installation site meets the space, power, and environmental guidelines found in the *Site Preparation for Origin Family and Onyx2* manual (P / N 007-3452-xxx).

6.6.1 Server System Installation Summary

See Table 6-2 for a quick installation checklist for the Origin2000 systems.

Table 6-2 Quick Origin2000 System Installation Checklist

Tasks	Location in Manual
1. Read the safety precautions.	See Section 6.1, “Safety.”
2. Unpack the system.	See Section 6.2, “Unpacking a Deskside System.”
3. Connect the terminal.	See Section 6.6.2, “Terminal and Keyboard Cabling.”
4. Connect the keyboard.	See Section 6.6.2, “Terminal and Keyboard Cabling.”
5. If you have a rack system, set up the System Controller connections.	See Section 6.7, “Setting Up the System Controller Communications (Rack Systems Only).”
6. Power on the system and verify normal operation.	See Section 6.8, “Powering On and Verifying Operation.”
7. Load the front-loading drives.	See Section 6.9, “Loading the Single Connector Assembly (SCA) Drives Into the System.”
8. Install optional XIO boards.	See Section 6.10, “Installing XIO Boards.”
9. Connect the printer.	See Section 6.11.1, “Printer Cabling.”
10. Connect the modem (optional).	See Section 6.11.2, “Modem Cabling.”
11. Load the software.	See Section 6.12, “Software Installation.”
12. Install the CrayLink Interconnect cabling (if required).	See Section 6.14, “CrayLink Interconnect Cabling.”
13. Setup a multirack configuration (if required).	See Section 6.16, “Setting Up Multiple-Rack Configurations.”
14. Connect to Ethernet.	See Section 6.17, “Connecting the System to Ethernet LAN.”

6.6.2 Terminal and Keyboard Cabling

The Origin2000 system server configuration requires a customer-supplied ASCII terminal and keyboard. The terminal connects to the system using standard PC-style serial 9-pin to 25-pin cables (see Note). Make sure that the terminal is set to the following parameters:

- 9,600 baud, full duplex
- XON / XOFF handshaking
- 8 bits
- 1 stop bit

Note: You cannot use earlier Silicon Graphics serial cables.

6.6.2.1 Connecting a Terminal to a Deskside

Attach a terminal to a deskside system as shown in Figure 6-7 using the tty_1 Console port.

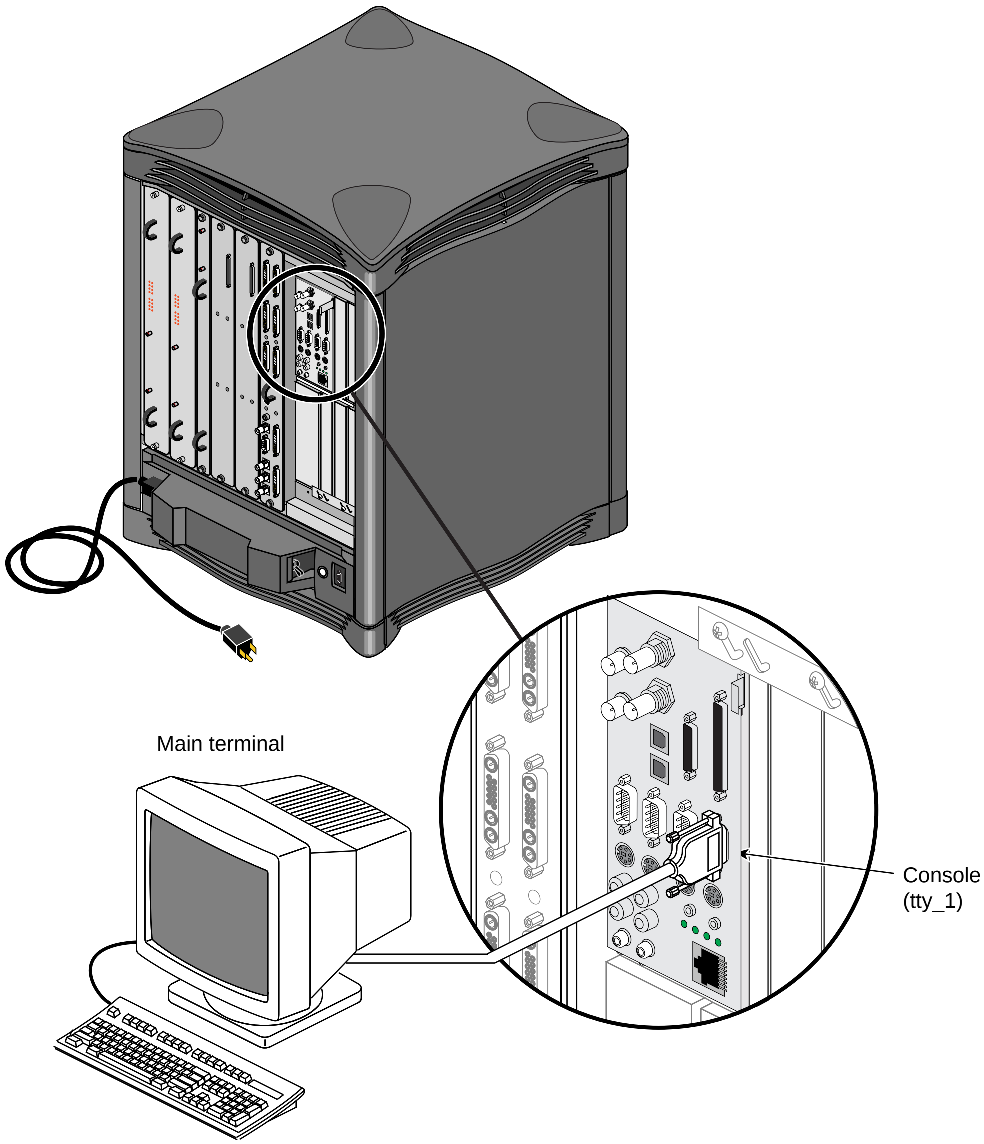

Figure 6-7 Connecting a Terminal to a Deskside System

6.6.2.2 Connecting a Terminal to a Rackmount System

The rackmount system requires a DIN-8-to-DB9 cable, a DIN-8-to DB9 converter, and a PC-style serial (9-pin-to-25-in). cable. These cables or equivalent types, should be provided with the system shipment. It is highly recommended that you connect a terminal to a rackmount system as follows:

1. Connect the tty_1 (Console) port on the BaseIO board to the COM 4 port on the MMSC using a DIN-8-to-DB9 cable (see Figure 6-8). This enables remote access to the BaseIO console port through the MMSC.

   Note: You can connect a terminal directly to the tty_1 (Console) port on the BaseIO board using a 9-to-25-pin PC-style cable. However, you will not be able to access this port remotely as a result.

2. Connect the COM 1 (console) connector on the MMSC to the terminal using a 9-to-25-pin PC-style cable (with the DIN-8-to DB9 converter)—see Figure 6-8.
3. Connect a diagnostic terminal to the COM 5 (alternate console) using a 9-to-25-pin PC-style cable (with the DIN-8-to DB9 converter)—see Figure 6-8.

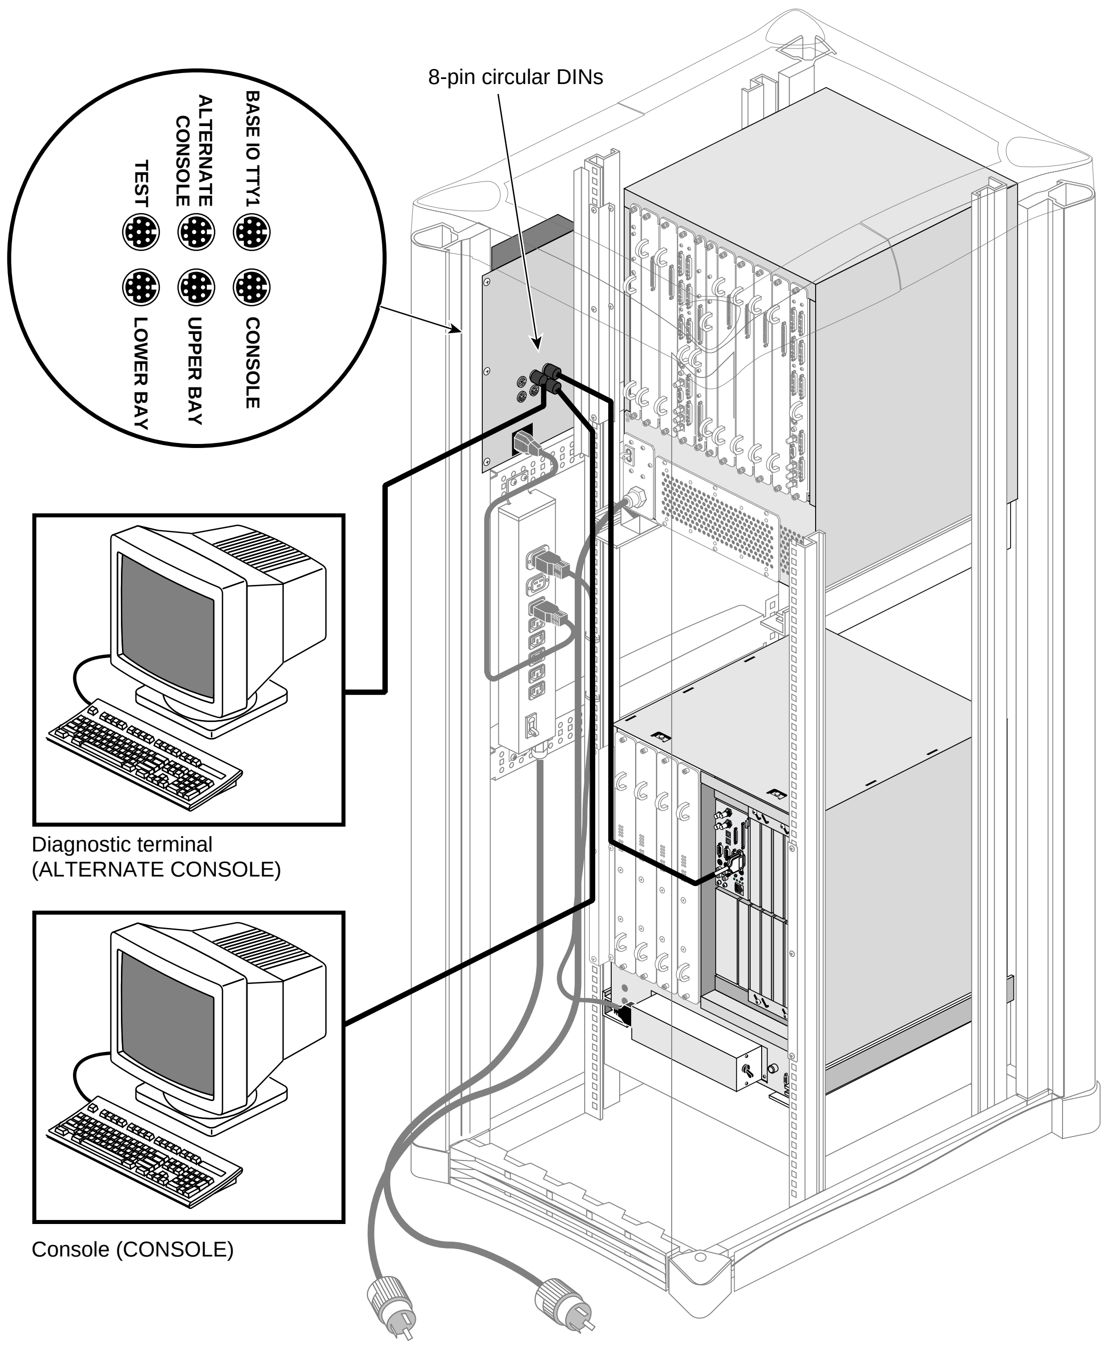

Figure 6-8 Connecting a Terminal (Rackmount System)

6.7 Setting Up the System Controller Communications (Rack Systems Only)

If you have a rackmount system, you must connect the MSC(s) to the MMSC and verify that the multimodule display is connected to the MMSC.

6.7.1 Connecting the MSC to the MMSC

The single module System Controller(MSC) also known as the entry-level System Controller—ELSC, communicates with the multimodule System Controller, also known as the full-feature System Controller—FFSC, through a serial line connection. The serial cable must have a nine-pin D-SUB (female) connector on one end and an eight-pin circular DIN (male) connector on the other end. Follow these instructions to connect the MSC to the MMSC.

1. Go to the rear of the chassis and plug the serial cable into the nine-pin D-SUB connector near the module power switch (see Figure 6-9).
2. Plug the other end of the serial cable into the MMSC (lower- or upper- "bay") connector as applicable.
3. Plug in the second serial cable from the second MSC into the MMSC, as required.
4. When you completed connecting the terminals and MSCs to the MMSC, your cabling scheme should be similar to Figure 6-10).

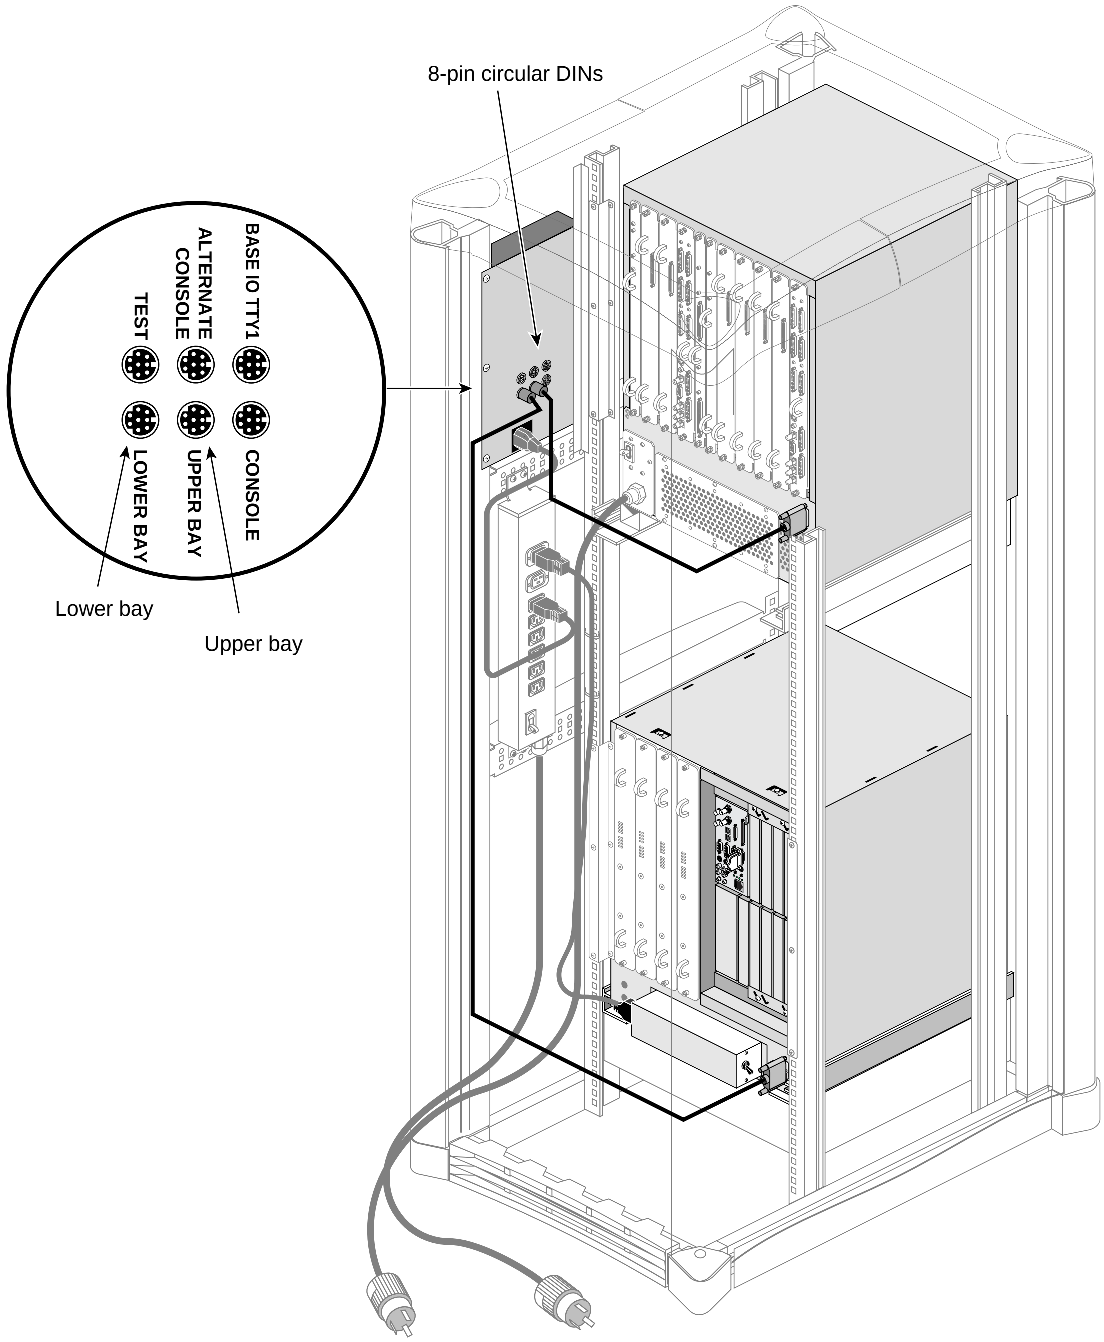

Figure 6-9 Connecting the MSCs to the MMSC

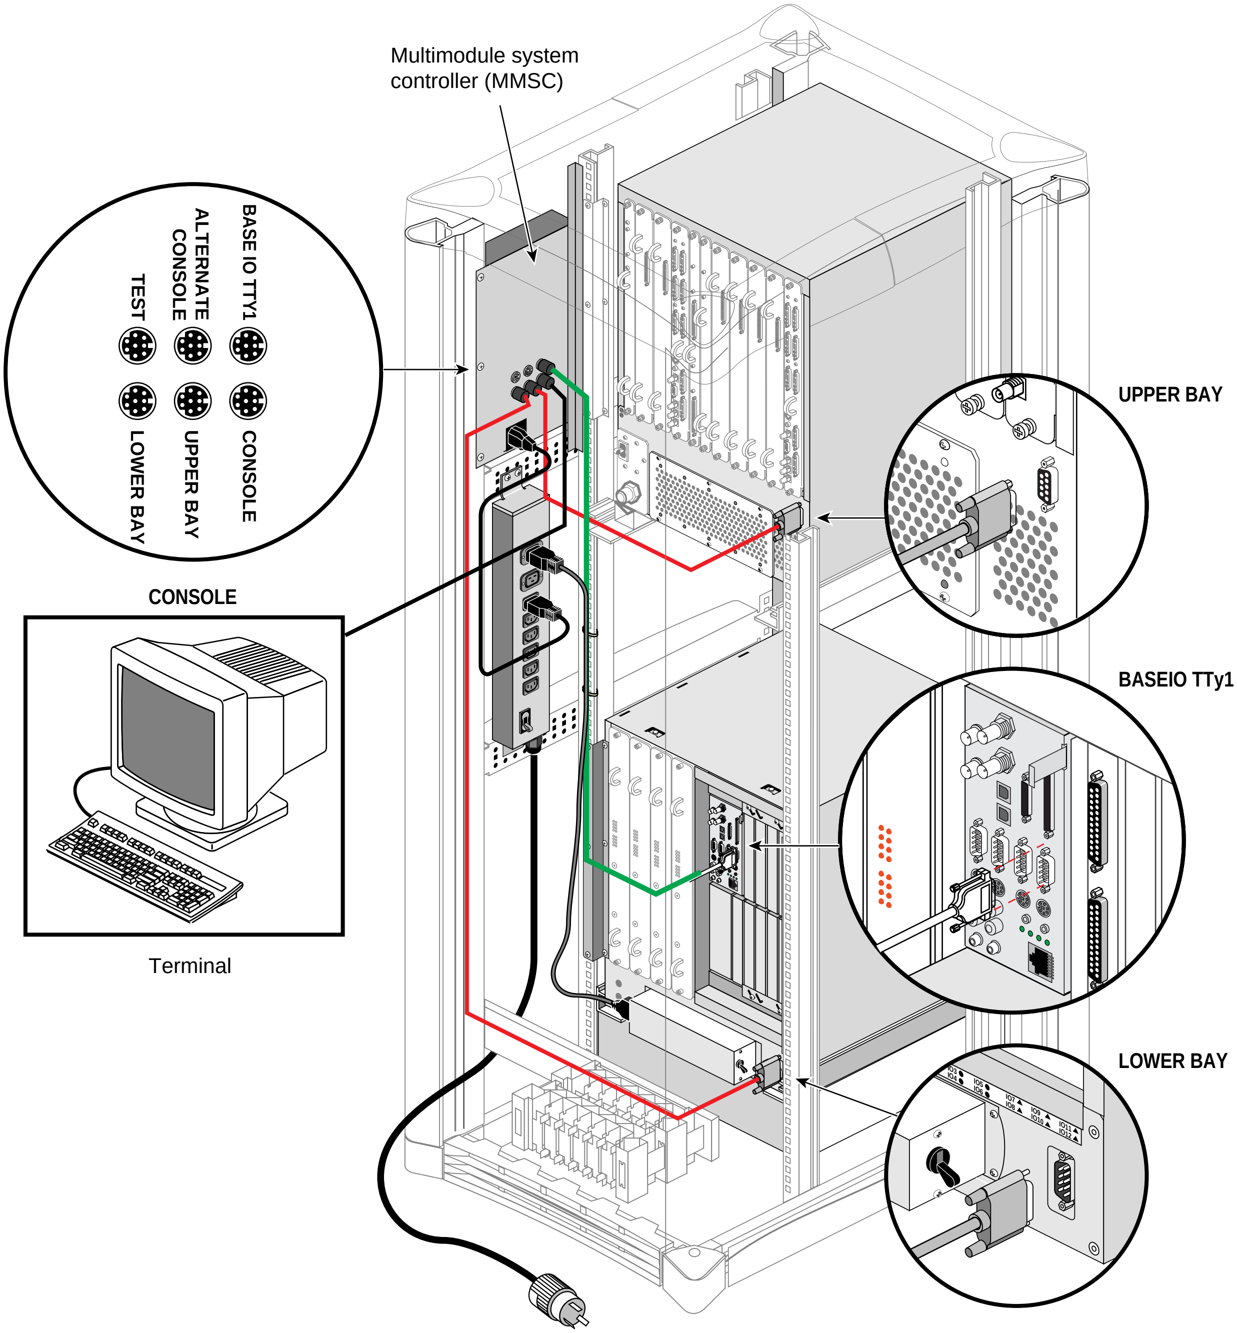

Figure 6-10 All Connections to MMSC

6.7.2 Connecting the Multimodule Display

If the multimodule display is not already connected to the MMSC, follow these procedures to connect the display. The multimodule TFP (totally flat-panel) display connects to the MMSC through a video cable that routes through the side of the chassis (see Figure 6-11).

1. Remove the side panel of the chassis (see “Removing the Rackmount Side Panels” in Chapter 10).
2. Remove the cable door as required (see “Removing the Cable Cover Door” in Chapter 10).
3. Pull away the front bezel of the display housing to access the internal video connector (see Figure 6-11).
4. Plug the cable into the video connector and reattach the bezel.

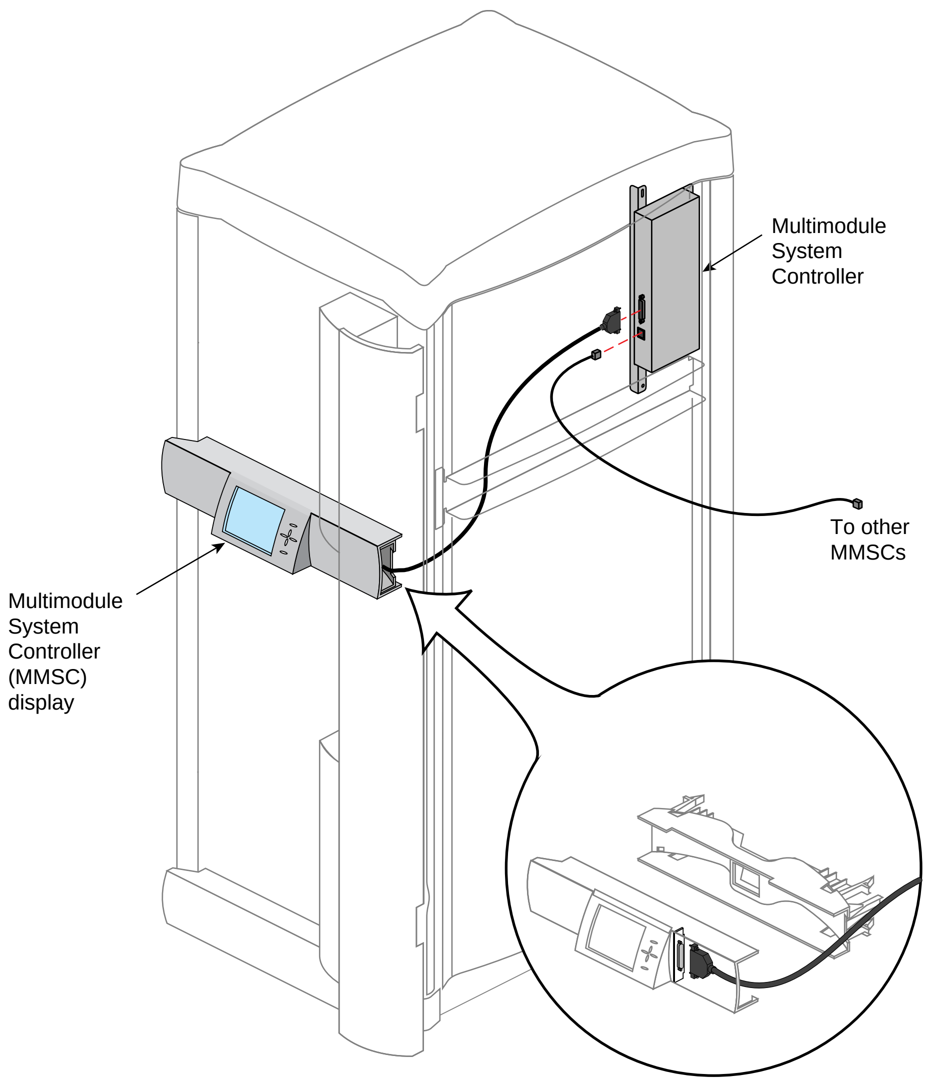

Figure 6-11 Connecting the Multimodule Controller Display to the MMSC (Side Panel Removed)

6.8 Powering On and Verifying Operation

At this stage of the installation, you should have only the keyboard, mouse, and terminal or monitor connected. *Do not* connect additional devices to the chassis at this time. If something is wrong with the system, it is easier to spot the possible problem. Follow these procedures to power up the system.

1. Plug in the system and peripheral power cords to the applicable power source. If you have a rack system, be sure to plug in the module power cords into the PDU (see Figure 6-10).

   Note: The Onyx2 graphics module and PDU require separate 220 power sources, see Figure 6-12.

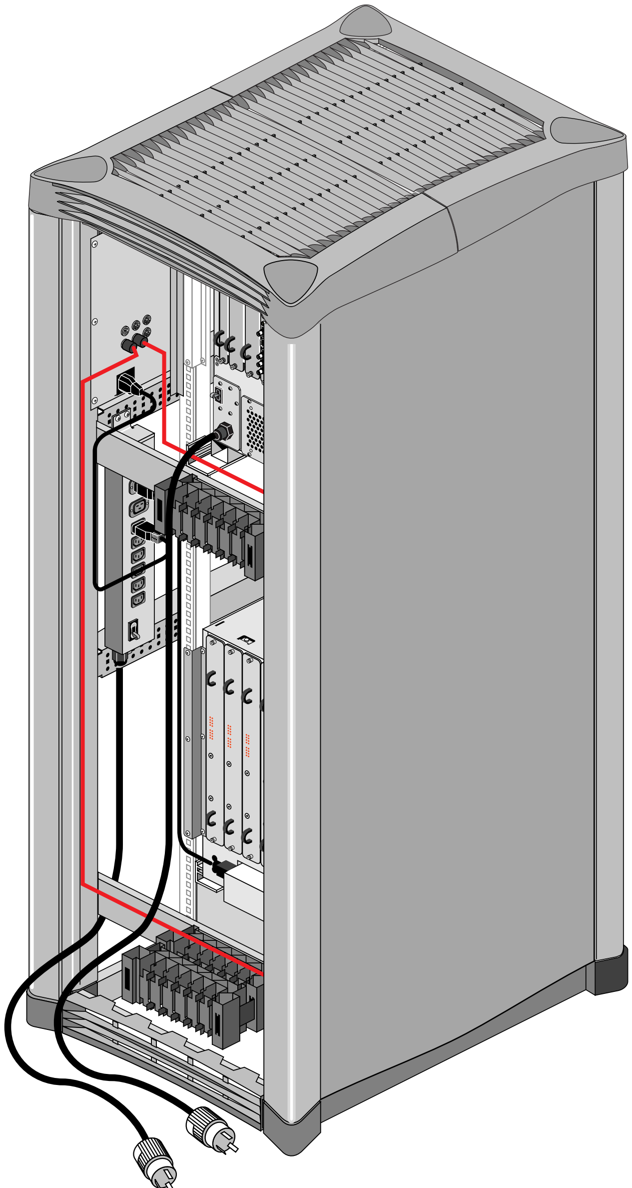

Figure 6-12 Separate Power Sources Required for Onyx2 Graphics Module and PDU

2. Power on the graphics monitor or ASCII terminal.
3. Turn the System Controller key switch to the Diag position (see Figure 6-13).

   Note: The System Controller powers on in either in the Diag or On position; however, it is recommended that you select Diag. If a problem occurs on power-on, the Diag mode helps to isolate the cause by displaying status messages on a diagnostic terminal connected to the MMSC **COM 5** port or to the **Diagnostic Port** on the MSC.

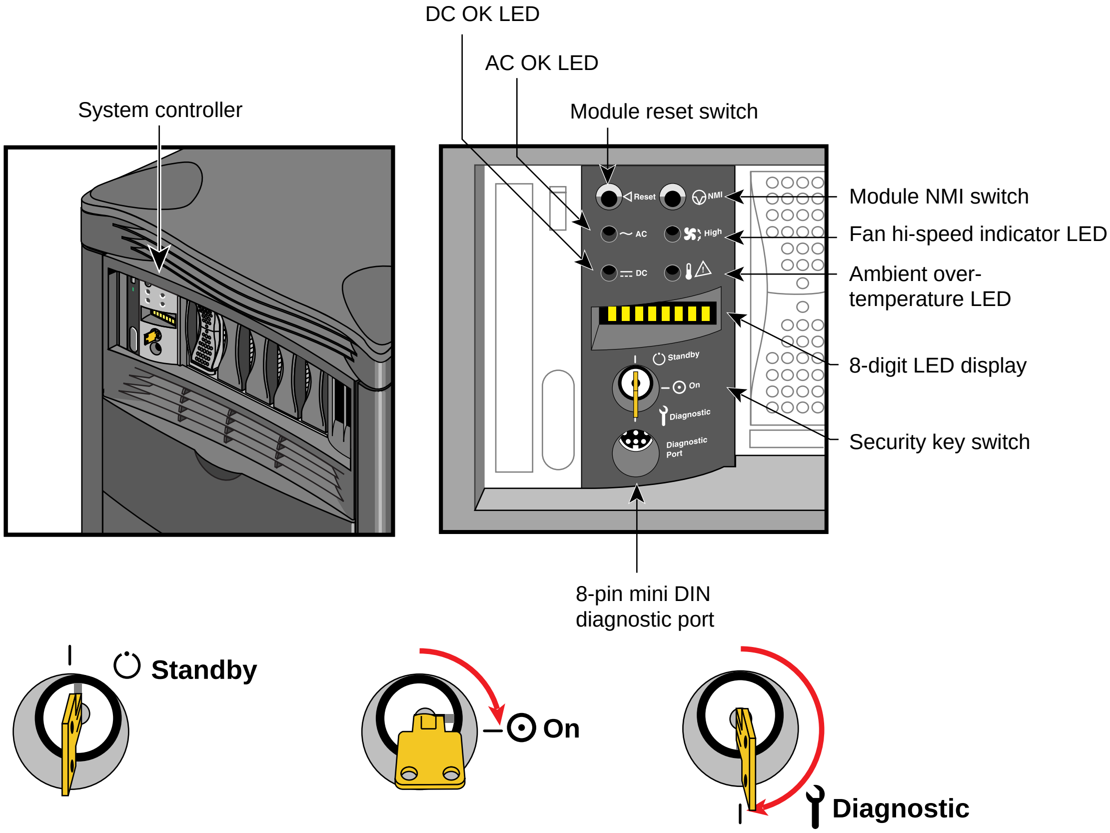

Figure 6-13 Module System Controller Key Positions

4. Turn on the power switch(es) on the back of the unit.
   - If you have a rack system, turn on the PDU first and then turn on the individual modules (see Figure 6-14).
   - If you have an Onyx2 rack system, it is a good habit to power on the graphics module before the server module (see Figure 6-15). This helps enable the operating system to recognize the graphics subsystem when IRIX is booting.
   - If you have a deskside, see Figure 6-16.

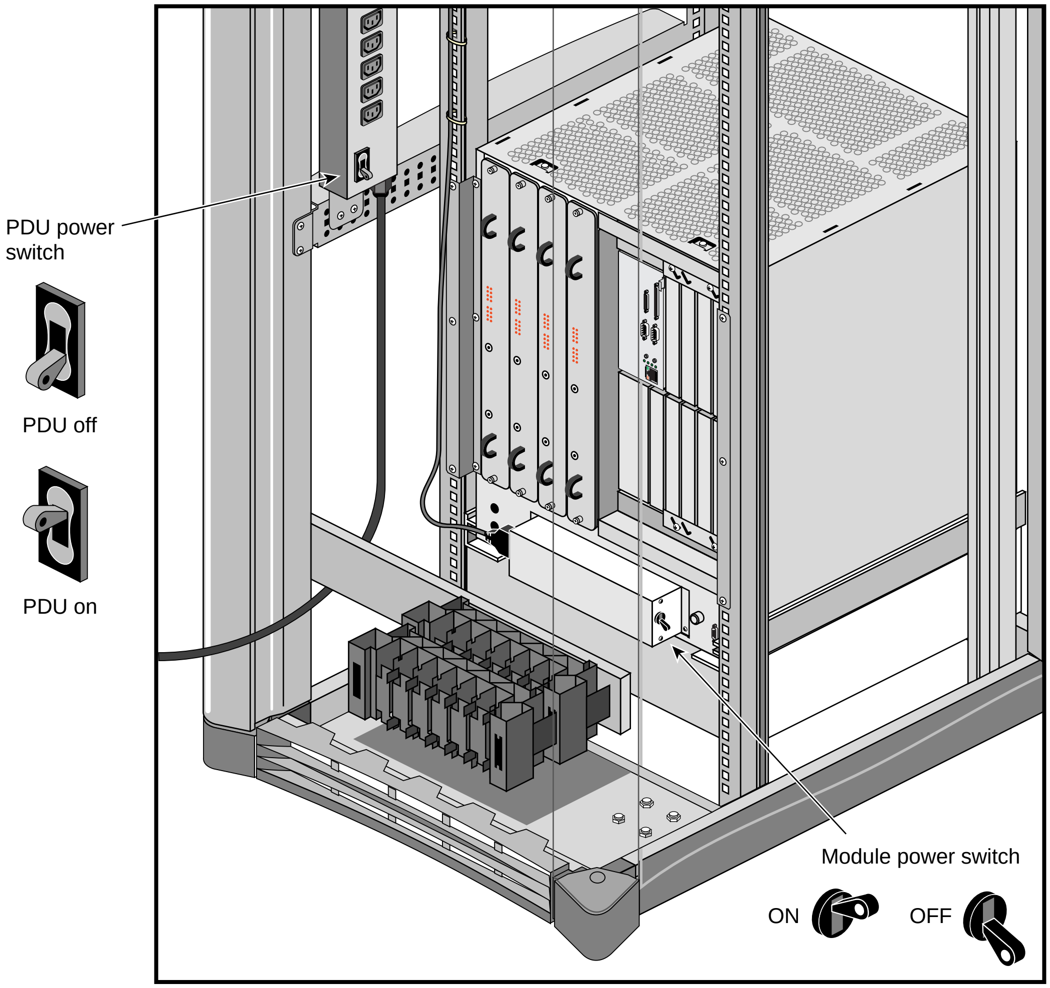

Figure 6-14 Turning On the PDU and Module(s) on a Rack System

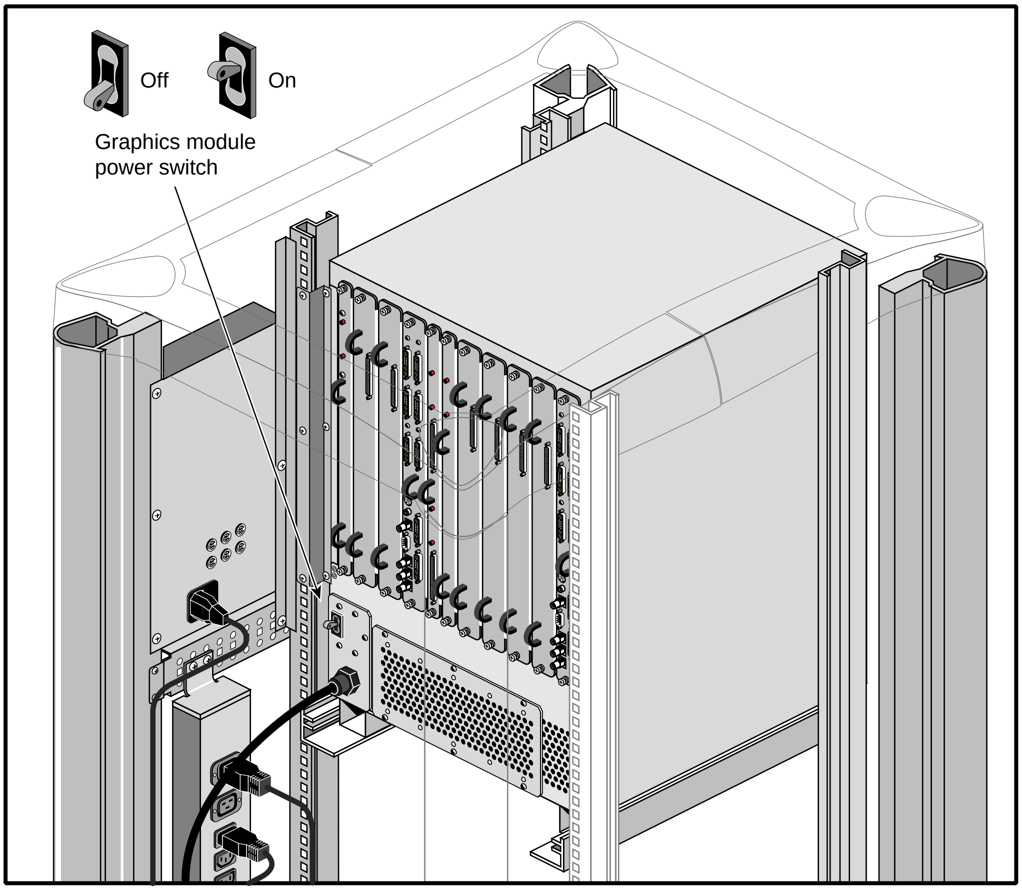

Figure 6-15 Onyx2 Graphics Module Power Switch Location

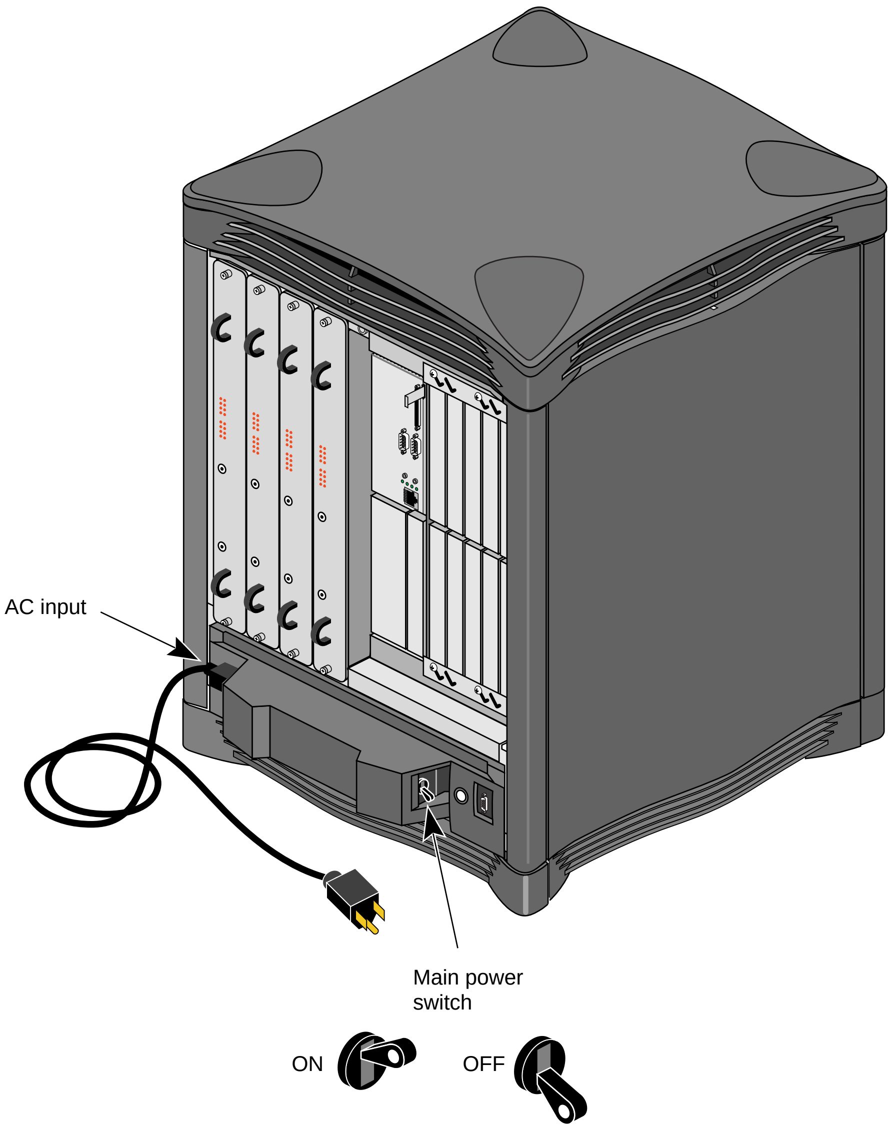

Figure 6-16 Powering On the Deskside System (Rear View of Chassis)

5. If you have a rack system, boot the system using the MMSC display.

   - Select the Focus menu on the display using the direction keys, then press the Enter key. Make sure that the All option is highlighted

   - Go to the "Action" menu on the display using the direction keys (see Figure 6-17).

   - Select the Power Up option and press the Enter key.

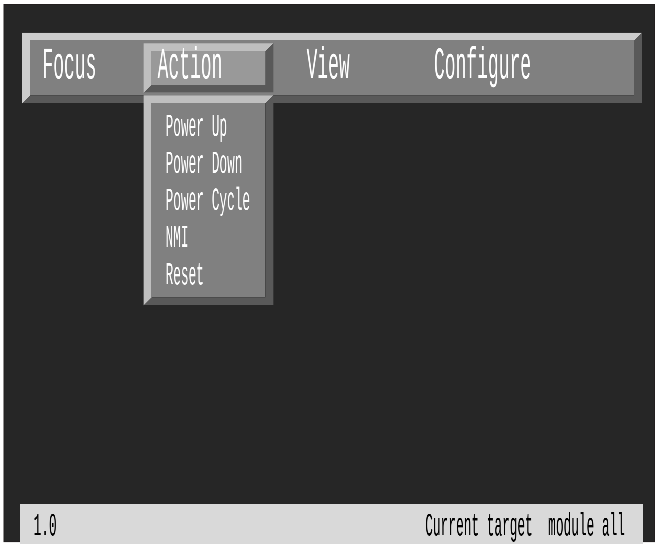

Figure 6-17 Action Menu Selection for MMSC Display

6. If desired, use the System Controller menu and display to watch the boot process status display.

   `POWER ON`

   `VERS x.xx` (PROM version number)

   `MOD` (module number, such as MOD 1, etc.—see Figure 6-18).

   Note: With a multimodule rack configuration, the module number assignment is arbitrary. For example during power-on, the upper module may display the number 0 and the lower number module may indicate 1. The next time you power on the system or change the hardware configuration, the module numbers can change. In a multimodule configuration, the numbers should always differ from module to module.

   A future release of the MSC PROM will enable you to assign fixed module numbers that do not change unless you manually alter them. Appendix D describes commands to reassign the module numbers.

   Module numbers are used by different software command programs (such as *hinv* and *hwgraph*) to identify and call out a specific module for status reporting (see Chapter 7, “Diagnostic Tools,”). The MMSC command set (see Appendix C) identify the modules according to the position in the rack—upper and lower.

7. Before leaving the customer site, turn the MSC keyswitch to the On position.

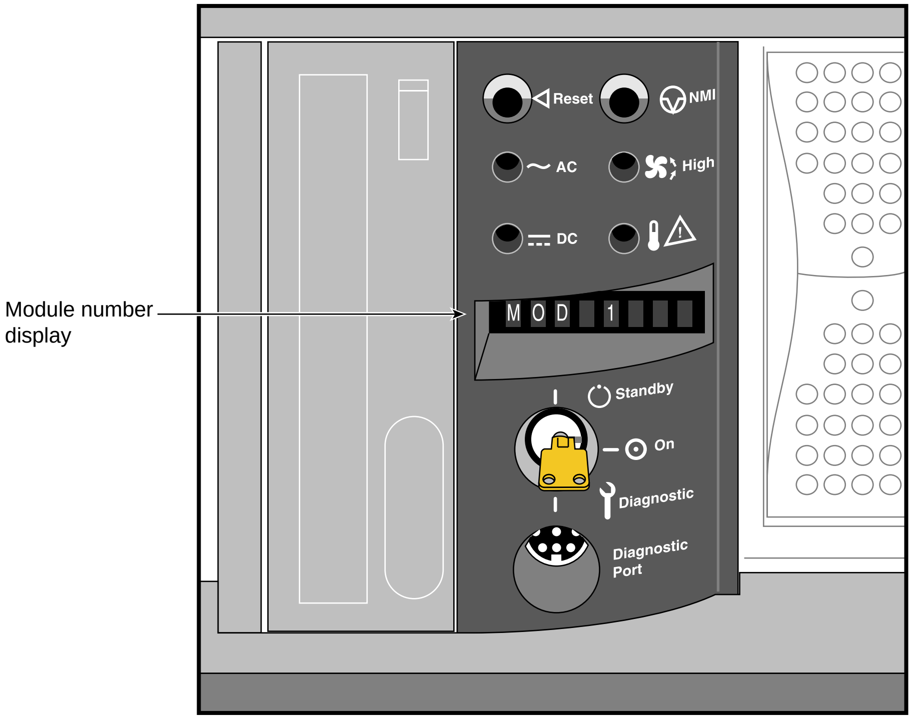

Figure 6-18 Module Number Display at Power-Up

8. When the power-on diagnostics have been completed, you will see this message on the monitor or terminal:

```
Starting up the system ...
To perform System Maintenance instead, press <Esc>
```

9. Go the PROM monitor menu.

10. At the monitor prompt, type **hinv** to check and verify the hardware configuration of the system. You should get a display similar to the following:

```
FPU: MIPS R10010 Floating Point Chip Revision: 0.0
CPU: MIPS R10000 Processor Chip Revision: x.x
2 200 MHZ IP27 Processors
Main memory size: 128 Mbytes
Instruction cache size: 32 Kbytes
Data cache size: 32 Kbytes
Secondary unified instruction/data cache size: 1 Mbyte
Integral SCSI controller 0: Version QL1040A
  Disk drive: unit 1 on SCSI controller 0
Integral SCSI controller 1: Version QL1040A
Integral Fast Ethernet: ef0, version 1
```

See also “NIC Reporting” in Chapter 7 for other uses of the *hinv* command.

Note: Some hardware options (such as a network board) may not be recognized by the PROM monitor. These options are not listed by the *hinv* command at this time. However, after you load and boot the operating system, the *hinv* command (executed at the IRIX prompt) should list *most* of the system hardware options.

6.9 Loading the Single Connector Assembly (SCA) Drives Into the System

The Origin2000 and Onyx2 systems use an SCA sled assembly for mounting into the drive bay. Disk drive modules are aligned vertically at the front of the chassis, as shown in Figure 6-19. In the server systems and Onyx2 rackmount, note that the leftmost disk drive—the system drive—is oriented differently from the others. All the disks in an Onyx2 deskside system have the same orientation.

Note: The systems use single-ended SCSI drives only.

To insert a disk module, follow these steps:

1. If necessary, snap the handle to the open position so that it is centered, as shown in Figure 6-19.
2. If you are adding a drive, remove the drive filler plate that covers the drive slot you want to use.
3. Align the new disk module with the drive guide, as shown in Figure 6-19.
4. Gently but firmly slide the disk module on the guides over the pin.
5. When the disk module is in all the way, snap the handle to the right to the closed position, as shown in Figure 6-19. In the case of the system disk module, which is upside down relative to the other drives, snap the handle to the left.

The SCSI IDs for the SCA drives are hard-wired into the Origin2000 and Onyx2 midplanes. Figure 6-20 shows the assigned addresses. The SCSI drives are terminated directly on the midplane.

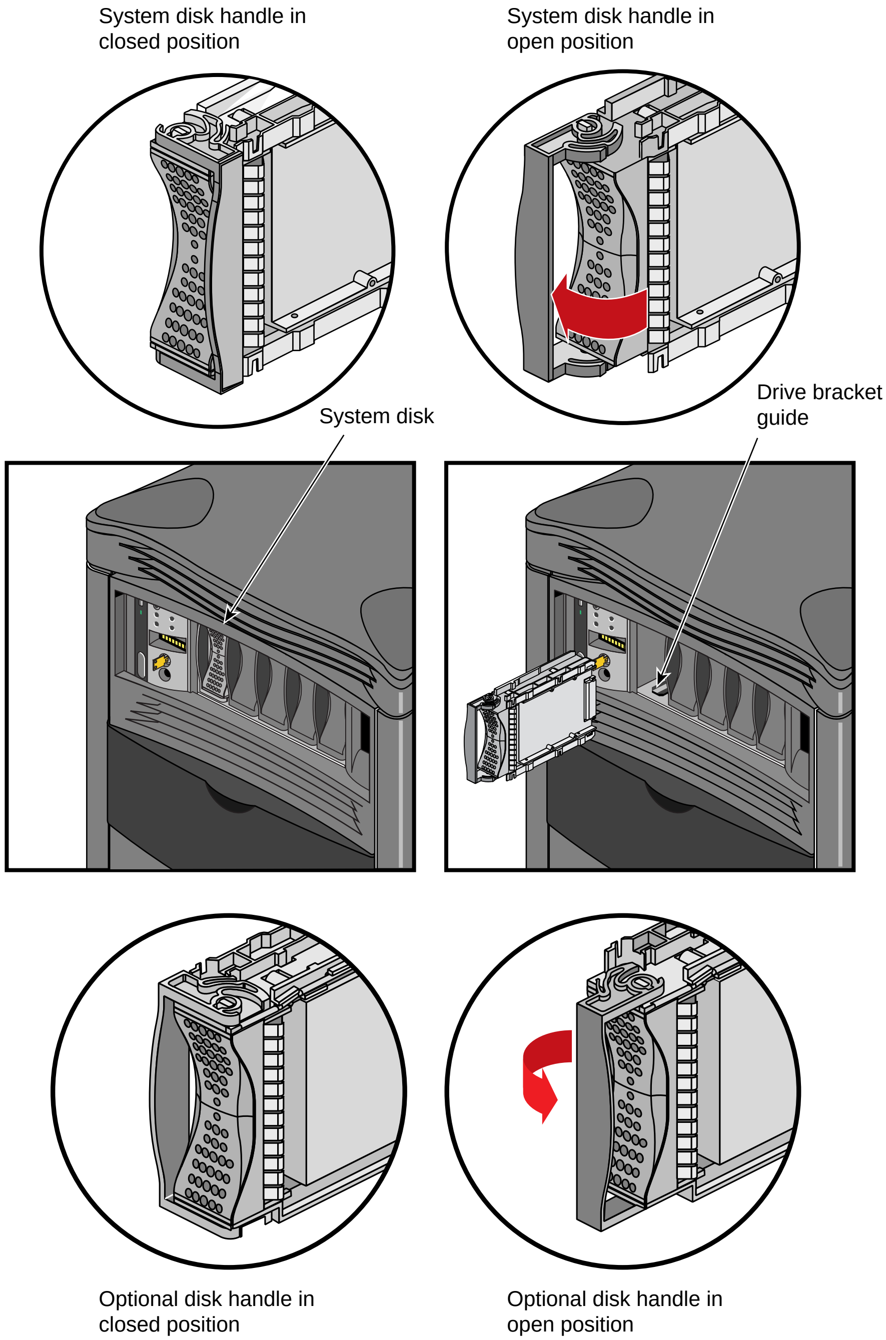

Figure 6-19 Installing an SCA Drive

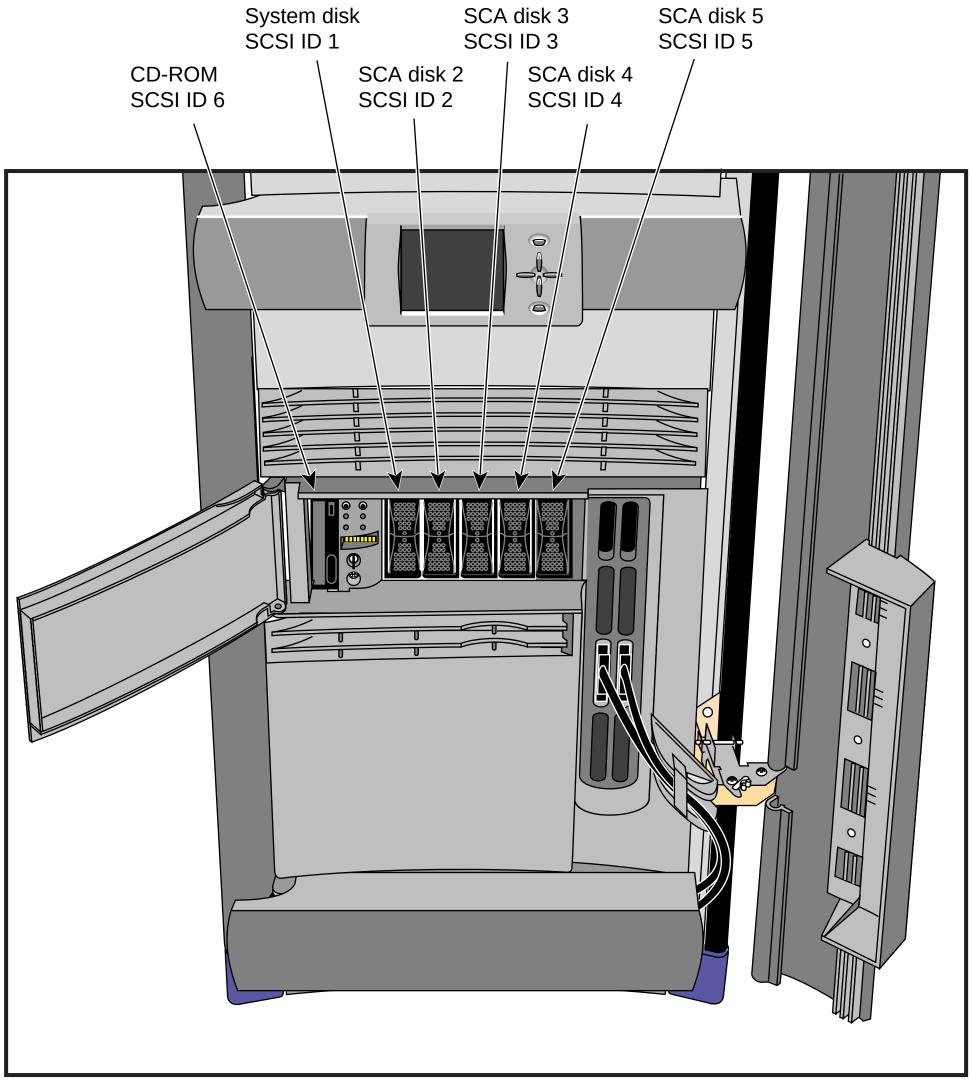

Figure 6-20 SCSI Hardwire Addresses for the Origin2000 and Onyx2 Chassis (Rackmount or Deskside)

6.9.1 Loading the CD-ROM

The CD-ROM installs vertically into the chassis. Follow these instructions to load the CD-ROM into the drive:

1. Press the eject button to release the CD-ROM carrier (see Figure 6-21).
2. Hold the CD-ROM disc around the edges and place it into the holder clips as shown in Figure 6-21.
3. Insert the CD-ROM carrier into the drive.

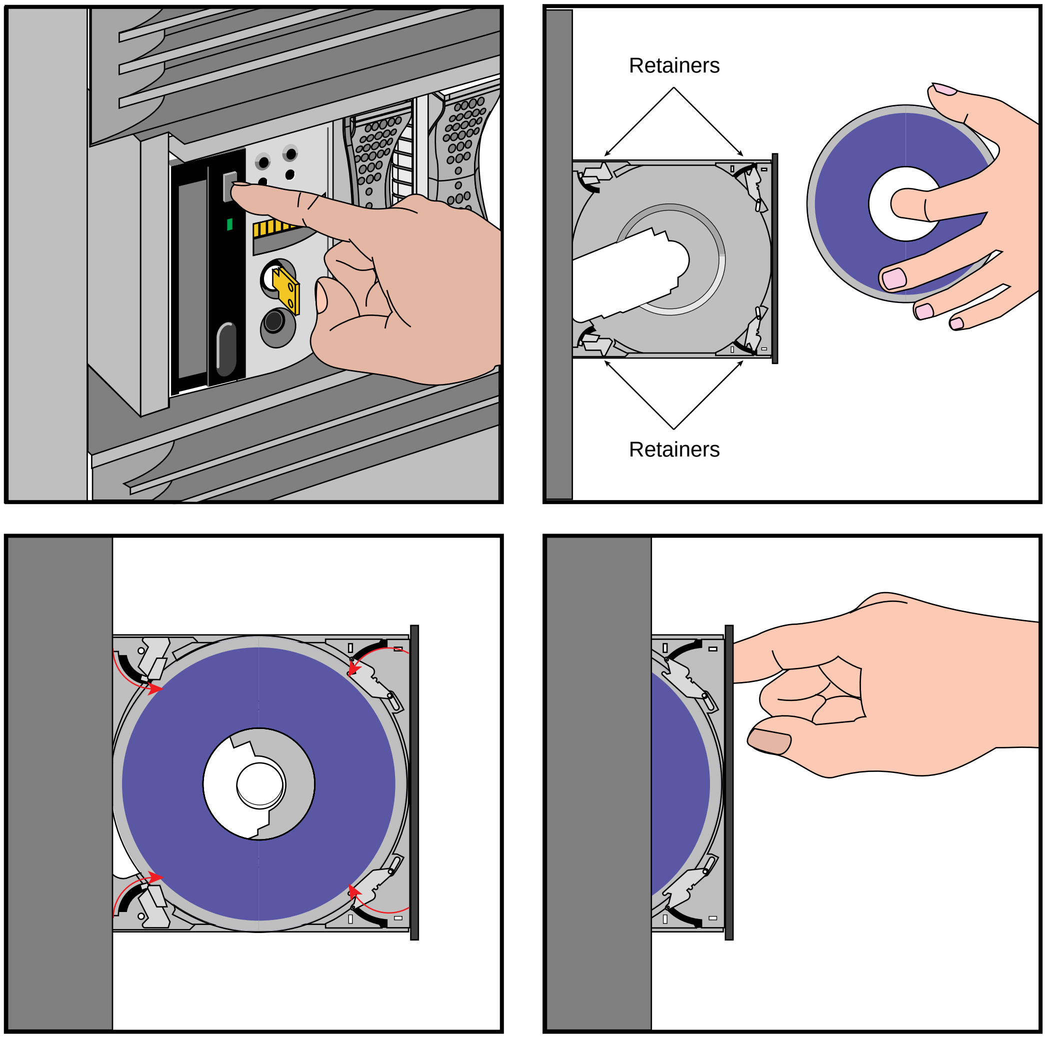

Figure 6-21 Installing a CD Into the Drive of an Origin2000 and Onyx2 Chassis

6.10 Installing XIO Boards

The XIO boards are installed in the rear of the chassis (see Figure 6-22). For information on installing these boards, see “Removing and Installing an XIO Board” in Chapter 10.

Caution: The XIO boards have strict installation guidelines that must be observed. Be sure to read “Removing and Installing an XIO Board” in Chapter 10 before attempting to install an XIO board into the chassis.

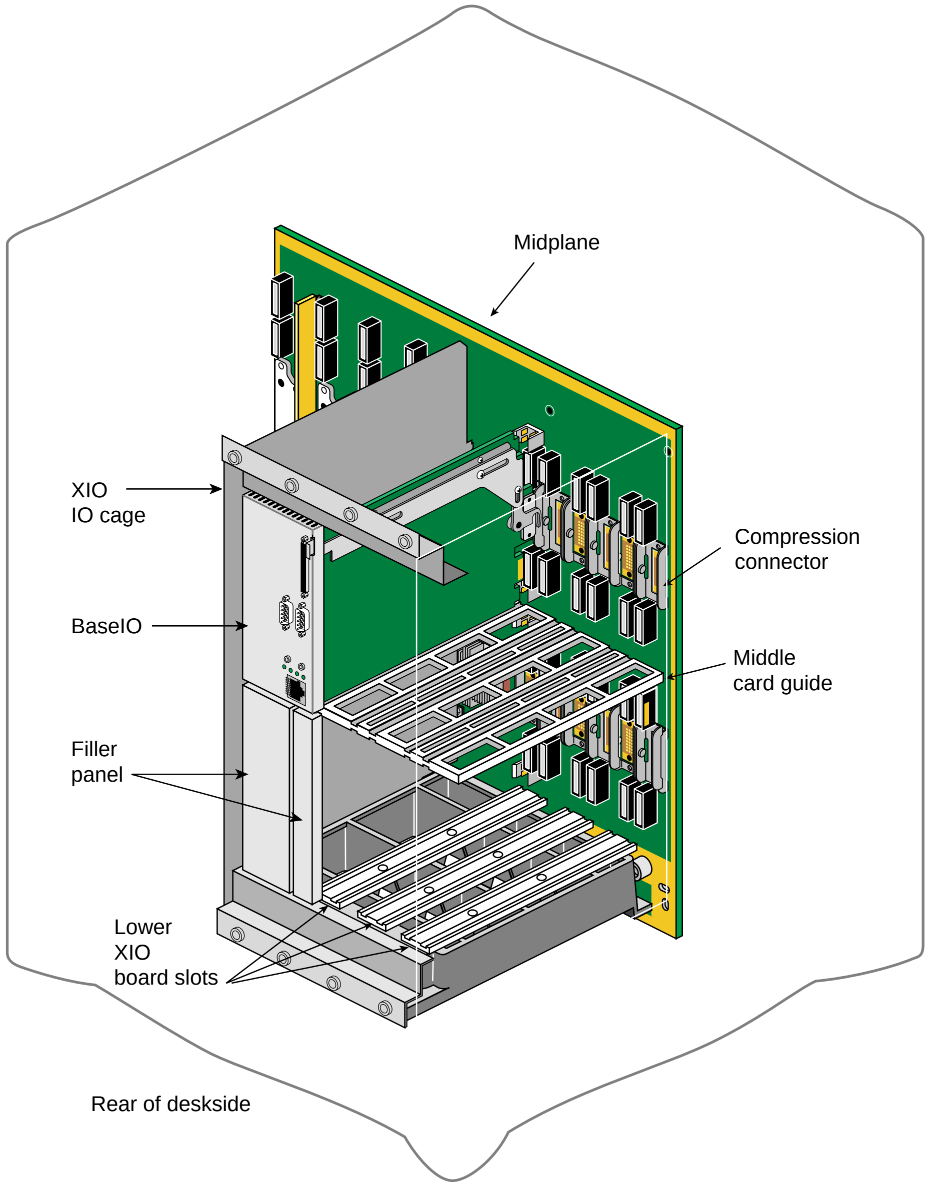

Figure 6-22 XIO Board Slots (Rear of Chassis)

6.11 Optional Equipment

This section describes how to connect optional equipment such as a modem and a printer to the Origin2000 and Onyx2 systems.

6.11.1 Printer Cabling

The serial printer connection is a 9-pin serial port located directly on the BaseIO board. See Chapter 3, "Chassis Tour" and Appendix B, "Connector Pinouts," for information on pin-outs.

Use the system manager's "Printers" tool to configure the software to recognize the printer. Refer to the *IRIX Admin: Peripheral Devices* manual or for more information about printer configuration.

6.11.2 Modem Cabling

To attach a modem to the system, connect a modem cable to one of the nine-pin serial connectors on the BaseIO panel. This connection may require a 9-pin to 25-pin adapter cable. Be sure that the pin assignments on the cable match the pin assignments on the modem. See Appendix B, "Connector Pinouts,"for information on pinouts.

Note: This product requires the use of external shielded cables to comply with Part 15 of the FCC rules. Serial cables from different vendors may not be compatible.

Refer to the *IRIX Admin: Peripheral Devices* manual for more information about modem configuration.

6.12 Software Installation

Refer to the *IRIX Admin: Software Installation and Licensing* manual for the detailed steps required to install the software.

Note: A copy of the IRIX operating system is supplied with the system on a compact disc. Place the CD in a secure place in case you ever need to reinstall the operating system.

6.13 Configuring the Monitor

At this point of the installation, you can now check or change the monitor configuration if desired.

Note: For information on using the *ircombine* program, see the *InfiniteReality Video Format Combiner User's Guide* (P / N 007-3279-xxx).

6.13.1 Checking the Monitor Resolution

Use the following command to determine the screen resolution:

```
/usr/gfx/gfxinfo
```

You should see a message similar to the following:

```
Graphics board 0 is “HILO” graphics.
Managed (“:0.0”) 1280x1024
Display has 2 channels
8 GEs (of 8), occmask - 0x0f
2 RM7 boards (of 2) 1/1/0/0
Texture Memory: 64MB/64MB/-/-
Medium pixel depth
32K cmap
```

6.13.2 Changing the Monitor Resolution

After you load IRIX, you can use the *setmon* command to change the resolution or video output format (VOF). This example shows you how to change the VOF to 1024 x 768:

```
su
setenv DISPLAY :0
/usr/gfx/setmon -x 1024x768_60
/usr/gfx/stopgfx;/usr/gfx/startgfx
```

Note: The *setmon* command loads the VOF value into EEPROM so that the system boots in the specified display mode.

For additional information on monitor and graphics settings, see the following reference (man) pages:

- setmon(1g)
- gfxinfo(1g)
- gfxinit(1g)
- stopgfx(1g)
- startgfx(1g)

The *setmonitor* command has been superseded by the *XSGIvc* library.

6.14 CrayLink Interconnect Cabling

This section describes how to install the CrayLink Interconnect cable on 16P, 24P, 32P, and 64P configurations.

Caution: You cannot install CrayLink Interconnect cabling on a standalone deskside system because of power grounding requirements. The power differential between two interconnected modules should not exceed 500 millivolts; otherwise severe damage can result to boards and other components inside the chassis. The power distribution unit (PDU) inside the rack provides a common ground source for the modules. In addition, groundstraps are installed on multirack configurations to help provide common grounding across the racks as described in Section 6.16.

6.14.1 CrayLink Interconnect Cable and Router Handling Considerations

The CrayLink Interconnect cables contain delicate copper fibers and should be handled with care. Observe the following guidelines:

- Avoid overbending the CrayLink Interconnect cables (not more than 1.25-inch radius).
- Avoid stepping on the CrayLink Interconnect cables.
- Avoid "hotplugging" in or removing CrayLink Interconnect cables while both Origin2000 modules are up and running; plugging in or removing a CrayLink Interconnect cable hangs or crashes the entire system when the modules are powered on.
- Avoid touching, bumping, or scraping the compression connector pads on the Router board. This prevents a build-up of contaminants and possible damage to the connector pads.
- Be sure to align the cable connectors properly to the router ports before tightening down the bolts. If the connector is installed crookedly, you can damage the entire cable.

6.14.2 CrayLink Cabling Configuration Illustrations

Figures 6-24 through 6-35 show two views of each configuration (16P, 24P, 32P, and 64P), both with and without Xpresslinks. The first view provides an overview of cabling runs, and the second shows the exact route the each CrayLink cable must follow. Refer to Table 6-4 to find the figures that illustrate the cabling configuration that you need.

Table 6-3 CrayLink Cabling Overview and Routing Illustrations

Figure	Configuration Shown	Note
Figure 6-23	CrayLink cables and router board ports	
Figure 6-24	16P cabling overview	
Figure 6-25	16P cabling routing	
Figure 6-26	16P with Xpresslinks cabling overview	Xpresslinks cross
Figure 6-27	16P with Xpresslinks cabling routing	Xpresslinks cross
Figure 6-28	24P with Xpresslinks cabling overview	Xpresslinks cross
Figure 6-29	24P with Xpresslinks cabling routing	Xpresslinks cross, keep top two comb slots unused
Figure 6-30	32P cabling overview	
Figure 6-31	32P cabling routing	Keep top two comb slots unused
Figure 6-32	32P with Xpresslinks cabling overview	Xpresslinks cross
Figure 6-33	32P with Xpresslinks cabling routing	Xpresslinks cross, keep top two comb slots unused, connect icon pairs in illustration

Table 6-3 (continued) CrayLink Cabling Overview and Routing Illustrations

Figure	Configuration Shown	Note
Figure 6-34	64P cabling overview	
Figure 6-35	64P cabling routing	Keep top two comb slots unused, connect icon pairs in illustration
FIgure 6-36	24P Onyx2 configuration with two graphics pipes	Link 2 CrayLink crossover

6.14.3 CrayLink Cabling Installation Procedure

As shown in Figure 6-23, each router board has three ports labeled Link 1, Link 2, and Link 3. The Origin modules shown in Figures 6-23 thorough 6-35 each have two router boards. Figure 6-23 shows a Router A and Router B to simplify textual descriptions.

1. Find the drawing set that applies to your installation.
2. Check to make sure that you have sufficient quantities of CrayLink cables, up and down bails, bail clips and comb covers to complete the installation.
3. Arrange the CrayLink cables into groups.
4. The 58” cables for the vertical cabling runs are installed first. Connect one end of the 58” cable to the Router A’s Link 1 router port on the upper module. String the cable vertically, and connect the other end of the cable to Router A’s Link 1 on the lower module. Repeat this procedure for each additional rackmount chassis.

Note: Leave the top two cable comb slots empty for future expansion.

5. The 84” or 108” cables for the diagonal Xpresslinks are installed next (if applicable). Check the illustrations showing Xpresslinks. Connect one end of the cable from Router A’s Link 2 port to Router B’s Link 2 port.

Note: Xpresslinks must crossover.

6. The 58” cables for the horizontal cabling runs on 32 P and 64 P configurations are installed last. Connect one end of the 58” cable to the Router A’s Link 3 router port on Rack 1. String the cable horizontally, and connect the other end of the cable to Router A’s Link 3 port on Rack 2. Repeat this procedure for each additional rackmount chassis.

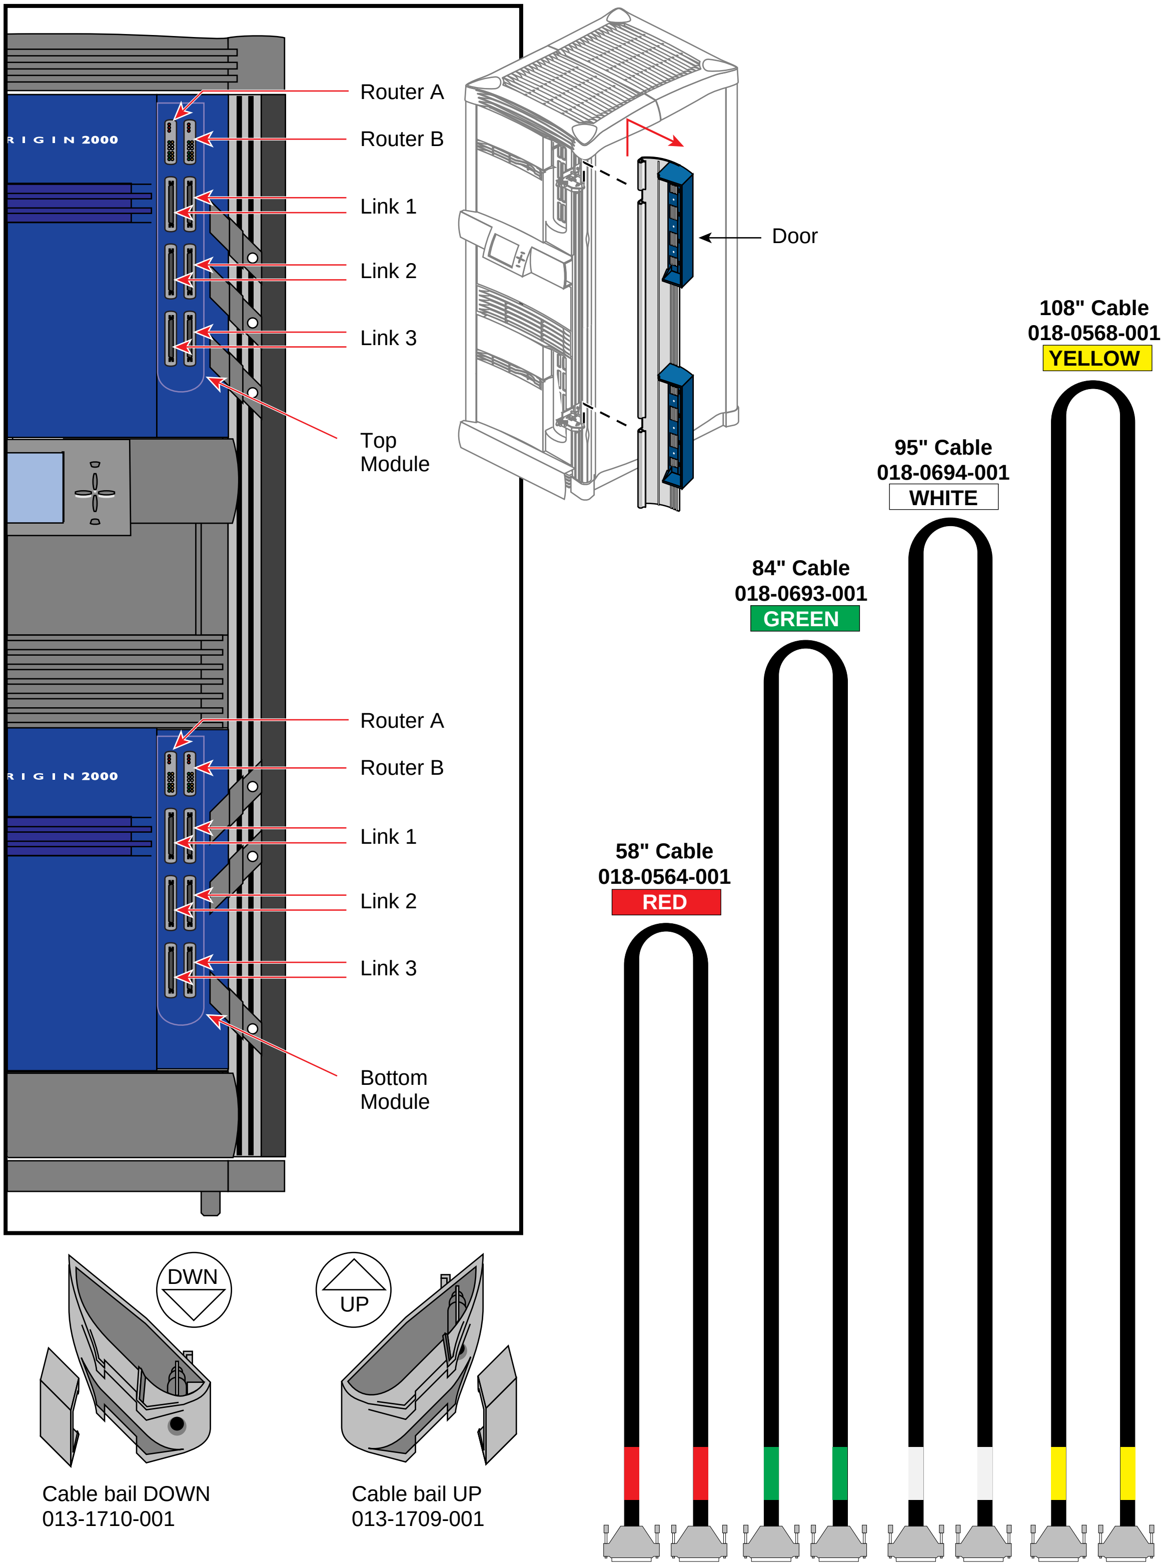

Figure 6-23 CrayLink Cables and Router Board Ports

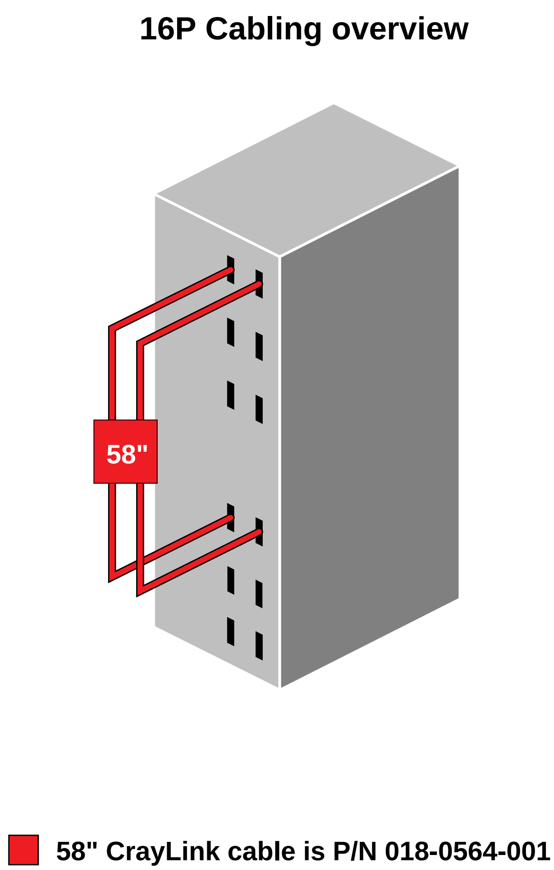

Figure 6-24 16P Cabling Overview

16P configuration

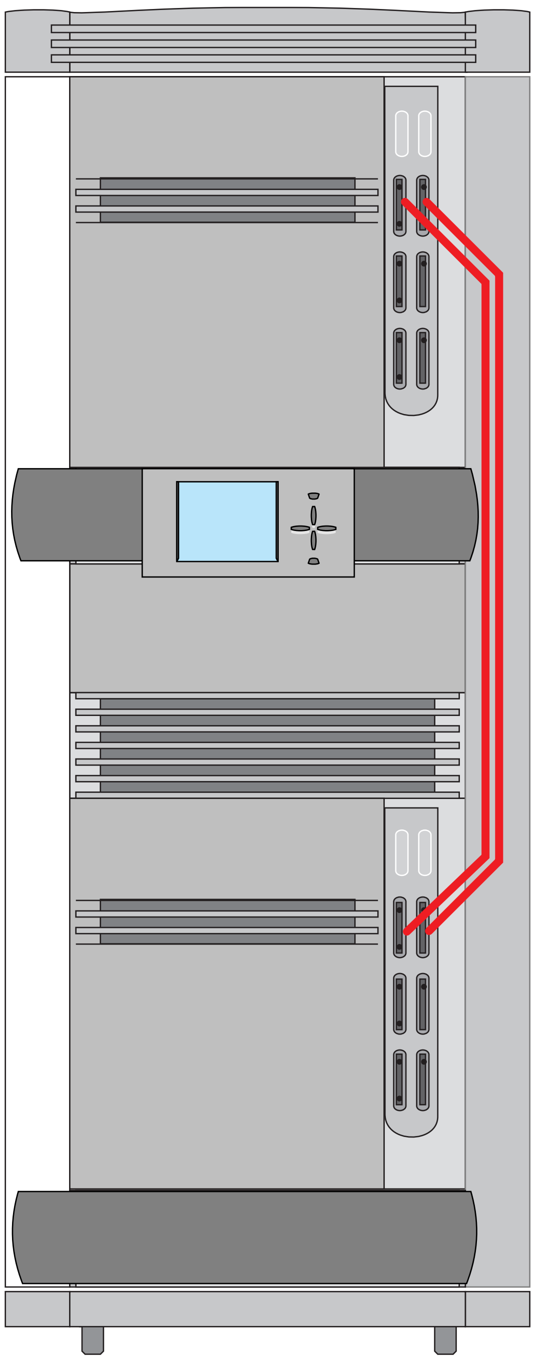

58" CrayLink cable is P/N 018-0564-001

Figure 6-25 16P Cabling Routing

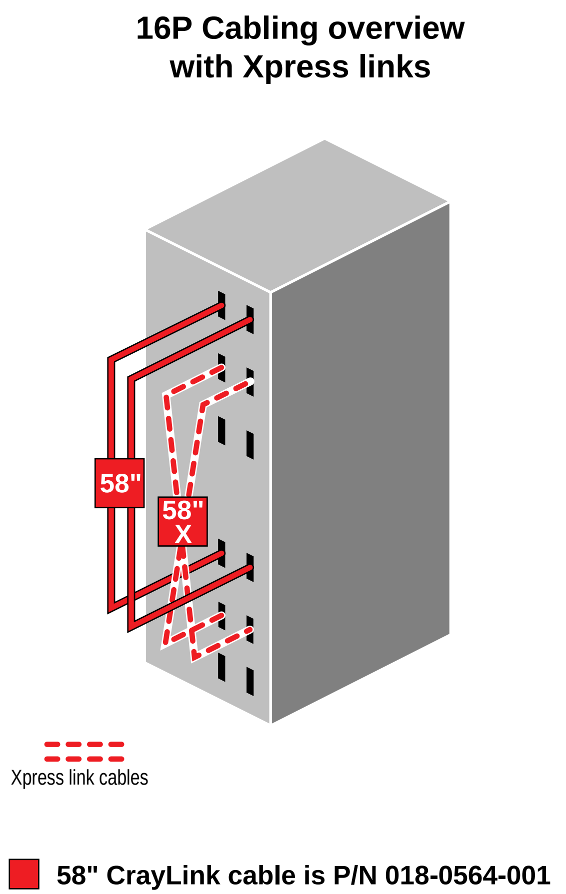

Figure 6-26 16P with Xpresslinks Cabling Overview

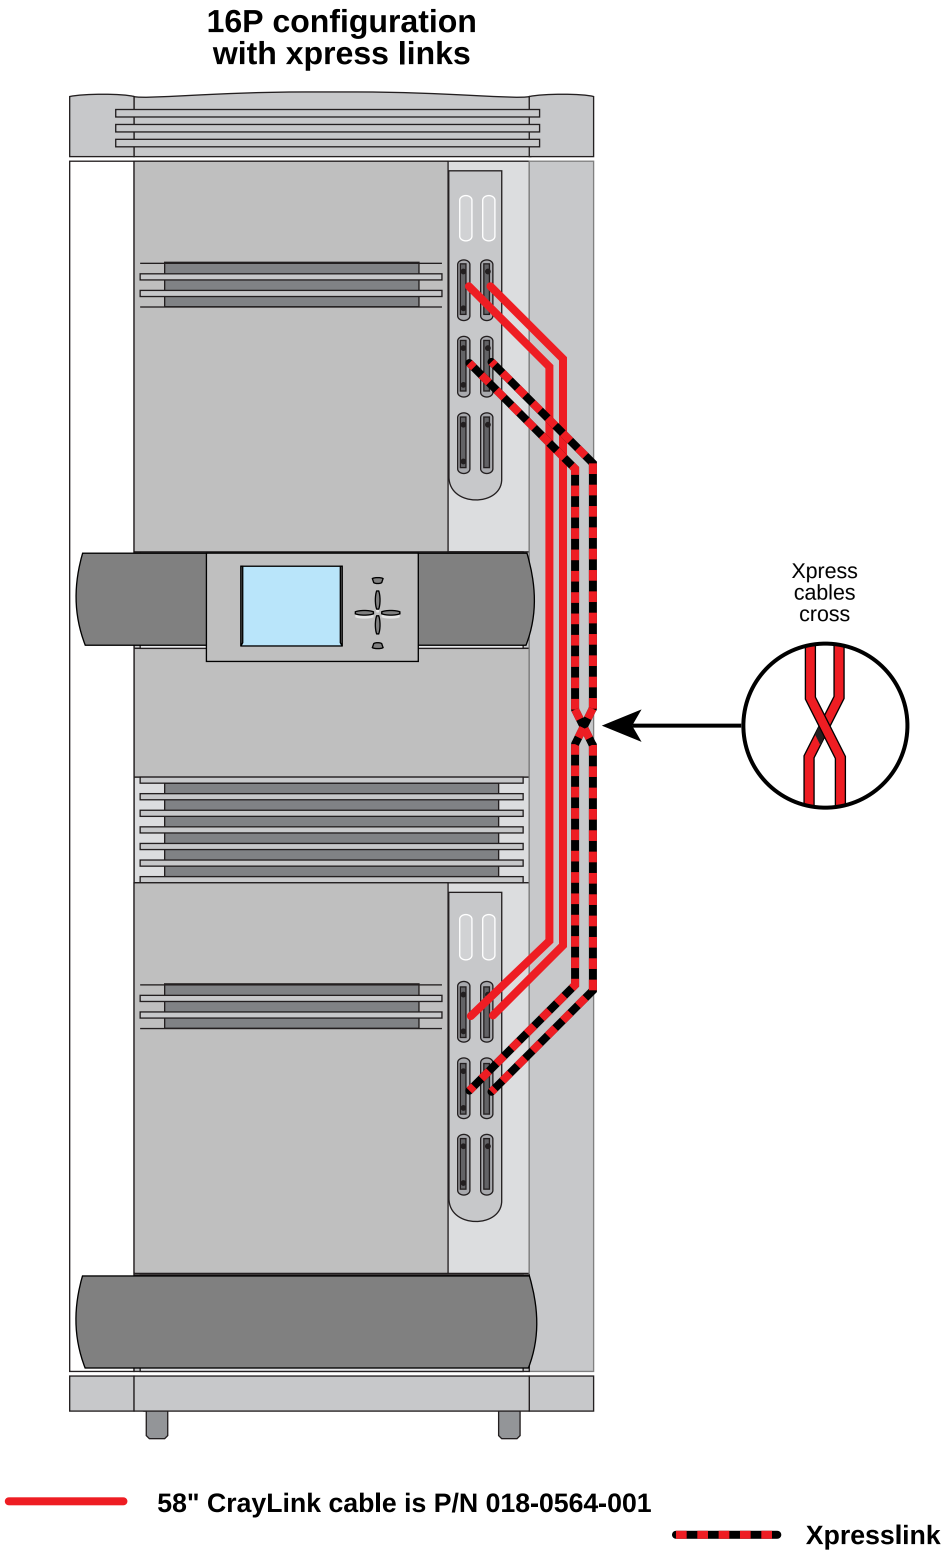

Figure 6-27 16P with Xpresslinks Cabling Routing

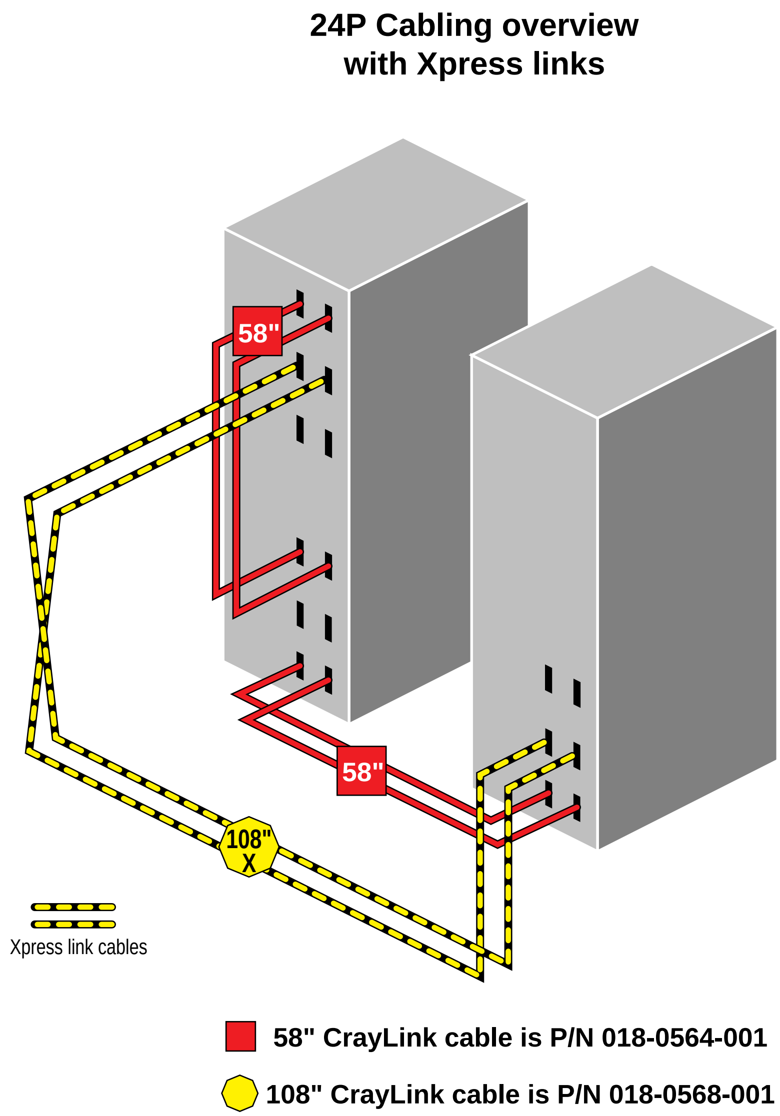

Figure 6-28 24P with Xpresslinks Cabling Overview

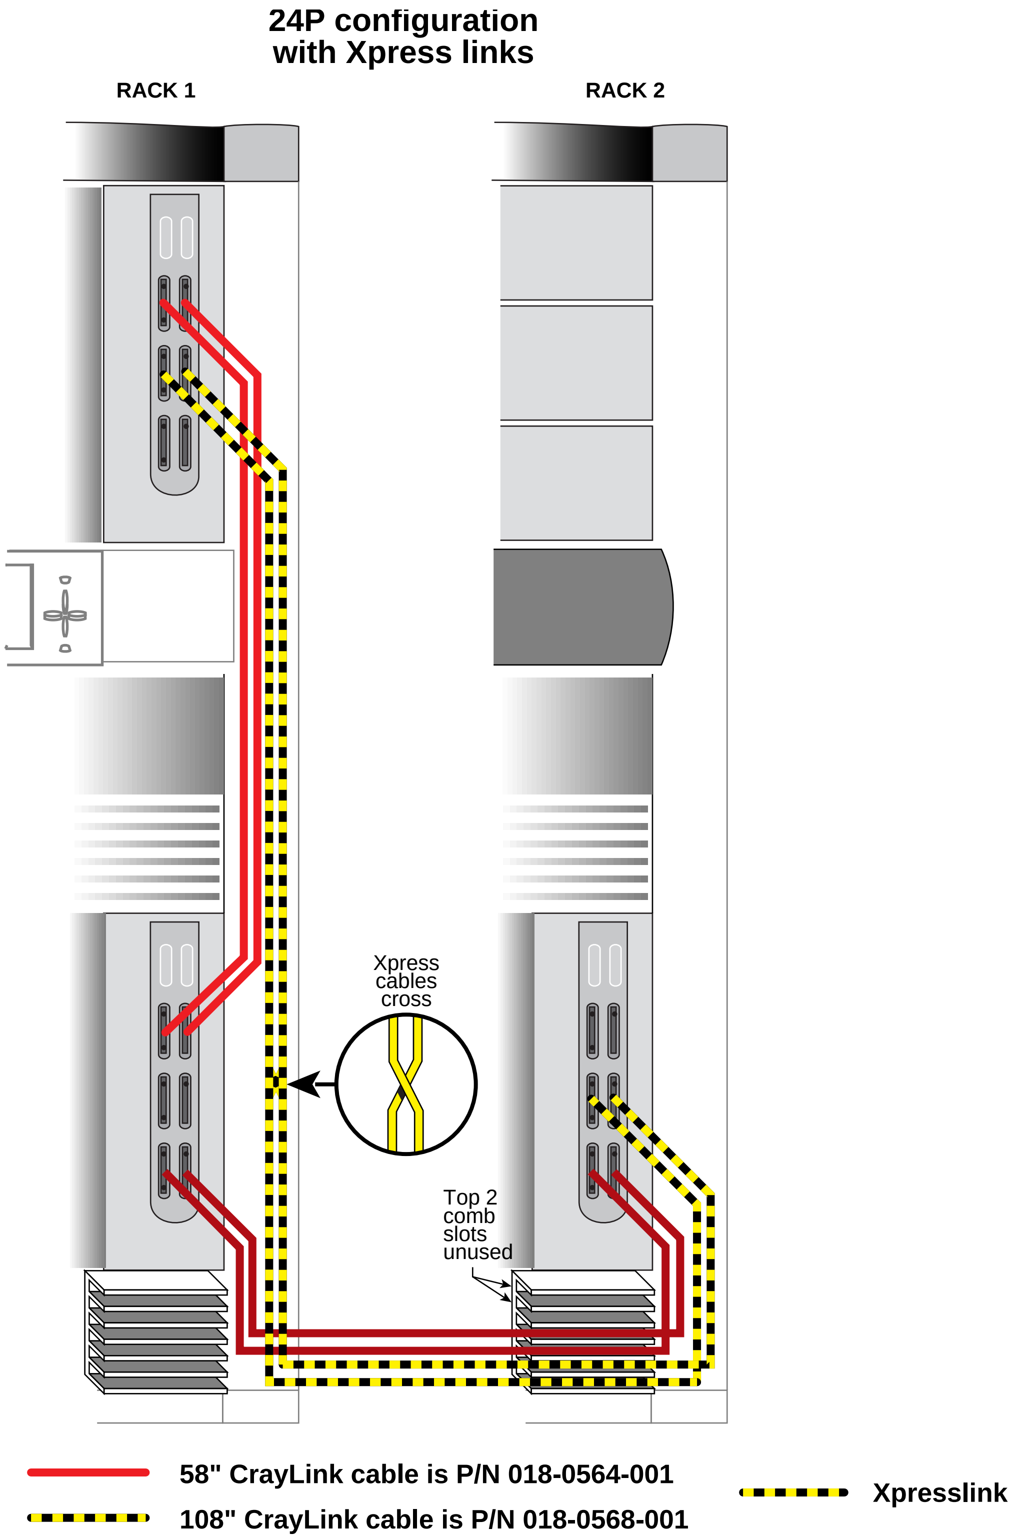

Figure 6-29 24P with Xpresslinks Cabling Routing

32P Cabling overview

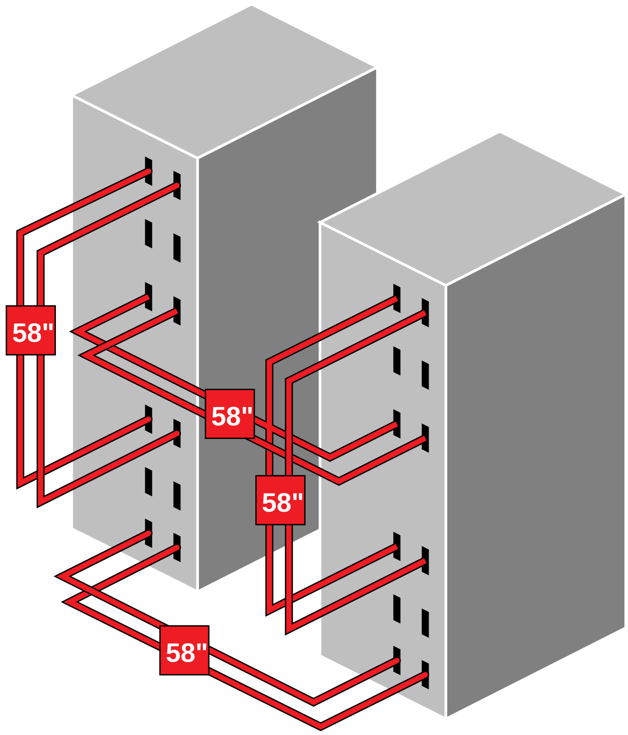

58" CrayLink cable is P/N 018-0564-001

Figure 6-30 32P Cabling Overview

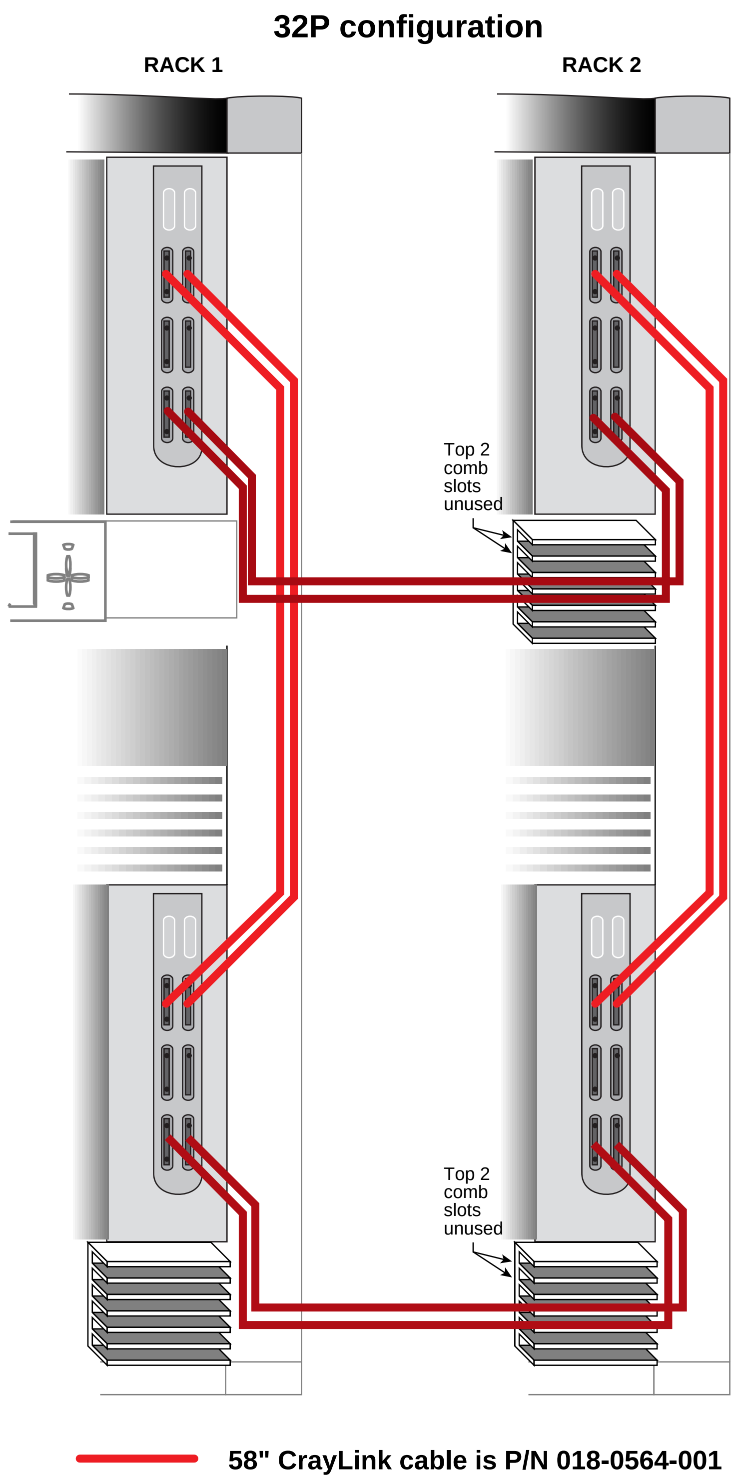

Figure 6-31 32P Cabling Routing

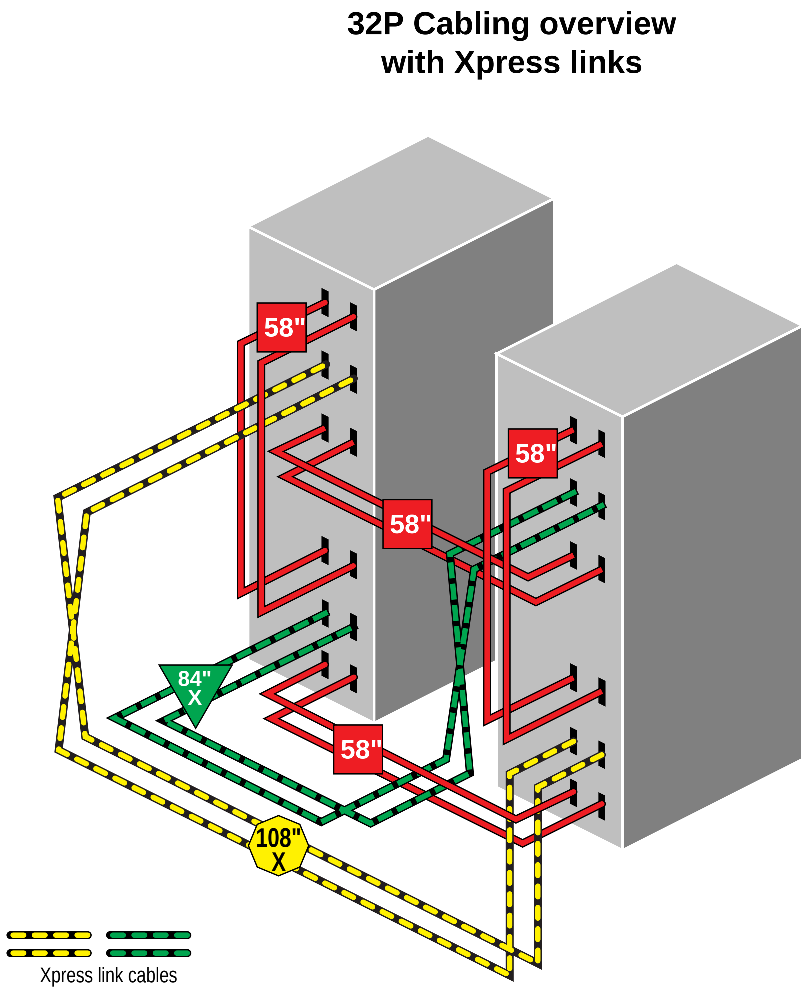

58" CrayLink cable is P/N 018-0564-001

84" CrayLink cable is P/N 018-0693-001

108" CrayLink cable is P/N 018-0568-001

Figure 6-32 32P with Xpresslinks Cabling Overview

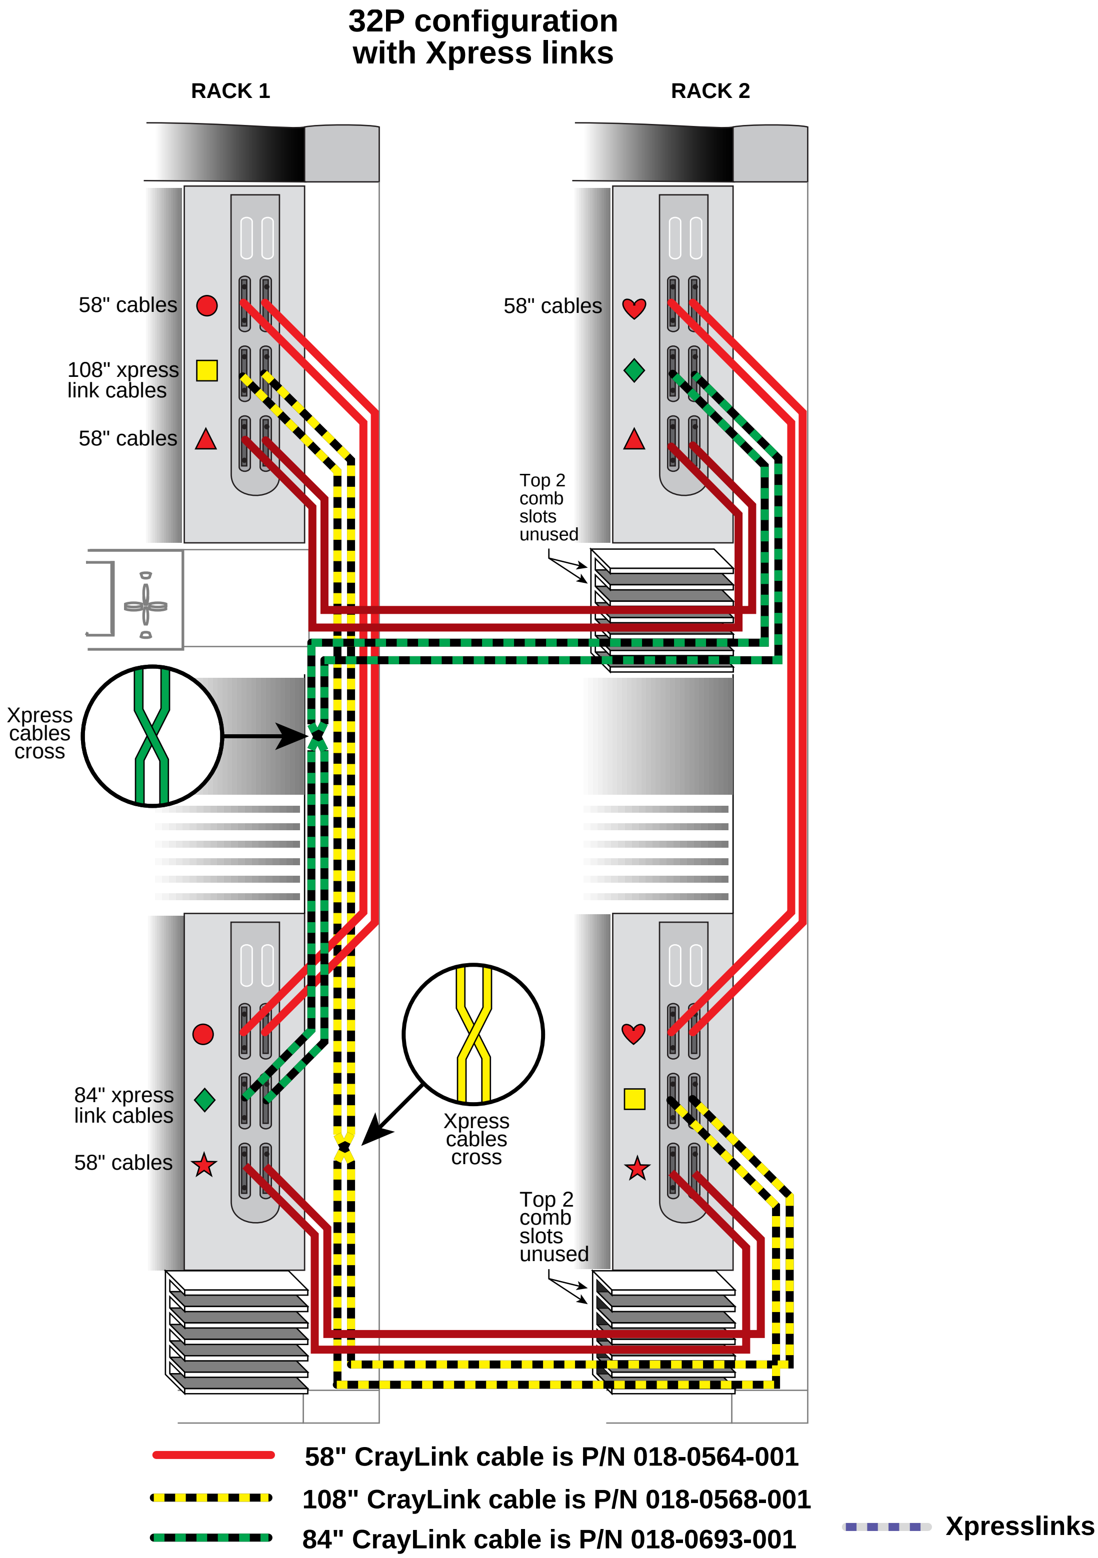

Figure 6-33 32 P with Xpresslinks Cabling Routing

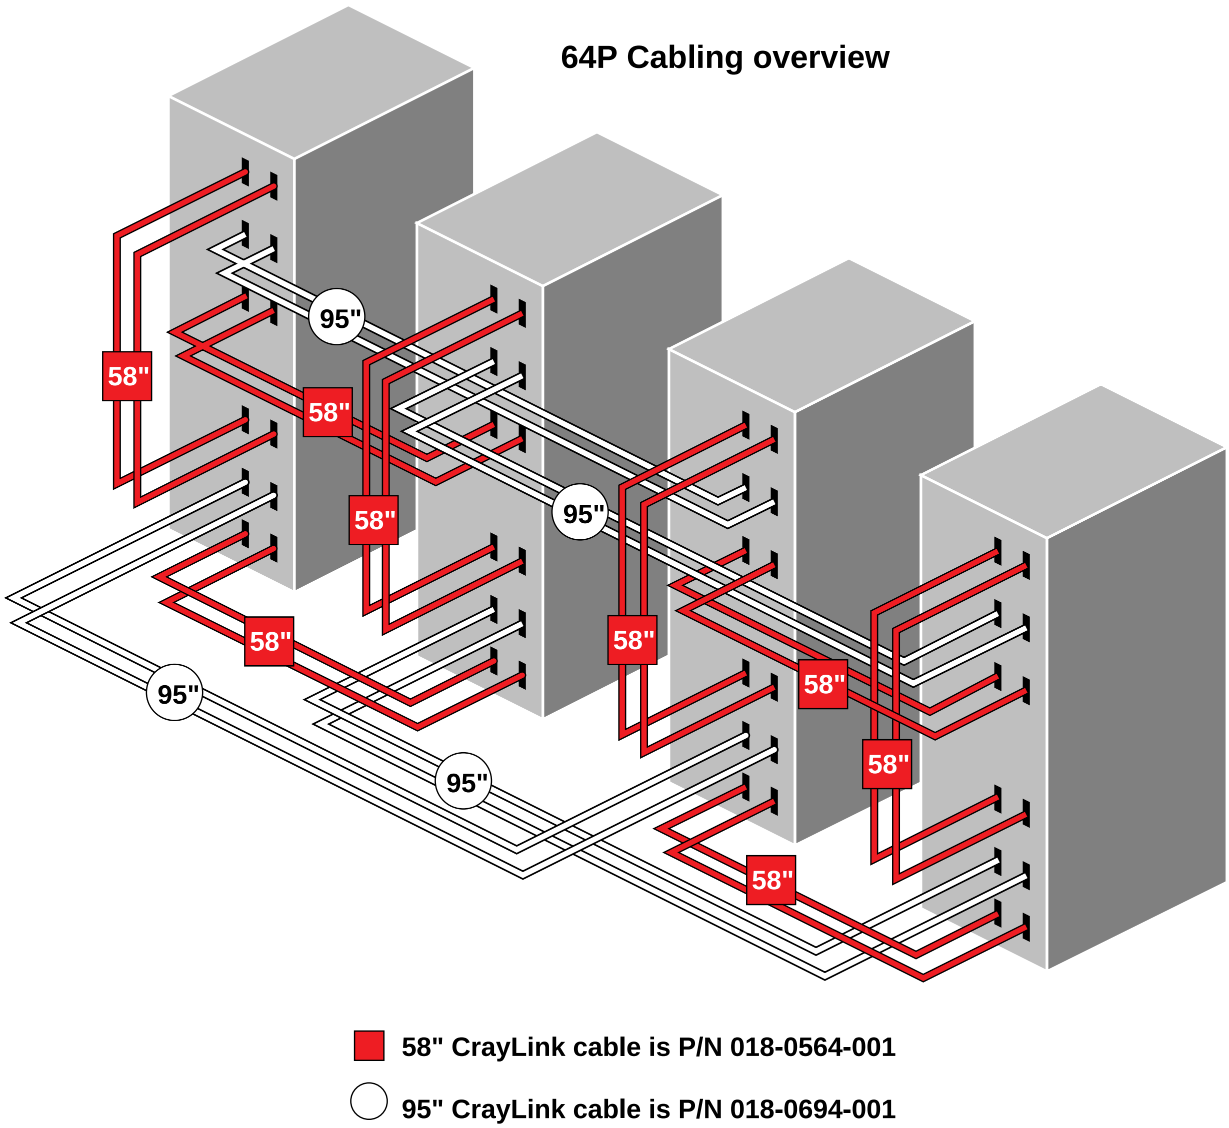

Figure 6-34 64P Cabling Overview

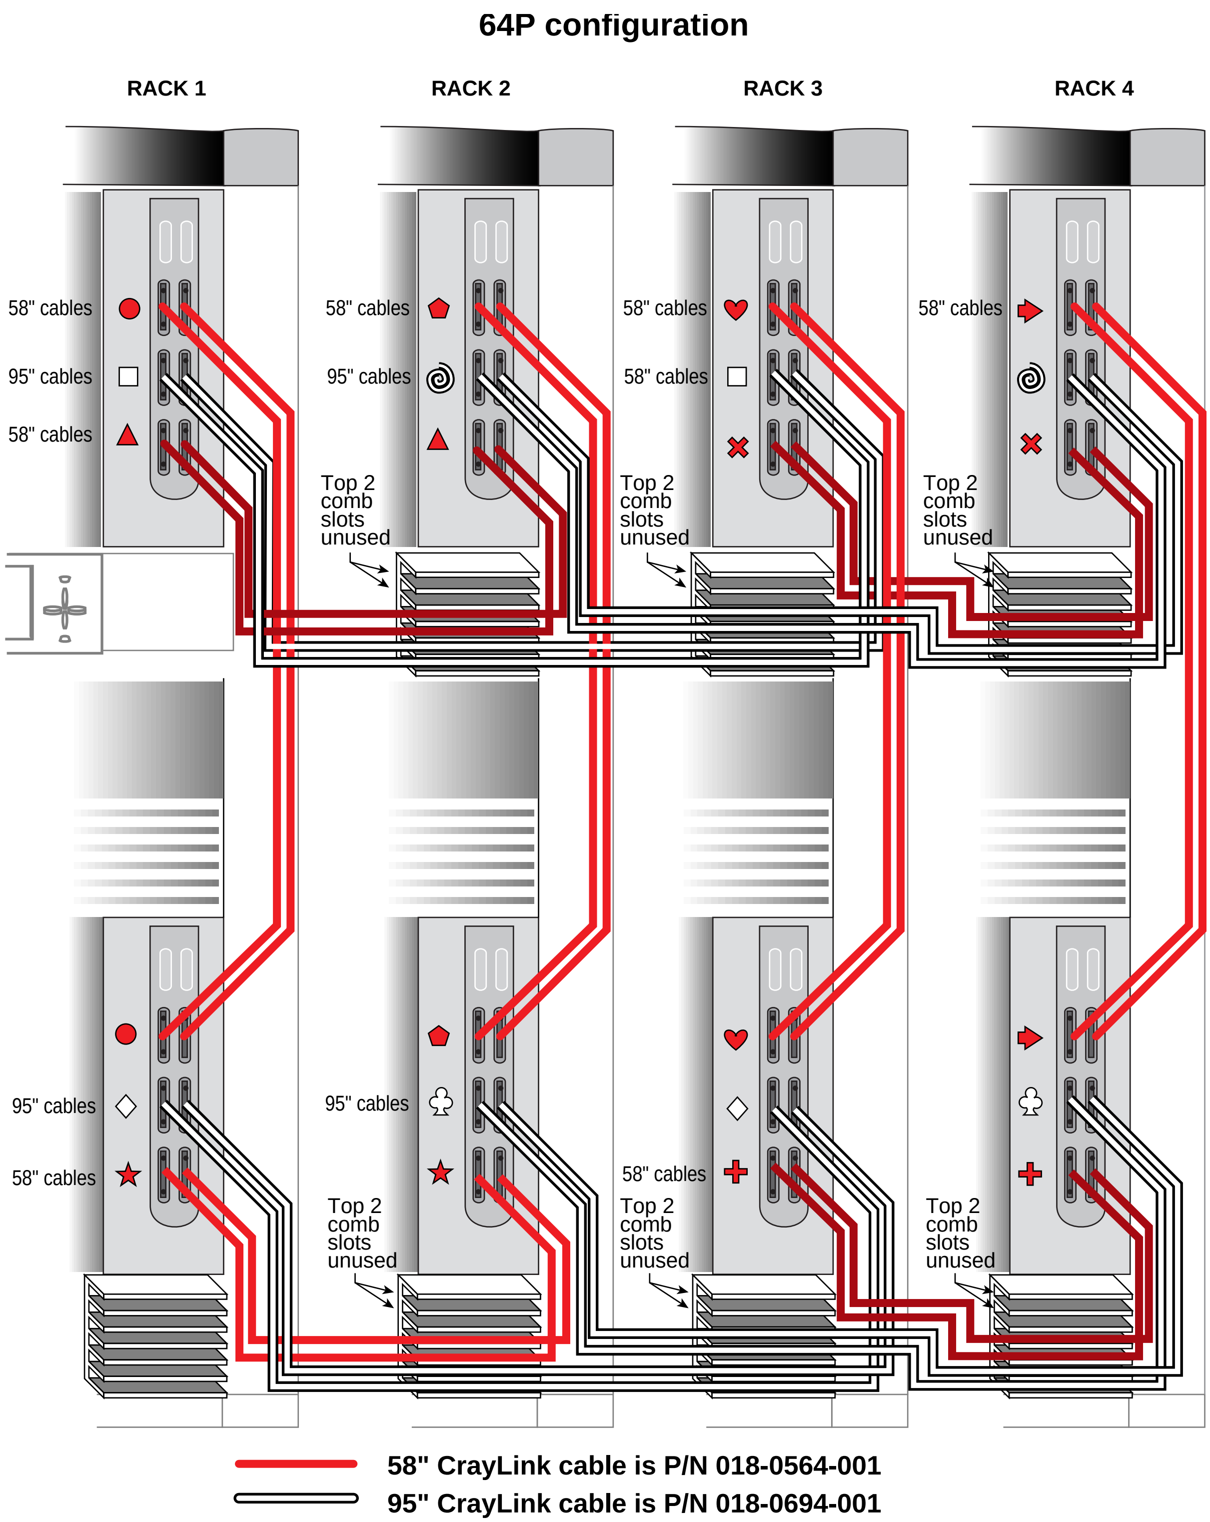

Figure 6-35 64P Cabling Routing

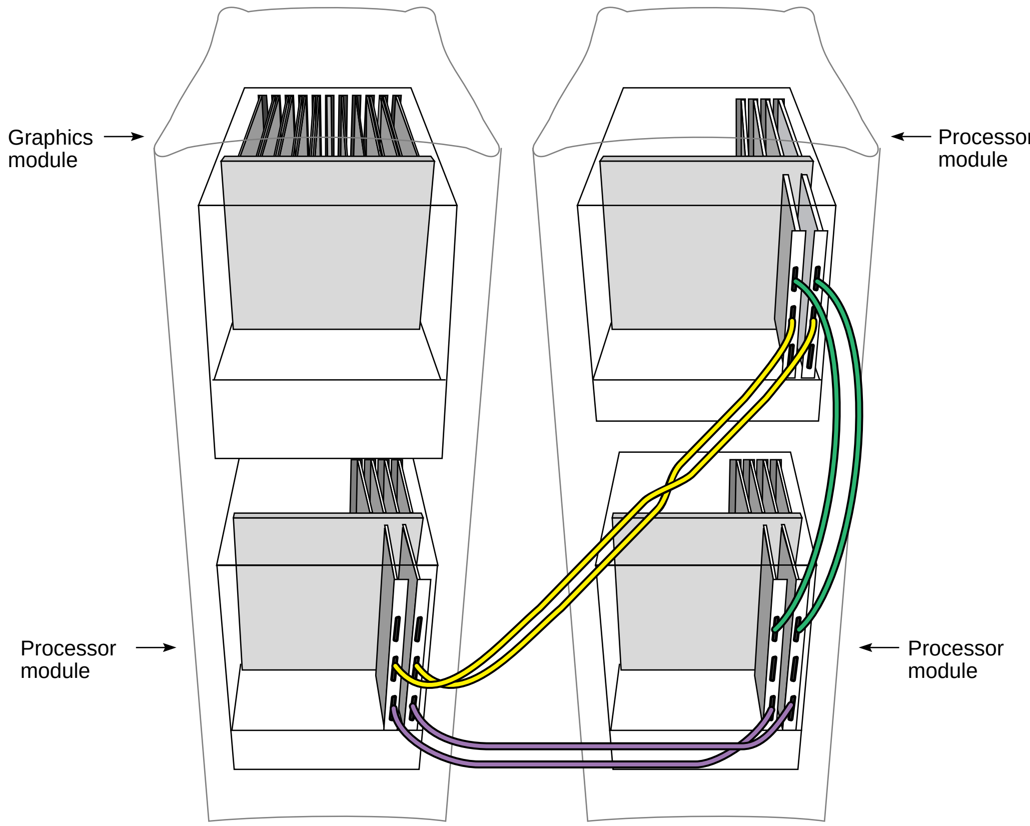

Figure 6-36 24P Onyx2 Configuration with Two Graphics Pipes

6.14.3.1 CrayLink Interconnect Cabling Managementt

The following illustrations provide an eight-step overview of the CrayLink cabling procedure. Follow these steps to install CrayLink Interconnect cables between modules.

Caution: Be sure to power off the modules before you install the CrayLink Interconnect cables.

1. Remove the cable door by lifting it up slightly and pulling it out. Set the door aside and out of the way.
2. Find the appropriate CrayLink Interconnect cable.

   Caution: Be sure the connectors are properly aligned with the router ports before you tighten the connector screws. You can destroy the cable if you do not install the connectors correctly.

3. Install a cable bale marked DWN (down) into the upper channel. Insert the tabs of the bale into the cutout and slide the bale down until it locks into place. There is a spring-loaded pin in the back of the bale that automatically engages when the bale is slid into position.
4. Insert a cable clip over the cable and cable bale.
5. Install an UP cable bale into the lower channel cutout.
6. Insert a cable clip over the cable and cable bale.
7. Starting with the top module, insert the excess cabling into the rubber cable slit on the right side of the chassis. Insert one cable at time.

   Caution: Make sure not to put a kink into the cables as you push them into the slit.

8. After you finish inserting one cable into the slit of the rubber cable trough, push the second cable into the rubber slit on top of the first cable. This helps the cables are from coming out of the slit.
9. Reinstall the cable door after you are done cabling.
10. Power on the rack as described in Section 6.8, "Powering On and Verifying Operation."
11. Check that the applicable green LEDs on the router boards are illuminated to determine if the cables have been properly installed (see "Interpreting Router Board LEDs" in Chapter 7 for additional information).
12. If they are not working, *power off* the system and recheck the CrayLink Interconnect cable connections.

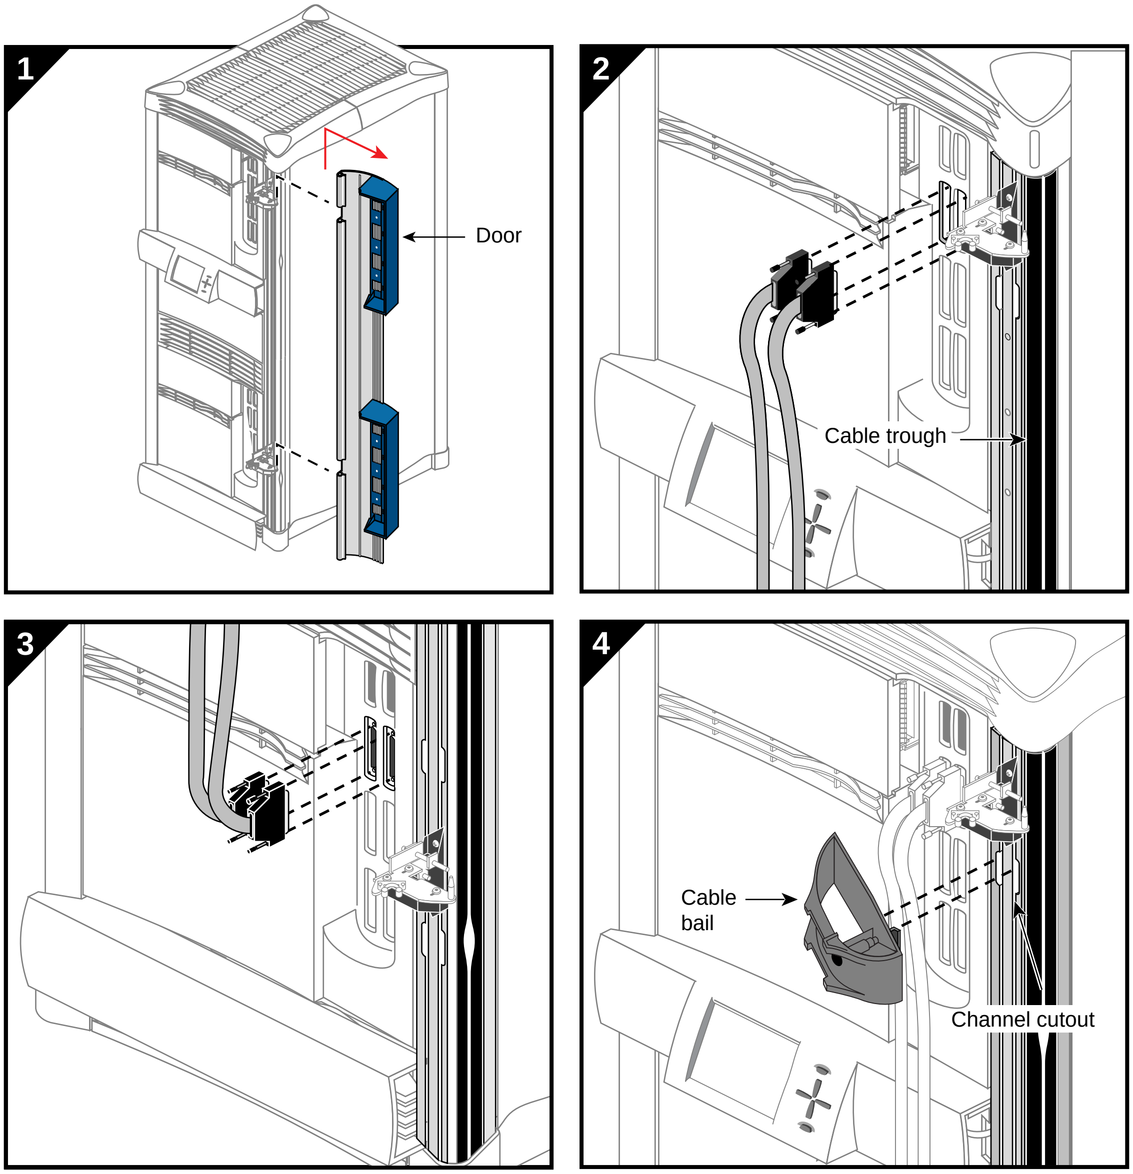

Figure 6-37 Installing Cables Between the Top and Lower Modules in a Rackmount System

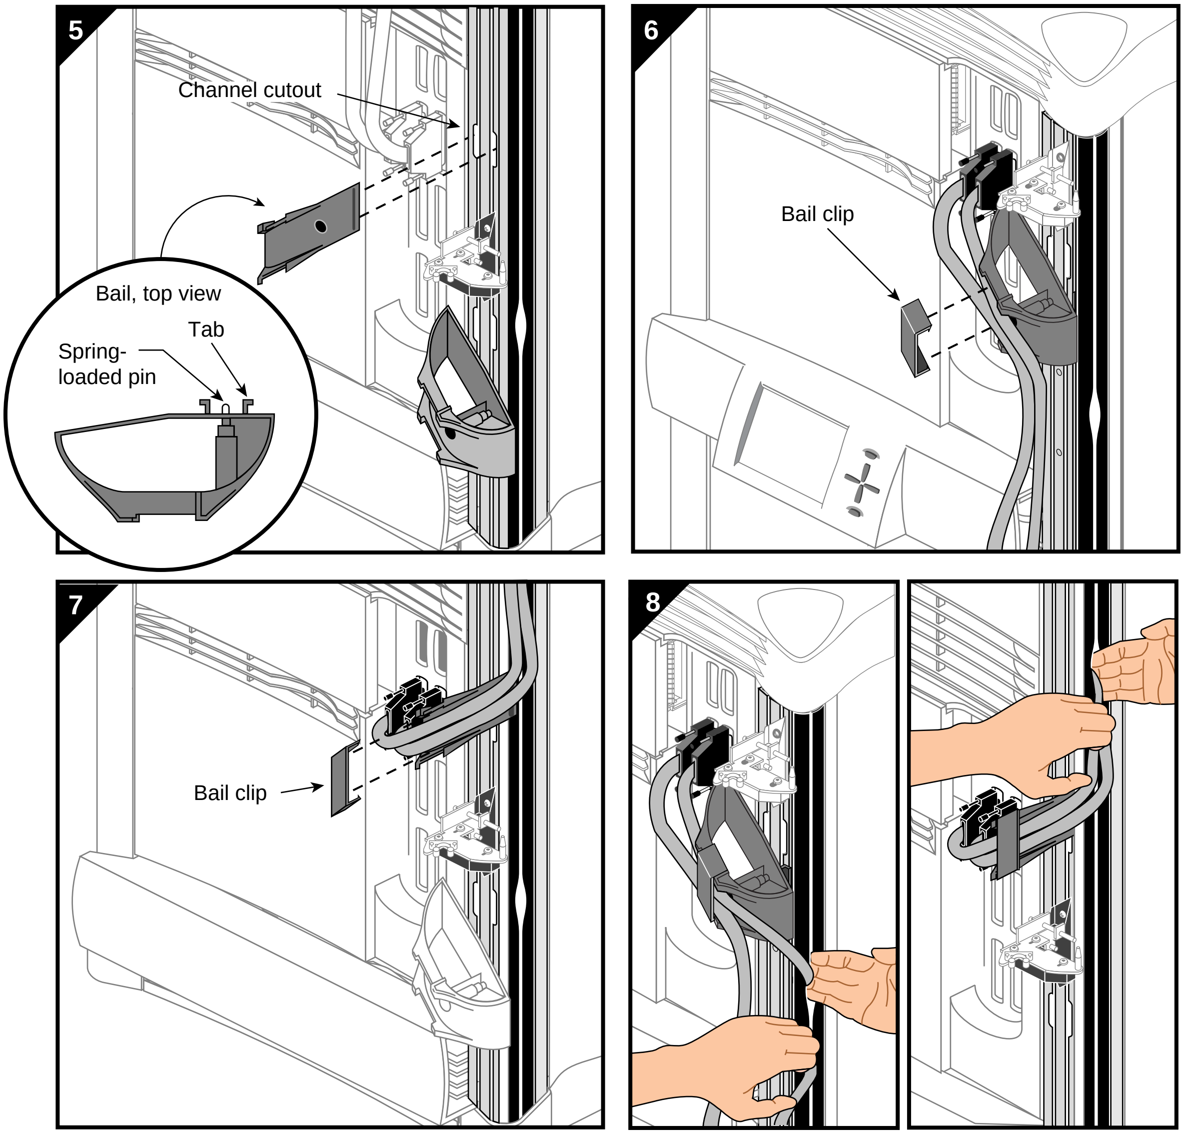

Figure 6-38 Installing the Cable Management Hardware

1. If not already done, remove each of the cable doors in front of the rack (see "Removing the Cable Cover Door" in Chapter 10 and Step 1 in Figure 6-37). Set these doors aside, out of the way.
2. Install an additional DWN (down) cable bale as shown in Figure 6-39 in each of the upper modules.
3. Check Figure 6-32 to Figure 6-34 to determine which size cable to use between modules. In most instances, it will be a 65" blue cable.
4. Pull out the upper and lower cable comb doors.
5. Important! Route the cables into the specific comb slots as shown in Figure 6-32 to Figure 6-39. Tuck the cables neatly in place into the comb.
6. Cover the upper cable combs with the comb cover after you have finished cabling.

7. Install an additional DWN (down) cable bale as shown in Figure 6-40.
8. Check Figure 6-32 to Figure 6-34 to determine which size cable to use between modules. In most instances, it will be a 65” (blue cable).
9. Important! Route the cables into the specific comb slots as shown in Figure 6-32 to Figure 6-40. Tuck the cables neatly in place into the comb.
10. Cover the upper cable combs with the comb cover after you have finished cabling.
11. Reinstall the cable doors to hide the CrayLink Interconnect cabling.

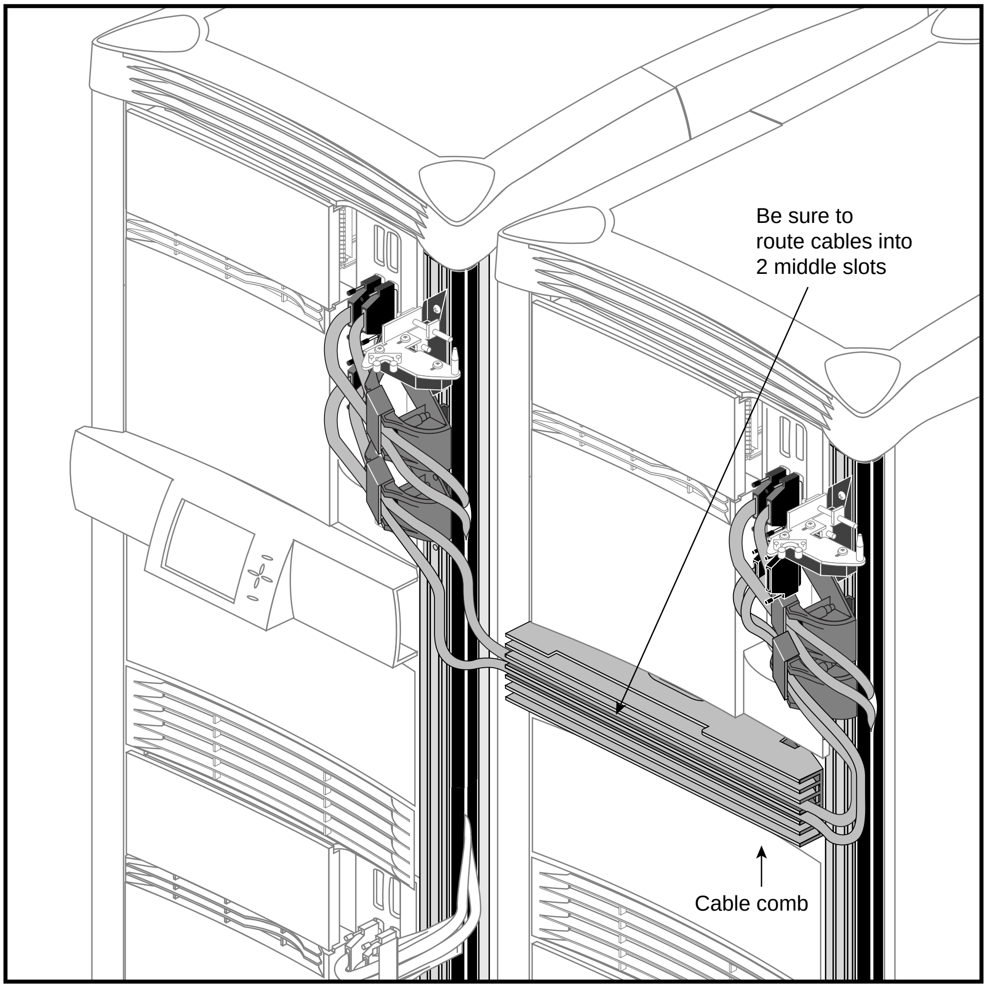

Figure 6-39 Installing CrayLink Interconnect Cabling Between Adjacent Upper Modules

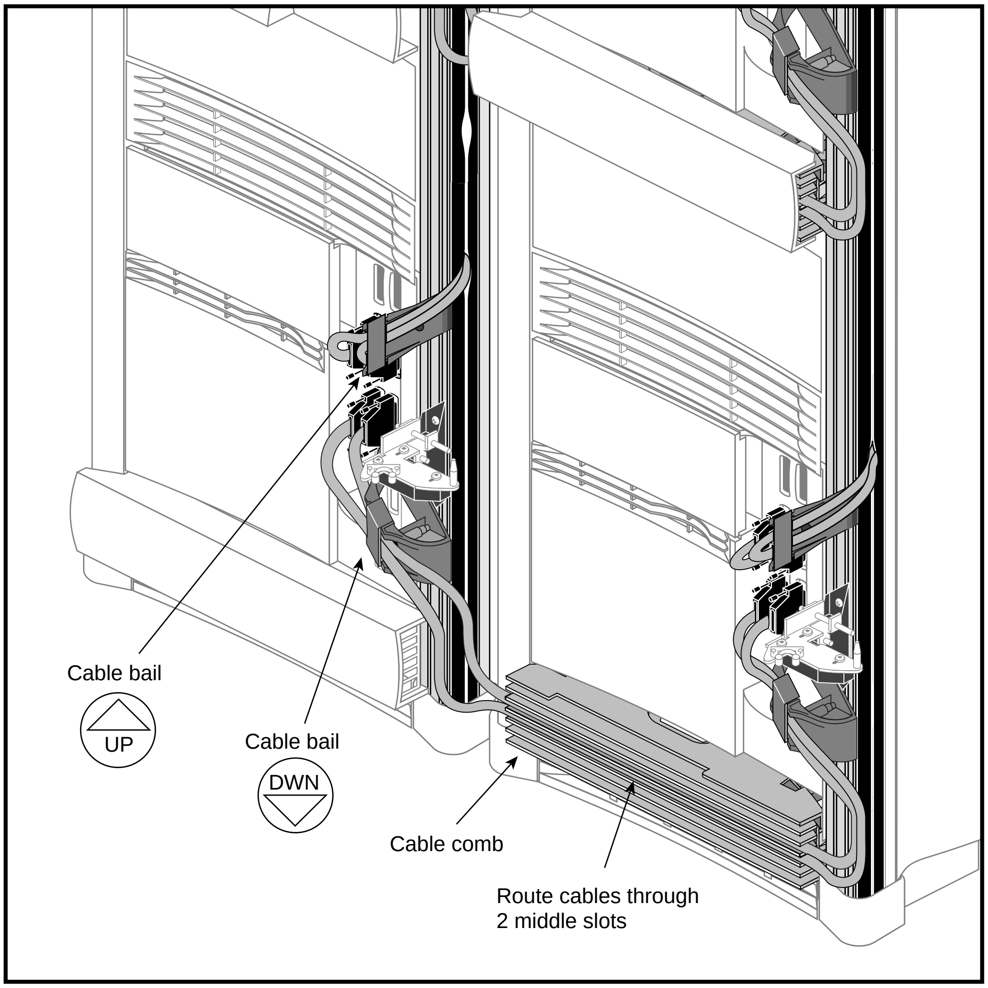

Figure 6-40 Installing CrayLink Cabling Across Adjacent Bottom Modules

6.15 Connecting the Crosstown Cable in Onyx2 Racks

If you have an Onyx2 rack system, you must connect the Ktown board in the graphics module to a Crosstown board in the XIO cardcage of the server module using a Crosstown cable (see Figure 6-41 and the following procedures).

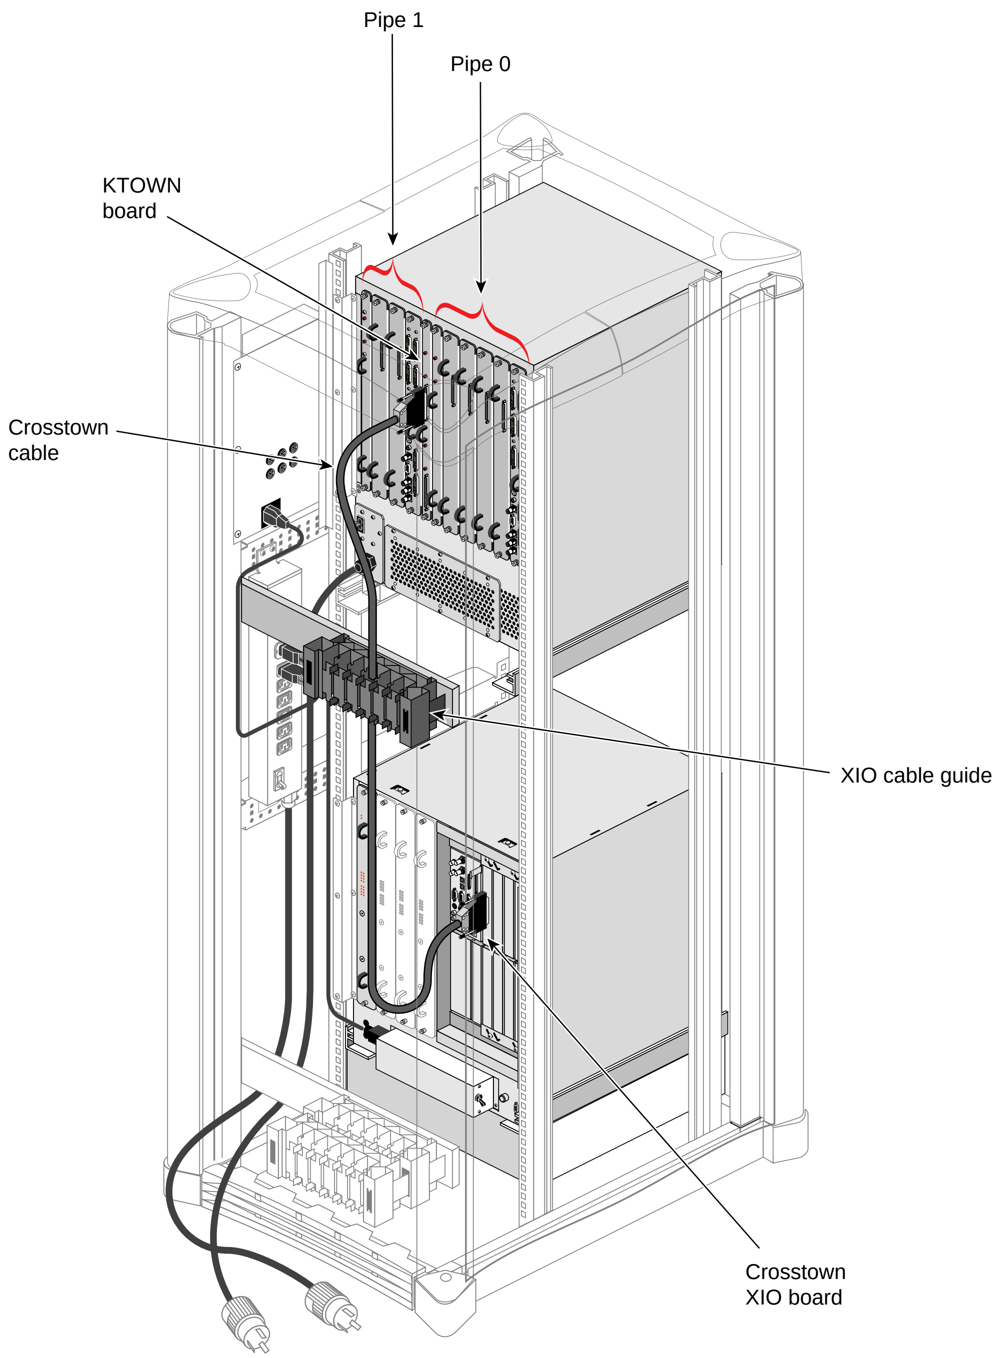

Figure 6-41 Connecting the Graphics Module to the Server Module in an Onxy2 Rack

Follow these procedures to connect the graphics module to the server module.

1. Select a crosstown cable. If you are not sure what the differences are between a crosstown cable and a CrayLink Interconnect cable, see “CrayLink Interconnect, Xpress Link, and Crosstown Cables” in Chapter 3.
2. If it is not already in place, install a Crosstown board into the XIO cardcage of the server module. See “Removing and Installing an XIO Board” in Chapter 10 for additional information.
3. Plug one end of the Crosstown cable into either port 0 or port1 of the Ktown board. Port 0 corresponds to graphics pipe 0 (the board set to the right of the Ktown board as you face the back of the system—see Figure 6-41. Port 1 corresponds to graphics pipe 1 (the board set to the left of the Ktown board).
4. Plug the other end of the crosstown cable into the Crosstown board located in the XIO cardcage.
5. Route the cable into the XIO cable guide management hardware. See “Removing and Installing an XIO Board” in Chapter 10 for directions on securing the cable as required.

6.16 Setting Up Multiple-Rack Configurations

To set up multiple racks, bring up the systems individually as outlined in Table 6-1 and Table 6-2 and make sure that each rack system is properly operating before attempting to cable the systems together.

6.16.1 Position the Racks

In a multirack configuration, you must always place the rack with the multimodule display in the leftmost position (see Figure 6-42). This helps ensure proper routing of the CrayLink cables. In addition, the racks should sit flush with one another with little or no spacing between them so that a grounding strap can be installed between the racks as described in Section 6.16.2. The flushed or *nested* racks also give the appearance of a single system configuration.

Follow these instructions to position the racks. Remember to observe the guidelines outlined in the *Site Preparation for Origin Family and Onyx2* manual (P / N 007-3452-xxx).

1. Remove the side panels in between the racks (see “Removing the Rackmount Side Panels” in Chapter 10). Only the outside racks should have a side panel (see Figure 6-42). The leftmost rack should have a left side panel (but not a right side panel), and the rightmost rack should have a right side panel (but not a left side panel).
2. Place the rack with the multimodule display in the leftmost position (see Figure 6-42).
3. Push the racks together so that they sit flush with one another.

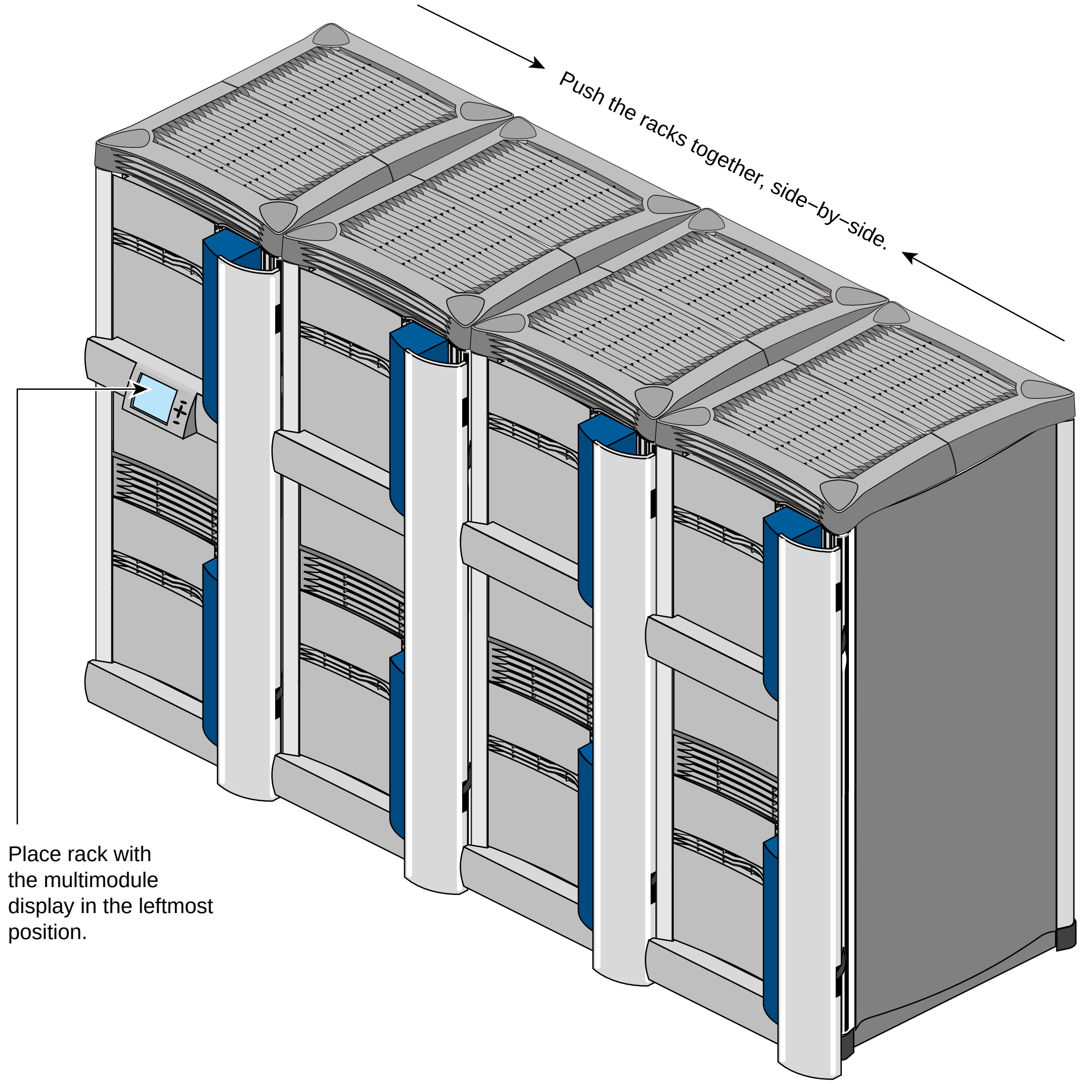

Figure 6-42 Positioning Racks in a Multirack Configuration (Four Racks)

6.16.2 Installing a Ground Strap Between the Racks

After the systems have been positioned, install ground straps between the racks as shown in Figure 6-43. This provides additional grounding to support CrayLink Interconnect cabling.

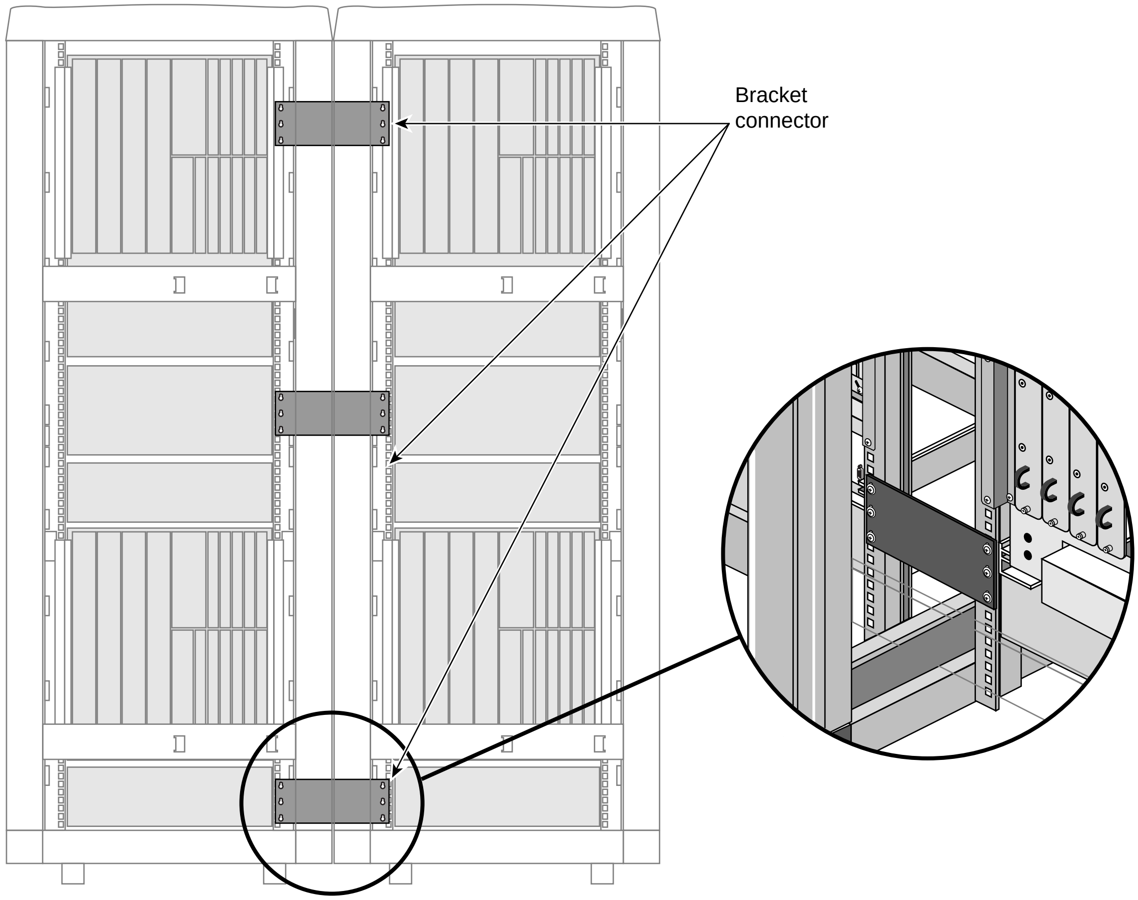

Figure 6-43 Installing Ground Straps Between Racks

6.16.3 Connecting the MMSCs

If you have two racks, connect the Ethernet ports on the MMSCs using an RJ-45 cable (see Figure 6-44). This enables the multimodule System Controllers to communicate with each other. If you have three or more racks, use an Ethernet hub to provide a private network between the racks. To connect three or more racks, follow these instructions.

1. On the rack housing the private network hub, connect the MMSC to the hub.
2. Connect the hub to the MMSCs in the remaining racks.

You do not need to install additional software or establish IP addresses or host files to set up the MMSC Ethernet connection. This private network provides a “plug-and-play” operation; no additional configuration is required.

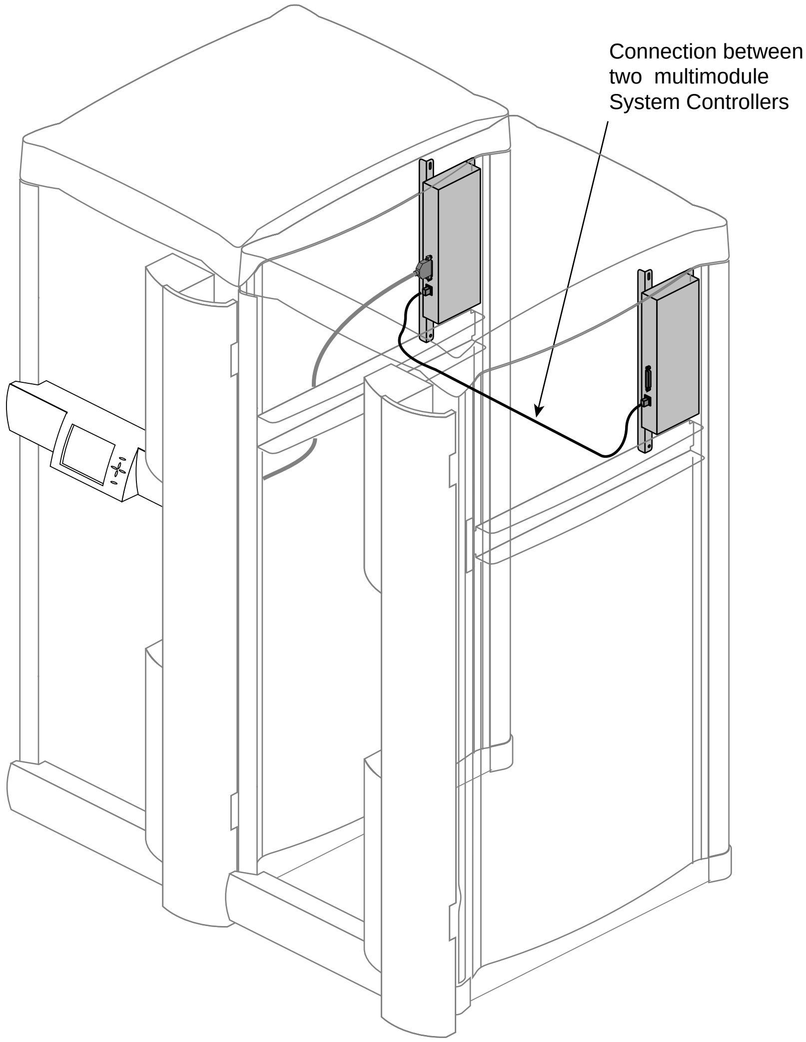

Figure 6-44 MMSC Connection to a Second MMSC (in a Two-Rack Configuration)

Caution: Power off the racks before you install the CrayLink Interconnect cables.

6.17 Connecting the System to Ethernet LAN

After you have set up and configured the systems, installed the software, and connected the CrayLink cables (if required), you can then install the Ethernet cabling. The Origin2000 and Onyx2 systems come with a RJ-45 (10 / 100-Base-T) Ethernet connector. Follow these instructions to connect an Ethernet drop to your system.

1. Locate the Ethernet line, then route it to the rear of the chassis.
2. Plug the cable into the applicable connector (see Figure 6-45 and Figure 6-46).

Note: Consult the *IRIX Admin: Networking and Mail* manual (P / N 007-2860-xxx) for additional information on setting up the network.

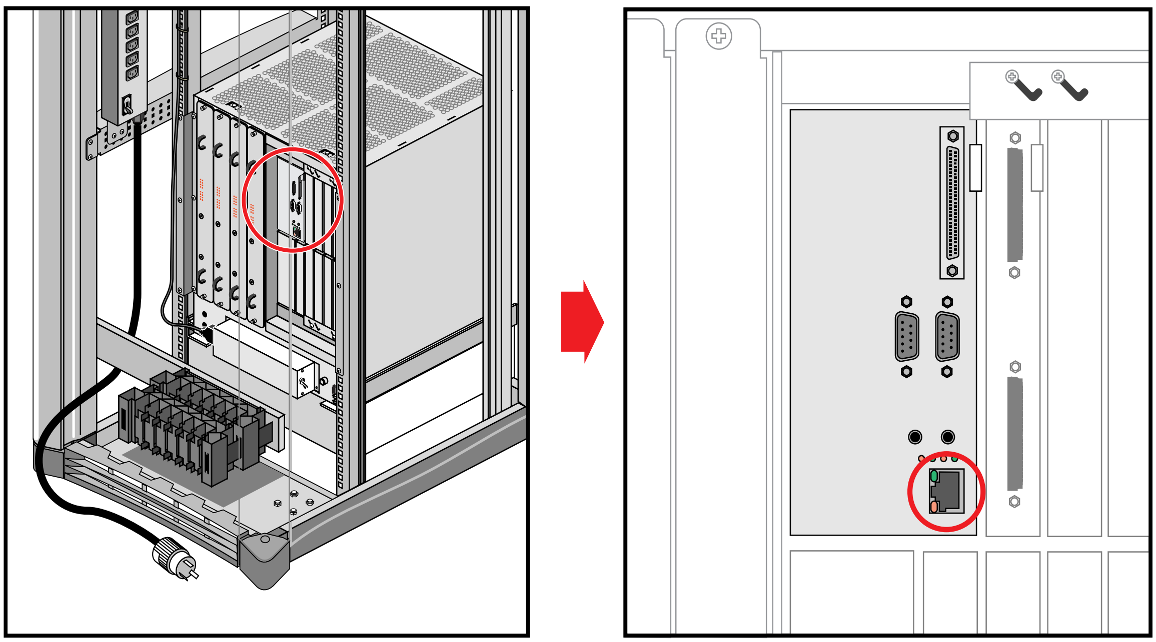

Figure 6-45 Connecting Ethernet to the Origin 2000 System (Rackmount Chassis)

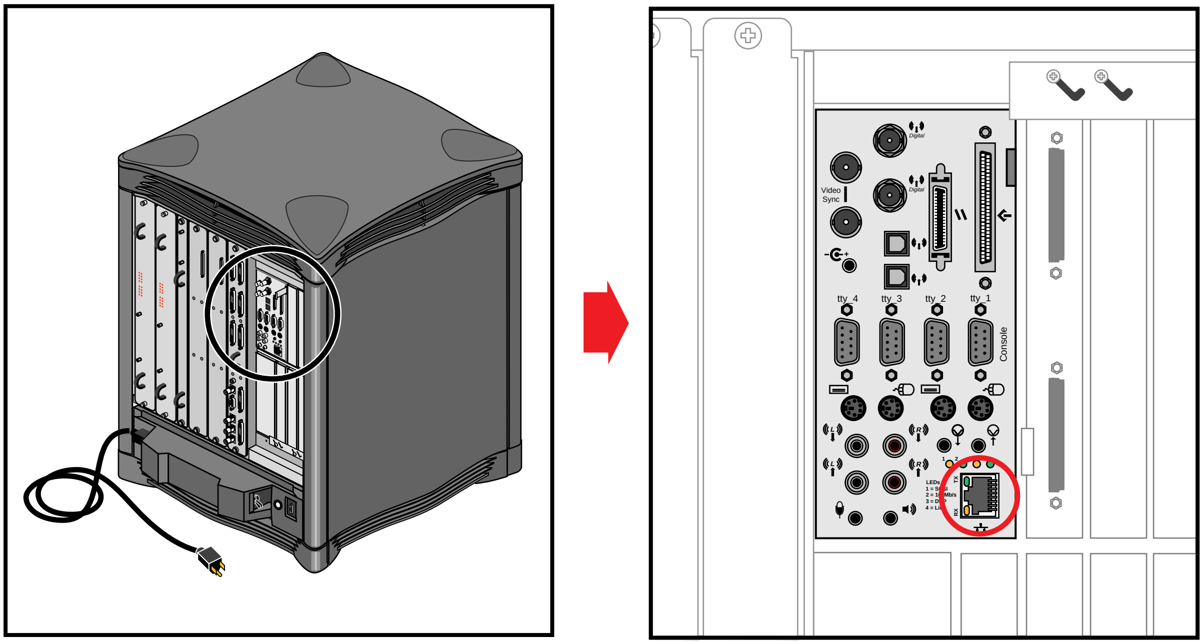

Figure 6-46 Location of Ethernet Connector on Onyx2 System (Deskside Chassis)

Chapter 7

Diagnostic Tools

The Origin2000 and Onyx2 provide a number of diagnostic, verification, and status reporting tools to aid in the troubleshooting and monitoring of system operation and resources. This comprehensive package includes the following categories of support:

- LED status reporting
- number-in-a-can (NIC) inventory report
- diagnostics
- field service software toolkit
- hardware graph inventory program
- flashing the PROMs
- updating the MMSC firmware

7.1 Interpreting the Node Board LEDs

Each Node board has two side-by-side columns of eight discrete LEDs. The left column provides a status value from CPU (slice) A, and the right column provides a status value from CPU B. The columns have the most significant LED bit on top (see Figure 7-1).

Interpreting the binary state of the LED bits is counterintuitive. If an LED is *on*, this indicates a *zero* bit; if an LED is *off*, this indicates a *one* bit. In Figure 7-1, CPU A shows hex value 0xf8, while CPU B displays 0x03.

During the boot process, the CPU changes the LEDs before each phase of initialization and during system operation, as needed. If a CPU were to hang during any phase, the residual LED value helps indicate which phase hung and pinpoint the failing component (for example, the R10000 data cache).

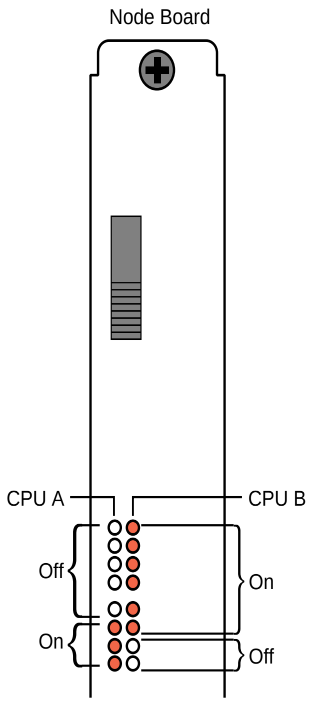

Figure 7-1 Node Board LED Example State

7.1.1 Boot Status LEDs

During system boot, the Node board LEDs are constantly updated with values indicating the boot progress. If the system were to crash during any phase, the LEDs would indicate what it was doing at the time (see Table 7-1). Hex values 0x00 to 0x7f are set aside for boot status information. In addition, diagnostics values are displayed on the Origin2000 and Onyx2 module System Controller display during boot. Further along into the boot process, a console becomes available to report more detailed information on failures.

Table 7-1 Boot Progress LED Values

LED	Name	Phase
0x00	RESET	
0x01	INITCPU	Initializing R10000 GPRS, FPRS, and COPO
0x02	TESTCP1	Testing R10000 COP1 registers
0x03	RUNTLB	Switch to mapped mode
0x04	TESTICACHE	Test R10000 primary instruction cache
0x05	TESTDCACHE	Test R10000 primary data cache
0x06	TESTSCACHE	Test secondary cache
0x07	FLUSCHCACHES	Flush all caches

Table 7-1 (continued) Boot Progress LED Values

LED	Name	Phase
0x08	CKHUBLOCAL	
0x09	CKHUBCONFIG	
0x0a	INVICACHE	Invalidate R10000 primary instruction cache
0x0b	INVDCACHE	Invalidate R10000 primary data cache
0x0c	INVSCACHE	Invalidate secondary cache
0x0d	INMAIN	Succeeded in jumping to main()
0x0e	SPEEDUP	About to increase PROM access speed
0x0f	SPEEDUPOK	Increased PROM access speed
0x10	INITDCACHE	
0x11	INITICACHE	
0x12	INITCOP0	
0x13	FLUSHTLB	
0x14	CLEARTAGS	
0x15	CCLFAILED_INITUART	
0x16	HUBINIT	
0x17	HUBCFAILED INTUART	
0x18	N0CLOCK_INITUART	
0x19	HUBINITDONE	
0x1a	MSCPROBE	About to probe for presence of MSC
0x1b	JUNKPROBE	About to probe for presence of junk UART
0x1c	DONEPROBE	Done probing for presence of MSC
0x1d	UARTINIT	About to initialize selected UART
0x1e	UARTINITDONE	Done initializing selected UART
0x1f	CKHUBCHIP	
0x20	PODMAIN	
0x21	PODLOOP	About to enter POD mode, C portion
0x22	POPDPROMPT	Just about to enter POD prompt loop
0x23	PODMODE	About to enter POD mode, assembler portion
0x24	LOCALARB	Performing local arbitration (CPU A / B)
0x25	SCINIT	
0x26	BMARB	

Table 7-1 (continued) Boot Progress LED Values

LED	Name	Phase
0x27	BMASTER	
0x28	BARRIER	About to perform first local barrier
0x29	CKPDCACHE1	
0x2a	MAKESTACK	About to configure Dex mode stack and data
0x2b	MAIN	Reached main()
0x2c	LOADPROM	
0x2d	CKSCHACHE1	
0x2e	CKBT	
0x2f	INSLAVE	
0x30	PROMJUMP	
0x31	NMI	
0x32	INV_IDCACHES	
0x33	INV_SCACHE	
0x34	WRCONFIG	
0x35	RTCINIT	About to initialize Hub real-time counter
0x36	RTCINITDONE	Done initializing Hub real-time counter
0x37	LOCK	
0x38	BARRIEROK	First local barrier succeeded
0x39	LOCKOK	
0x3a	FPROMINIT	
0x3b	FPROMINITDONE	
0x3c	JUMPRAMU	About to jump to UALIAS space
0x3d	JUMPRAMUOK	Jumped to UALIAS space
0x3e	JUMPRAMC	About to jump to cached space
0x3f	JUMPRAMCOK	Jumped to cached space
0x40	STACKRAM	About to test stack area of memory
0x41	STACKRAM	Done testing stack area of memory
0x42	SLAVEINT	
0x43	SLAVECALL	
0x44	SLAVEREND	
0x45	LAUNCHLOOP	About to enter slave launch loop

Table 7-1 (continued) Boot Progress LED Values

LED	Name	Phase
0x46	LAUNCHINTR	Received launch interrupt
0x47	LAUNCHCALL	Calling launched function
0x48	LAUNCHDONE	Launched function returned
0x49	UARTBASE	
0x4a	MDIRINIT	About to initialize Hub MD and SIMM controls
0x4b	MDIRCONFIG	About to probe and configure memory size
0x4c	12CINIT	About to initialize PCF858412C chip
0x4d	12CDONE	Done initializing PCF858412C chip
0x4e	CONFIG_INIT	
0x4f	IODISCOVER	About to discover Hub I/O
0x50	HUB_CONFIG	
0x51	ROUTER_CONFIG	About to write Router cfg info into KLCONFIG
0x52	INITII	About to initialize I/O section of Hub
0x53	CONSOLE_GET	About to probe I/O section for console
0x54	INITIIDONE	Done initializing I/O section of Hub
0x55	STASH1	About to save Hub error registers
0x56	STASH2	About to clear Hub error registers
0x57	STASH3	About to enable error checking
0x58	STASH4	Done enabling error checking
0x59	IODISCOVER_DONE	Done discovering Hub I/O
0x5a	NMI_INIT	About to initialize NMI handler area
0x5b	TEST_INTS	About to test Hub interrupts

7.1.2 Failure Modes

In addition to indicating boot progress, the LEDs indicate fatal hardware problems found during diagnostics the PROM performs in each boot phase. If a fatal problem is found, the CPU sets the LEDs to a failure value between 0x80 and 0xa1 (as shown in Table 7-2) and automatically disables itself.

Table 7-2 Failure LED Values (for Node Boards)

LED	Name	Reason
0x81	CP1	R10000 COP1 register test failed

Table 7-2 (continued) Failure LED Values (for Node Boards)

LED	Name	Reason
0x82	RESTART	Restart master unable to load BaseIO PROM
0x83	ICACHE	R10000 primary instruction cache test failed
0x84	DCACHE	R10000 primary instruction cache test failed
0x85	SCACHE	Secondary cache test failed
0x86	UNUSED1	
0x87	BADGDA	
0x88	ECC	
0x89	XTL BMISS	
0x8a	UTL BMISS	
0x8b	KTL BMISS	
0x8c	GENERAL	
0x8d	NOTIMPL	
0x8e	CACHE	
0x8f	OS	
0x90	HUBINTS	
0x91	HUBLOCAL	
0x92	HUBCONFIG	
0x93	HUBCLOCK	
0x94	HUBCHIP	
0x95	HUBUART	
0x96	HUBCCS	
0x97	MAINRET	Main() returned
0x98	NOMEM	Node board has no local memory
0x99	IC2FATAL	
0x9a	DISABLED	CPU is disabled by environment variable
0x9b	NONMI	
0x9c	COREDEBUG	Can't set CORE debug registers
0x9d	IODISCOVER	
0x9e	HUB_CONFIG	Failed writing Hub info into KLCONFIG
0x9f	ROUTER_CONFIG	Failed writing Router info into KLCONFIG
0xa0	HUBII_INIT	Hub I/O initialization failed

Table 7-2 (continued) Failure LED Values (for Node Boards)

LED	Name	Reason
0xa1	CONFIG_INIT	Failed initializing KLCONFIG

7.1.3 Exception Error Reporting

The codes in Table 7-3 indicate that an exception occurred so early in the PROM boot process that no TTY device was available to display pertinent information.

Table 7-3 Early Exception Failures LED Values

LED	Name	Reason
0x81	CP1	R10000 COP1 register test failed
0x82	RESTART	Restart master unable to load BaseIO PROM
0x83	ICACHE	R10000 primary instruction cache test failed
0x84	DCACHE	R10000 primary instruction cache test failed
0x85	SCACHE	Secondary cache test failed
0x86	UNUSED1	
0x87	BADGDA	
0x88	ECC	
0x89	XTL BMISS	
0x8a	UTL BMISS	
0x8b	KTL BMISS	
0x8c	GENERAL	
0x8d	NOTIMPL	
0x8e	CACHE	
0x8f	OS	
0x90	HUBINTS	
0x91	HUBLOCAL	
0x92	HUBCONFIG	
0x93	HUBCLOCK	
0x94	HUBCHIP	
0x95	HUBUART	
0x96	HUBCCS	
0x97	MAINRET	main() returned

Table 7-3 (continued) Early Exception Failures LED Values

LED	Name	Reason
0x98	NOMEM	Node board has no local memory
0x99	IC2FATAL	
0x9a	DISABLED	CPU is disabled by environment variable
0x9b	NONMI	
0x9c	COREDEBUG	Can't set CORE debug registers
0x9d	IODISCOVER	
0x9e	HUB_CONFIG	Failed writing Hub info into KLCONFIG
0x9f	ROUTER_CONFIG	Failed writing Router info into KLCONFIG
0xa0	HUBII_INIT	Hub I/O initialization failed
0xa1	CONFIG_INIT	Failed initializing KLCONFIG

7.2 Interpreting Router Board LEDs

The Rack Router board has two side-by-side columns of six discrete LEDs and three power LEDs on the top of the board (see Figure 7-2). The left-side (green) LEDs indicate the connection status of router ports one through six. They light up if a port has successfully maintained a connection with another device. The left-side LEDs turn off if the port is disconnected (physically or due to severe connection errors). The right-side (yellow) LEDs are controlled by system software and typically light up when there is traffic across the corresponding link. Even though ports four through six are on the Origin2000 and Onyx2 rackmount module midplane, the connection lights are still useful because they can indicate if a Node board is improperly seated.

Note: Only the Rack Router boards have status LEDs on the front bezel. The Null and Star Router boards do not have any LEDs.

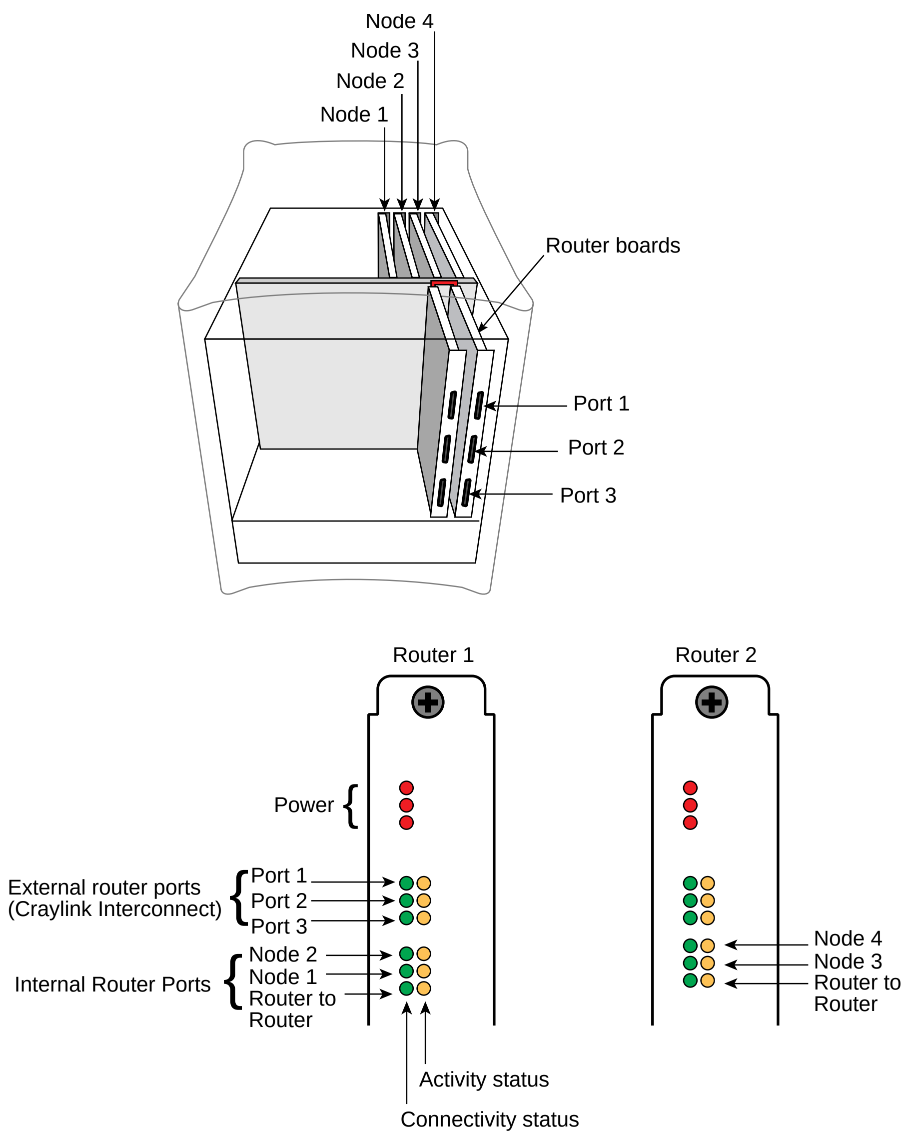

Figure 7-2 Router Board LEDs

Figure 7-3 provides an example of a deskside configuration with three Node boards installed. In this example, the LED for Router board port three is lit, along with the LEDs corresponding to Node boards 1, 2, and 3. Notice that the LED for node 4 is off because that board is not present.

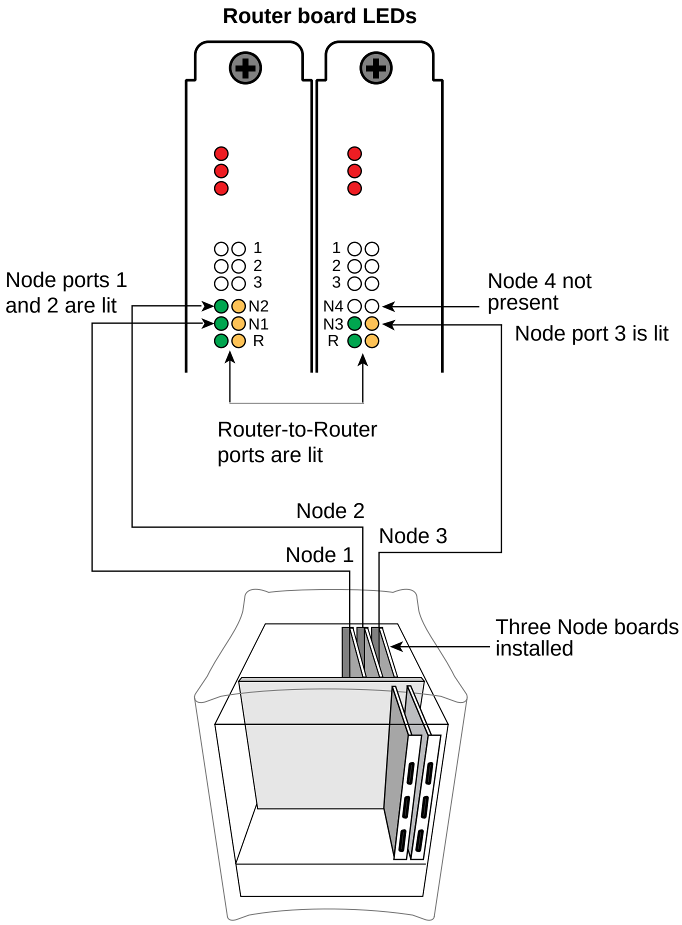

Figure 7-3 Router Board LED State Example in a Deskside System

7.3 NIC Reporting

The number-in-a-can (NIC) chip is used extensively throughout the Origin2000 and Onyx2 systems. There is a NIC on each type of board in the system, including the Node card, Router card, BaseIO card (with two NICs), and midplane.

The NIC contains a 48-bit number that is permanently laser-burned during manufacturing. The number is guaranteed to be unique from all other NICs. The NIC on the Node card serves as a serial number useful for software licensing. The NIC on the midplane helps the software recognize when different modules are added or removed from the system, allowing it to reconfigure appropriately.

In addition to the 48-bit number, NICs also contain several pages of nonvolatile memory that can be read or written by the system. This memory is used to store additional information, such as

- board type

- board revision number
- board serial number (as it appears on the bar code)
- board-specific information (e.g., on the BaseIO card, the Ethernet MAC address)

In cases where there are two valid ConsolePaths in a system, the system may determine which ConsolePath to use by comparing the removable midplane NICs on the applicable modules. Refer to "Correcting ConsolePath and Console-moving Problems" in Chapter 10 for more information.

Note: At boot time, the "nicprobe" test reads the NIC contents and does a few sanity tests to make sure the programmed NIC contents are actually correct.

7.3.1 NIC Types Used in the Origin2000 and Onyx2 Systems

Two types of NICs are implemented in the Origin2000 and Onyx2 systems. One is a tiny chip that is soldered directly into the board. The other is a removable, coin-shaped device that can be reinstalled on replacement boards. The midplane and BaseIO boards have both types of NICs. Whenever these boards are replaced, make sure that you remove the coin-shaped NICs from the old boards and reinstall them into the replacement boards (see Chapter 10, "General Replacement Procedures," for additional information).

7.3.2 Reading the NIC

The NICs can be read using the PROM monitor with the *hinv -v* (verbose) or *hinv -mv* (manufacturing verbose) command. Under IRIX, use the *hinv -m* command to obtain NIC information.

Here is an brief sample display of an *hinv -v* or *hinv -m* command:

```
8P/12MIDPLANE board: barcode CDZ048   part 013-1547-002 rev B
IP27 Board: barcode CDz180  part 030-0733-003 rev  B
```

7.3.3 Checking GE14 Board Information

To obtain serial number, assembly number, and revision level information on the GE14 graphics board, enter the following:

```
/usr/gfx/gfxinfo -vv
```

You see a readout similar to:

```
GE:NIC #: 0000.0002.a400 (family: 0b)
   Serial #: CGP091
   Part #: 030-0681-003
   rev_code: A
   grp_code: 0x0
   capability: “\ff\ff\ff\ff”
   variety: 0xff
   name: GE14-4
```

7.4 Diagnostics for Origin2000 and Onyx2 Systems

There are three types of diagnostics available for the systems:

- power-on diagnostics
- micro-diagnostic kernel (MDK) diagnostics
- UNIX®-based diagnostics

7.4.1 Power-On Diagnostics

The power-on diagnostics (POD) are programmable read-only memory (PROM) resident routines that run automatically during the boot sequence. You can also run several of the power-on diagnostics manually from the command line of the PROM monitor.

The power-on diagnostics focus on the hardware needed to boot the operating system or the micro-diagnostic kernel. The routines that the power-on diagnostics use are stored in the IP27 PROM and the BaseIO PROM.

The POD diagnostics run in four modes: no diags mode, normal mode, heavy mode, and manufacturing mode. Table 7-4 describes the modes and the DIP switch settings used to set each mode.

Note: The DIP switches are located behind the MSC front bezel. For information on accessing these switches, see "Module System Controller (MSC)" in Chapter 3.

Table 7-4 POD Modes

Mode	Description
Normal 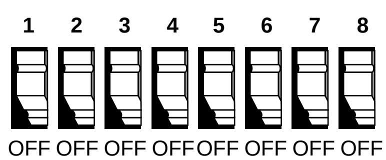	This mode performs testing to ensure that the operating system or MDK can run. Because of the time constraints necessary in the normal boot process, this mode performs a subset of the testing done by heavy mode. To select this mode, place all DIP switches in the off position.
No Diags 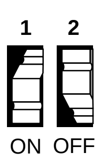	This mode does not perform any testing; only the initialization functions of the power-on diagnostics (for example, clearing memory) are performed. To set this mode, put DIP switch 1 in the on position and DIP switch 2 in the Off position. DIP switches 3 through 8 are off.

Table 7-4 POD Modes

Mode	Description
Heavy 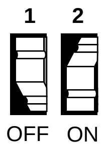	This mode performs more extensive and time-consuming testing than normal mode. Use this mode the first time that you power on a system in the field or when the normal mode detects errors. To select this mode, put DIP switch 1 in the off position and DIP switch 2 in the On position. DIP switches 3 through 8 are off.
Manufacturing 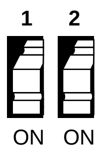	This mode performs the most extensive and time-consuming testing. Normally, this mode should be used only in the manufacturing environment. To select this mode, put both DIP switches in the on position. DIP switches 3 through 8 are off.

Power-on diagnostics return status information to the Node board LEDs, the module System Controller (MSC), a terminal connected to the MSC, and the system console.

7.4.2 Microdiagnostic Kernel Diagnostics

The microdiagnostic kernel (MDK) software provides an environment in which you can run diagnostic tests when the operating system is not running. Because MDK has complete control of the system, these diagnostics have more access to the hardware than the UNIX-based diagnostics.

Run MDK diagnostics from the PROM monitor by using the *bootp* command. These diagnostics return status information to the terminal from which you ran them.

7.4.3 UNIX-Based Diagnostics

The UNIX-based diagnostics run in the Origin2000 and the Onyx2 while the operating system is running in the mainframe. This enables you to test certain areas of the system while customer operations continue running.

Run UNIX-based diagnostics by entering the diagnostic name at the UNIX prompt. These diagnostics return status information to the terminal from which you ran them.

7.5 Field Service Software Toolkit

For the Origin2000 and Onyx2 systems, the Silicon Graphics field service toolkit contains the following support tools:

- System Verification Program (SVP)

The System Verification Program is an IRIX™-based software tool, combining utilities and test programs, that inventories the system hardware and software components, loads and runs a suite of test programs, and generates a set of files that report the results of this activity.

SVP helps verify the configuration and functionality of a newly installed system by a field system support engineer (SSE). SVP is a system acceptance test tool, and the test result is part of the system installation sign-off requirement for the SSE.

- Field Replaceable Unit (FRU) Analyzer

  The FRU analyzer is an IRIX-based software tool that provides the SSE with a starting point in resolving a system malfunction. When the operating system determines that a fatal error has occurred, the hardware error state is collected and passed to the FRU analyzer for interpretation. The FRU analyzer applies a set of rules to the hardware state and from those rules attempts to determine the cause of the fatal error.

  It evaluates which boards, if any, need to be replaced and assigns a percentage confidence level to each of its evaluations. Following analysis, the FRU analyzer reports failure information about the boards that scored the highest probability of failure.

- IRIX Crash Analysis Tool (ICRASH)

  ICRASH can greatly enhance an engineer's ability to analyze system core dumps. ICRASH provides additional information in the event of a system crash by sorting through the core dump file and reporting relevant information. ICRASH is also useful when interactively analyzing a functional system.

  ICRASH contains a wide range of functions that display data from the internal data structure of the IRIX kernel. The function output can be piped to such UNIX commands as *more*, *grep*, and *pg*. Online help provides information about each ICRASH function.

- Availability Monitor (AvailMon)

  AvailMon, together with ICRASH and the FRU analyzer, provides a technology platform for system availability and diagnostic data gathering and distribution. AvailMon is embedded in the system boot and shutdown processes. It differentiates controlled shutdowns, system panics, system hangs, power cycles, and power failures. A lightweight daemon tracks uptime and collects diagnostic information from ICRASH, SYSLOG, *hinv*, *versions*, and *gfxinfo*.

  All availability and diagnostic data for participating systems is archived in a Silicon Graphics database. Access to that database provides overall reliability data and a specific problem history for individual systems.

- Remote Access Tool (RAT)

  RAT is a software utility that resides on the multimodule System Controller (MMSC). It enables field and factory personnel to perform various hardware configuration and troubleshooting procedures on the computer system from a location other than the customer site. The user communicates with RAT through a standard telephone line and modem connection.

  RAT provides a command-line interface and ASCII menus for user control. Using simple keyboard entries, the user may power on / off or reset the system, access diagnostics, set and clear LEDs, read the system configuration, and check or modify vital signs (fan speeds, voltages, and temperature).

7.6 Hardware Graph

The hardware graph is a tool for inventorying the I/O devices of the Origin2000 and Onyx2 systems. Unlike *hinv*, the hardware graph is a UNIX filesystem whose branching character accommodates the possibility of multiple nodes, each with multiple I/O devices of several types. The hardware graph can also track the number of active graphic pipes in an Onyx2 rack or multirack configuration. The hardware graph keeps track of information in the kernel that is associated with the hardware.

Note: For information on how the pipes are numbered, see "Pipe Numbering in the Onyx2 Rack System" in Chapter 10.

Most of the hardware graph directories are much like their */dev* counterparts, but module numbers are persistent across reboots and hardware changes (until you change the module numbers).

To see the hardware graph, use the *ls* command. For example,

```
# ls /hw
console    mem       module     rdisk     ttys      scsi_ctlr   unknown
disk       kmem      mmem       null      scsi      ttys        zero
```

In this output, `module`, `rdisk`, `ttys`, `scsi`, `scsi_ctrl`, and `ttys` are subdirectories containing files. For example,

```
# ls /hw/ttys
tty4d1  tty4f1  tty4m1  ttyc1   ttyd1   ttyf1   ttym1
tty4d2  tty4f2  tty4m2  ttyc2
```

```
# ls /hw/scsi
sc1d2l0
```

```
# ls /hw/rdisk
dks1d2s0       dks1d2vh       root           volume_header
dks1d2s1       dks1d2vol      swap
```

```
# ls /hw/scsi_ctlr
0  1
```

To determine modules present in the system, enter

```
# ls /hw/module
1  2
```

This output indicates two modules.

Modules and slots are numbered with base 1; CPUs, nodes, and PCI slots are numbered with base 0. The hardware graph lists only devices actually present, skipping the numbers of unfilled slots.

Note: If a system detects an inconsistency in module numbering, the system prompts the user to enable automatic renumbering and may reassign the ConsolePath variable, which determines the console connection. For more information, refer to "Correcting ConsolePath and Console-moving Problems" in Chapter 10.

To determine I/O devices of the module, follow the directory structure. Figure 7-4 shows a simplified example of a */hw/module* filesystem with two modules; in the second module, there are two CPUs and one XIO board.

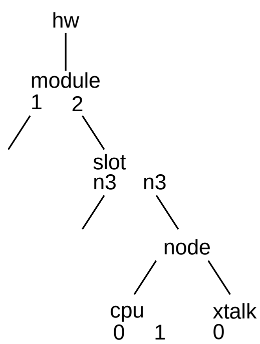

Figure 7-4 Example *hw* Module Filesystem

For example:

```
# ls /hw/module/1/slot/n4/node/link/cpu
0  1
# ls /hw/module/1/slot/n4/node/link/xtalk
0
```

To see connectivity of an I/O device, use *ls* in the */link* subdirectory with the *-l* flag. For example, determine the links present in the system:

```
# ls /hw/module/1/slot/n1/node/link
1  2  3  5  6
```

Then determine the connectivity:

```
# ls -l /hw/module/1/slot/n4/node/link
lrw-------    0 root      sys          28 Jul 29 01:22
/hw/module/1/slot/n1/node/link -> /hw/module/1/slot/r2/router
```

To determine a router link, use commands such as the following:

```
# ls /hw/module/1/slot/r2/router
1  2  3  5  6

# ls -l /hw/module/1/slot/r2/router/3
lrw-------    0 root      sys          28 Jul 29 01:22
/hw/module/1/slot/r2/router/3 -> /hw/module/1/slot/n2/node
```

7.6.1 The /etc/ioconfig Utility

Because the I/O devices are long (for example, `/hw/module/1/slot/n4/node/link/xtalk/0`) and do not look like previous device names, the *ioconfig* utility maps short names and controller numbers to long platform-dependent names. For example

The *ioconfig* utility keeps mappings in a file so that they do not change. It creates names only for devices that exist. To avoid confusion, controller numbers are not reused for different slots. The utility opens each disk's volume header in parallel and creates device nodes.

7.7 Flashing the PROMs

The IP27 PROMs need to be flashed whenever there is a Node board upgrade. Follow these procedures to flash the PROMs in a large system configuration. For smaller system configurations, use the *flashcpu*, *flashio*, or *flash* commands.

There are two methods to flash the PROMs:

- UNIX flash using *inst -f*
- through the IO6 command monitor using *flash -T*

If you use the UNIX flash and the PROM gets corrupted, you don't notice the problem until after you reset the system. It is therefore recommended that you use the IO6 PROM flash to flash multiple Node boards.

1. At the command monitor prompt, type **`flash -T`** to flash the first Node board.

   You should get a display similar to the following:

```
> c8000/d9ac0 ........
> d0000/d9ac0 ........
> d8000/d9ac0 ..
> d9ac0/d9ac0
> Programmed and verified
Flashing CPUPROM in module 2 slot n4
PROM log appears uninitialized (see initlog command)
promlog: Flashing log sectors, please wait.
promlog: Error writing header 0: Programming operation timed out
Cannot clear PROM log: Flash PROM error
Failed to initalize prom log
Flashing CPUPROM in module 2 slot n3
PROM log appears uninitialized (see initlog command)
promlog: Flashing log sectors, please wait.
promlog: Error writing header 0: Programming operation timed out
Cannot clear PROM log: Flash PROM error
Failed to initalize prom log
Flashing CPUPROM in module 2 slot n2
PROM log appears uninitialized (see initlog command)
promlog: Flashing log sectors, please wait.
promlog: Error writing header 0: Programming operation timed out
Cannot clear PROM log: Flash PROM error
Failed to initalize prom log
Flashing CPUPROM in module 1 slot n2
```

2. Turn the MSC keyswitch to the diagnostics position and press NMI switch to go to POD (power on diagnostics) mode.

   Note: The NMI switch issues a reset signal to all Node boards in the system.

3. Enter the cached mode using the *go cac* command.
4. Flash across to other Node boards one at a time.

7.8 Updating MMSC Firmware

Note: This information is from http://homegrown.engr/MMSC/commands.html and http://homegrown.engr/MMSC/flashing.html.

The firmware for the Origin2000 MMSC resides in nonvolatile storage on the MMSC itself. You can replace the existing firmware with a new firmware image by either of two methods:

- if there has been a catastrophic failure or if the nonvolatile storage that usually holds the MMSC is erased, use the serial downloader
- if the MMSC is operational and needs only to be upgraded to a new version, upgrade the firmware from IRIX

This section explains

- using the serial downloader
- upgrading MMSC firmware from IRIX
- using the *flash* command

7.8.1 Using the Serial Downloader

After certain catastrophic failures (or when an MMSC first arrives from the vendor), the nonvolatile storage that usually holds the MMSC firmware is erased. When an MMSC is powered on in this condition, it often shows a blank blue screen, or perhaps some PC-style BIOS configuration gibberish. Under these circumstances, you must flash the firmware image into nonvolatile storage using the serial downloader, a simple tool provided in the MMSC's BIOS.

This procedure requires a direct serial connection from either a PC or an IRIX system to the COM5 port on the afflicted MMSC. The serial connection should be set up for 19200/8/N/1 (only!). If you use a PC, it must have a terminal emulator capable of performing XMODEM or XMODEM-1K file transfers. If you use an IRIX system, it should have the *flashffsc* program if possible, although you can use an XMODEM file-transfer program

In addition, the MMSC must have a display attached to it. If the MMSC in question does not reside in the rack with the display panel, then you must unbolt the MMSC from the rack and move it closer to the display, or vice versa.

7.8.1.1 Starting the Serial Downloader

To start the serial downloader, follow these steps:

1. Make sure all users are logged off the system. ***Q I presume this is necessary***
2. Power on the MMSC and simultaneously hold down the bottom two buttons (**Down** and **Enter**) of the MMSC display. ***Q how are these buttons labeled?***

   This procedure is fairly awkward and usually requires two people, one to stand at the front of the system and hold down the bottom two buttons of the display, and the other to go to the back of the system, unplug, and then reconnect the MMSC from its power supply. There is actually a short delay before the buttons need to be pressed; if you are desperate or just gymnastically inclined, you might be able to power-cycle the MMSC and then run around to the front panel in time to hold down the buttons.
3. Hold down the display buttons until the serial downloader menu appears. It looks approximately like this:

   ```
   SBC-MMSC Serial Downloader 1.0
   (U)pload an application
   ```
4. Use either a terminal emulator or the *flashffsc* program to transfer the firmware image, as explained in the next sections.

7.8.1.2 Transferring the Firmware Image Using a Terminal Emulator

If the COM5 port of the MMSC is attached to a PC or other system capable of an XMODEM file transfer, you flash the firmware by uploading it into the MMSC. The menu in step 3 above should be echoed to the PC display; if it was not, press Enter once or twice, check the connections and baud rate, and make sure you are using a null modem cable (not one wired straight through) to connect the MMSC and PC.

Once you see the serial downloader menu from the MMSC, follow these steps:

1. Type U to upload.
2. Using either the XMODEM-CRC or XMODEM-1K file transfer protocol, instruct your terminal emulator to send the firmware image file, usually *MMSCfw.bin*.

   Note: *flashffsc* and the latest *MMSCfw.bin* file are at homegrown.engr:/usr/dist/flashffsc and homegrown.engr:/usr/dist/MMSCfw.bin.

   The transfer takes approximately 10 minutes. The display shows a series of dots to indicate progress; a typical firmware image takes 10-12 lines of dots to complete.

7.8.1.3 Transferring the Firmware Image Using flashffs

If the COM5 port of the MMSC is attached to a serial port on a system running IRIX, you can flash the image using the *flashffsc* program. The format of the command is

```
flashffsc -d -l /dev/ttydXX -f ffscfw.bin
```

where */dev/ttydXX* is the device name of the direct serial port you are using and *ffscfw.bin* is the file containing the firmware image you intend to load.

Run this program after the display shows the serial downloader menu described in step 3 in "Starting the Serial Downloader." It may take 10 to 20 seconds for the file transfer to begin; once it does, a series of dots is printed on both the MMSC display and the IRIX window. If these dots do not start after 30 seconds or so, check your connections as described in "Transferring the Firmware Image Using a Terminal Emulator."

7.8.2 Upgrading MMSC Firmware From IRIX

If the MMSC is operational and needs only to be upgraded to a new version, upgrade the firmware from IRIX using *flashffsc*. This command is assumed to be issued from a terminal that has access to one of the MMSCs associated with the system. To use this command, follow these steps:

1. Make sure all users are logged off the system.
2. Make the terminal access the MMSC by typing the MMSC escape character, typically Ctrl+T.

   To determine the MMSC escape character, type *esc ?*.
3. Type the appropriate *flash* command; for example, to flash a new firmware image onto rack 2, use

   ```
   % flashffsc -m
   ```

   The following text appears:

   ```
   Ready to transfer new image to full-feature system
   controller. To begin the transfer, type your ffsc
   escape character (normally CTRL-T) followed by the
   command:

       rack <rackid> flash

   where <rackid> is the identifier for the system
   controller you wish to upgrade.
   ```
4. Press the MMSC escape character (default Ctrl+T). The following message appears:

   ```
   FFSC> rack 2 flash
   Waiting for MMSC to initiate transfer...
   ```

   This command takes the firmware image from the default location */usr/cpu/firmware/ffscfw.bin*. To specify a different image location, use the **-f** option:

   ```
   flashffsc -m -f firmware_location.bin
   ```

Although this procedure is faster than using the serial downloader, it can still take ten or more minutes to run, depending on the speed of the connection between the IO6 and the MMSC. As with the serial downloader, a series of dots shows progress.

For full details of *flash*, see “Using the flash and reset_ffsc Commands.”

7.8.3 Using the flash and reset_ffsc Commands

flash [from system]

flash from console

Flashes a new firmware image into nonvolatile storage on the MMSC for the addressed rack. An error occurs if this command addresses more than one rack. The bay portion of the destination is ignored.

In the first form, *flash from system* (the words *from system* are optional), the image is provided to the MMSC from the console of a running IRIX system using the flashffsc(1m) command. The *flashffsc* IRIX command and the *flash* MMSC command must both be issued from the same terminal device. Typically, the *flashffsc* command is issued first with the **-m** option, an MMSC escape character is typed, and the *flash* command entered from the same terminal.

In the second form, *flash from console*, the image is read from the terminal device itself. This form is useful when the user terminal is in fact a terminal emulator (perhaps on a PC) capable of XMODEM file transfer. For this form of the *flash* command, issue the *flash* command to the MMSC first, and then arrange for an XMODEM or XMODEM-1K file transfer of the firmware image from the user terminal. The use of XMODEM-1K can decrease the transfer time by as much as 50%.

The *flash* command takes a long time to run. In addition to the time it takes to download the firmware image over a serial line (MMSC firmware images are approximately 1MB), it take an additional 100 seconds to clear out the flash RAM before the new image is installed.

Caution: Do not interrupt the MMSC while this command is in progress.

Early versions of the MMSC do not have sufficient flash storage to hold two separate images. Therefore, attempting to flash a new image wipes out the old firmware image, even if the attempt turns out to be unsuccessful. Do not reboot an MMSC that has tried and failed to flash a new firmware image. If failure does occur, download the MMSC firmware as explained in “Using the Serial Downloader.”

This command accepts an optional destination specification. For information on specifying destination and information on other MMSC commands, see Appendix C.

reset_ffsc Restarts the addressed MMSC(s). After a new firmware image is flashed onto the MMSC, this command can be useful in the process of picking up the changes. This command can also be used if the MMSC or system console has hung even though the system itself is still running. Of course, if this command is used for this second purpose, you should file a bug report or service call.

This command does not affect IRIX or any other part of the system. Only the addressed MMSCs are rebooted. For historical reasons, the command *reboot_ffsc* does the same thing as *reset_ffsc*.

This command accepts an optional destination specification.

Chapter 8

Memory Upgrades

This chapter provides information about memory upgrades for the Node board. The amount of memory on the Node board cna be increased in only two areas: .

- Main memory—provided through dual inline memory modules (DIMMs). All memory is accessable to all processors using a single, flat memory address space. The main memory DIMMs pairs implement 128 bits of memory along with additional bits of memory for ECC.
- Premium directory memory—provided through a separate set of DIMMs. This memory maintains cache coherency in configurations with more than 16 nodes. The premium directory memory is *not* needed in configurations with 16 or fewer Node boards.

  Note: Premium directory memory is not required for the Onyx2, because the largest configuration contains up to 12 Node boards (or up to 24 processors).

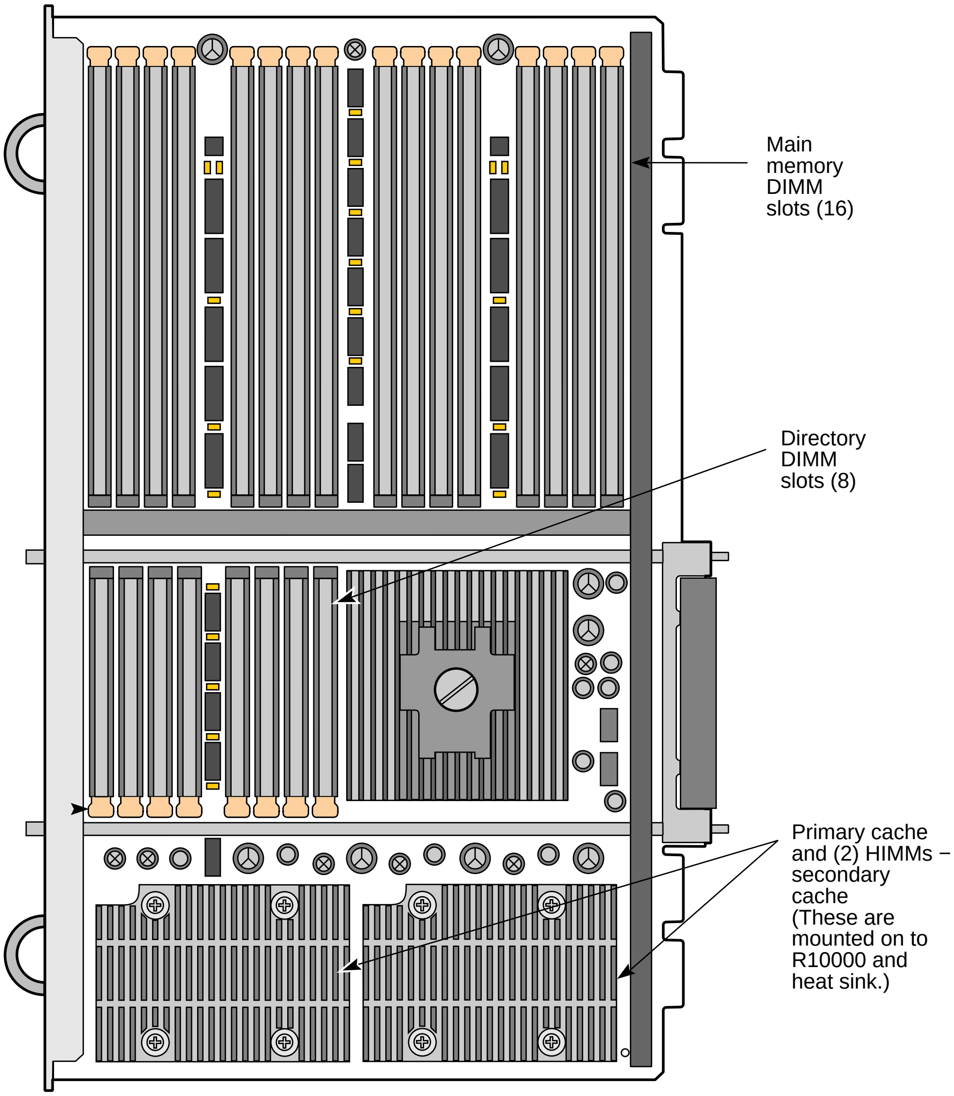

Figure 8-1 Memory Module Locations on IP27 Node Board

Note: There are 16 slots for main memory DIMMs and 8 slots for directory DIMMs.

8.1 Main Memory Installation

The main memory DIMMs come in three types (see Table 8-1 and Figure 8-2). A color code and a part number differentiate the DIMM types.

Table 8-1 Main Memory DIMM Types

DIMM Type	Part Number	Color Bar Code
32 MB	030-0759-xxx	Yellow

Table 8-1 Main Memory DIMM Types

DIMM Type	Part Number	Color Bar Code
64 MB	013-1372-xxx	Green
256 MB	013-1492-xxx	Red

Note: The 256-MB DIMM is not available for the first customer shipments.

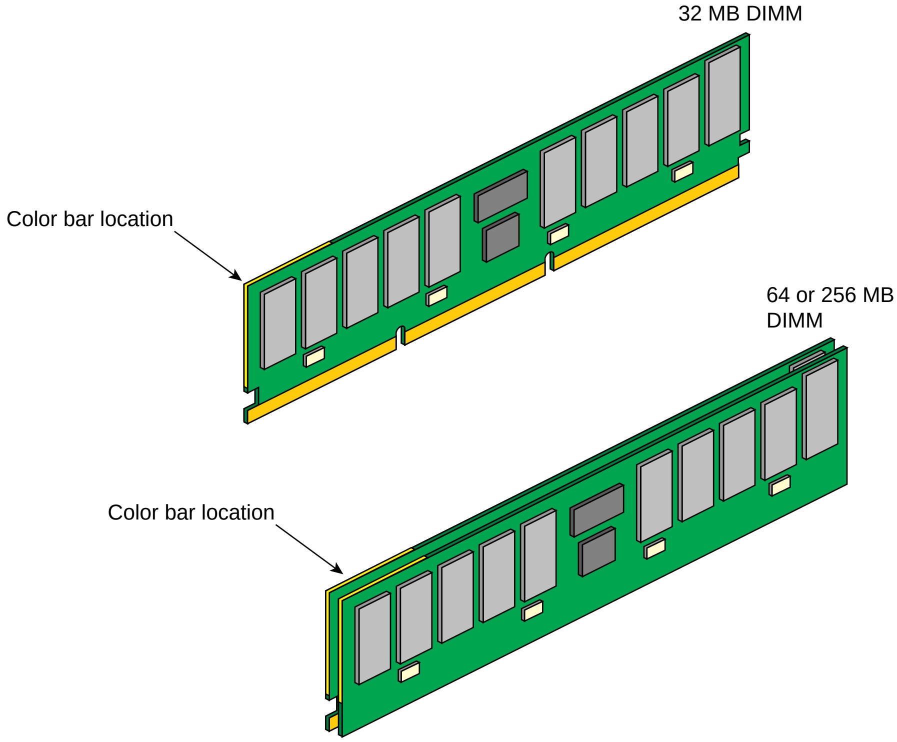

Figure 8-2 Main Memory and Directory DIMMs

Note: The 64- and 256-MB DIMMs are composed of two stacks of boards. The 32-MB DIMM is single board.

8.1.1 Installation Guidelines

Observe these guidelines when installing memory. Figure 8-3 shows the locations of the memory DIMMs and banks.

- Install DIMMs in pairs, a bank at a time (see Table 8-2 and Figure 8-4).

  Note: The minimum system configuration is 64 MBs using two 32-MB DIMMs.

- Populate the lowest banks (bank 0, MMXL0 and MMXH0) first, then proceed to the higher banks (MMYL7 and MMYH7)
- DIMMs in pairs must be of the same type.
- Between DIMM pairs, different DIMM sizes can be used (i.e., the largest DIMMs do not need to be in the lowest DIMM slots).

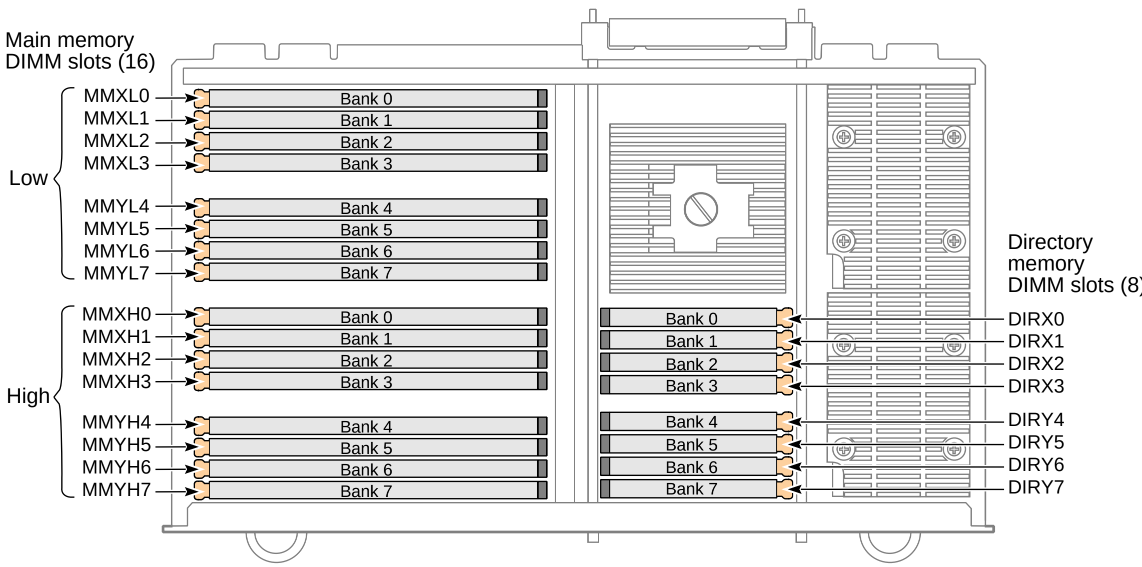

Figure 8-3 Memory DIMM and Bank Locations (Side View)

Table 8-2 Main Memory DIMM Locations

DIMM Pair	Memory DIMM Locations	Bank Number
First	MMXL0 and MMXH0	0
Second	MMXL1 and MMXH1	1
Third	MMXL2 and MMXH2	2
Fourth	MMXL3 and MMXH3	3
Fifth	MMYL4 and MMYH4	4
Sixth	MMYL5 and MMYH5	5
Seventh	MMYL6 and MMYH6	6
Eighth	MMYL7 and MMYH7	7

8.1.2 Installing Main Memory DIMMs

Use the following instructions along with Figure 8-4 and Figure 8-5 to install the main memory DIMMs on the Node board. A DIMM installation tool is needed to help push the DIMMs into the sockets. Review Section 8.1.1, "Installation Guidelines," before you begin.

Caution: Observe proper ESD practices (such as using a ground strap and an antistatic mat) when installing DIMMs.

1. Place each DIMM board directly into a single DIMM location with the notch at the ejector tab end. The DIMMs are keyed to prevent improper installation.
2. When the DIMM board is in position, use the DIMM installation tool to push it down into place. The DIMM tool helps provide sufficient force to install this memory board

   After the DIMM is properly installed, the ejector tab should be *up* (not down).
3. Install the DIMMs a bank at a time.

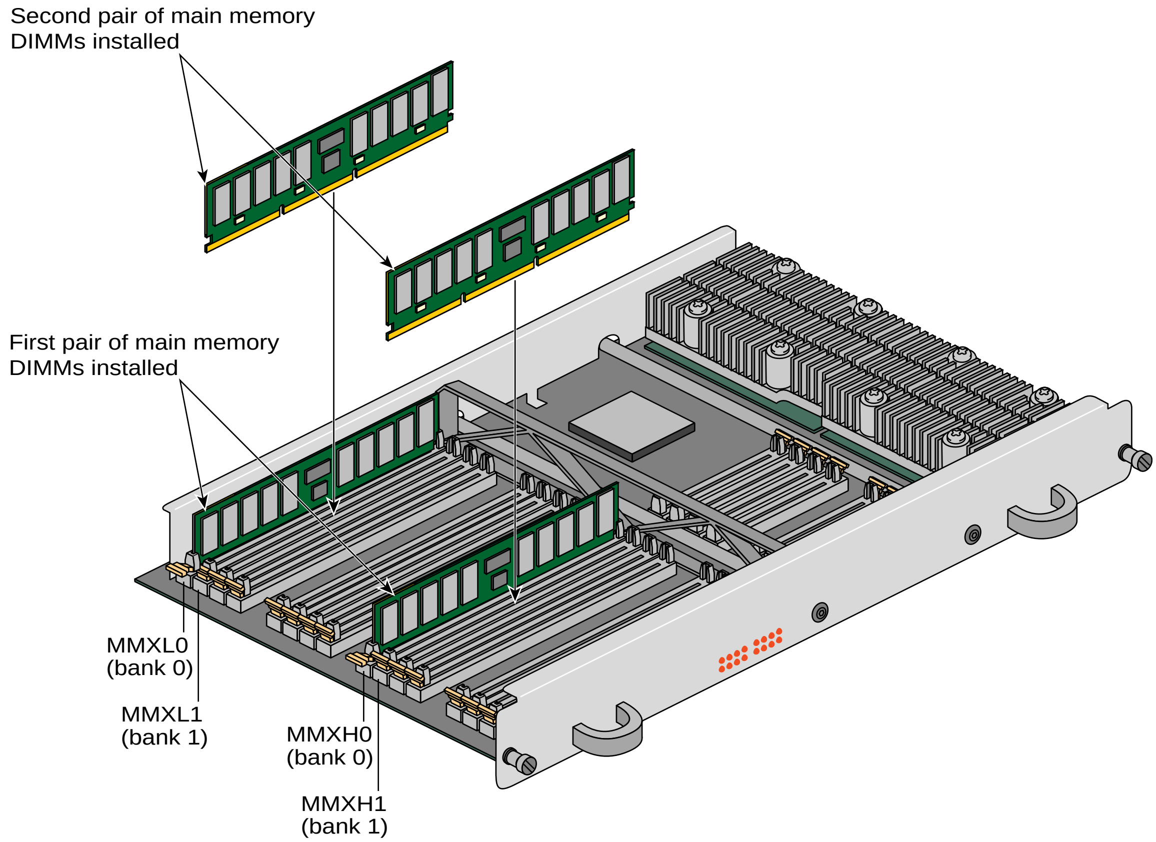

Figure 8-4 DIMM Installation Location

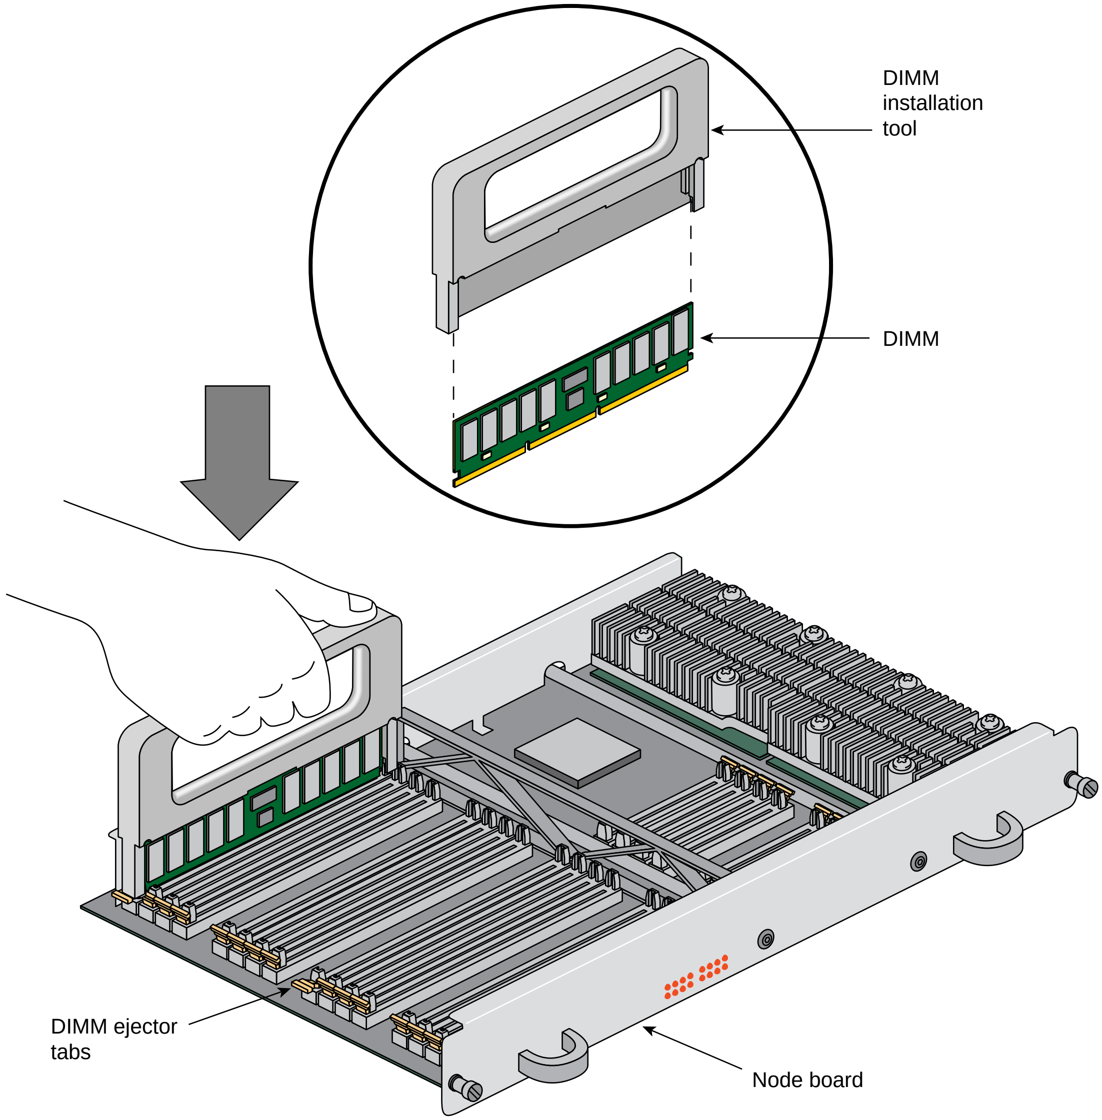

Figure 8-5 Installing DIMMs With the DIMM Installation Tool

8.1.3 Removing DIMMs

Follow these instructions to remove the DIMMs.

Caution: Observe proper ESD practices (such as using a ground strap and an antistatic mat) when removing DIMMs.

1. Push *down* on the DIMM ejector tab as shown in Figure 8-6.
2. Carefully remove the DIMMs and set them aside.

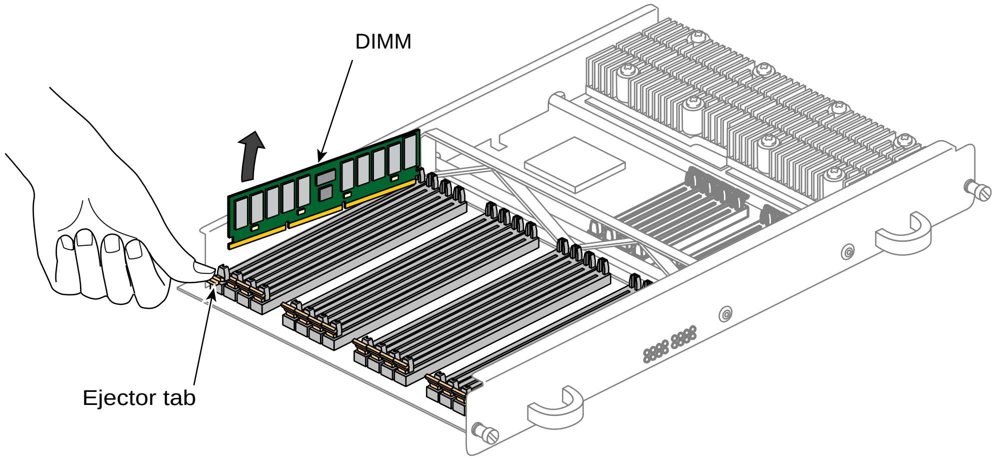

Figure 8-6 Main Memory DIMM Removal

8.1.4 Main Memory Interleaving

Interleaving enables greater utilization of bus bandwidth, as requests are evenly distributed to help mask read latency. This results in significant read and write access improvement.

Memory interleaving for the Origin2000 and the Onyx2 systems is automatically done through the Hub. In the CHALLENGE and Onyx systems, interleaving had to be aided manually by moving the single-inline memory modules (SIMMs) into different positions whenever the memory configurations changed. The Hub performs four-way interleaving.

8.2 Premium Directory Memory

The premium directory DIMMs are required only for configurations with *more* than 16 Node boards. For systems with 16 or fewer Node boards, directory memory is maintained using the directory memory bits that are part of the main memory DIMMs. The premium memory DIMMs implement the required extra 32 bits of directory width for configurations with more than 16 Node boards.

Note: The premium directory DIMMs are not required for any Onyx2 configuration.

Directory memory keeps track about which caches have a copy of the memory line and maintains cache coherency. Directory memory keeps counters per page and per node of the number of cache misses from that node.

8.2.1 Premium Directory Memory Installation Rules

In addition to the greater-than-16-node requirement, there are additional rules to observe when installing premium directory memory.

8.2.1.1 Install Main Memory and Directory Memory DIMMs in Triplets

You must add a premium directory DIMM for each main memory DIMM pair (see Table 8-3). In other words, the premium directory DIMMs and main memory DIMMs are installed in triplets. For example, if the Node board has main memory DIMMs in bank 0, you must install a corresponding directory DIMM in bank 0 in the premium directory memory slots.

Table 8-3 Main Memory DIMMs and Directory Memory DIMs Triplet Pairing

DIMM Triplet	Main Memory DIMM Location	Premium Directory Memory Location
MMXL0, MMXH0, DIRX0	Bank 0	Bank 0
MMXL1, MMXH1, DIRX1	Bank 1	Bank 1
MMXL2, MMXH2, DIRX2	Bank 2	Bank 2
MMXL3, MMXH3, DIRX3	Bank 3	Bank 3
MMYL4, MMYH4, DIRY4	Bank 4	Bank 4
MMYL5, MMYH5, DIRY5	Bank 5	Bank 5
MMYL6, MMYH6, DIRY6	Bank 6	Bank 6
MMYL7, MMYH7, DIRY7	Bank 7	Bank 7

8.2.1.2 All Node Boards Must Have Premium Directory Memory Installed

You must install directory DIMMs in *all* the Node boards in the configuration. For example, if Node 1 has 4 banks of main memory, you need to install four banks of directory DIMMs into their corresponding slots. If Node 2 has three banks of main memory, you need to install three banks of directory DIMMs into the corresponding slots. You must continue this process until all the Node boards in the configurations have the corresponding directory DIMM slots filled.

8.2.1.3 Directory DIMM Must Match Main Memory Pair Type

Table 8-4 shows the corresponding main memory and directory memory sizes. The 32- and 64-MB main memory DIMMs require a type-D directory DIMM. The 256-MB main memory DIMM requires a type-E directory DIMM.

Table 8-4 Corresponding Main Memory and Directory Memory Colors

Memory Size	Total Memory Per DIMM Pair	Memory Color	DIMM Color
32 MB	64 MB	Yellow	White
64 MB	128 MB	Green	White
256 MB	512 MB	Red	Blue

8.2.2 Directory DIMM Installation

The premium directory DIMMs can be installed without the use of a DIMM installation tool (see Figure 8-7). These DIMMs are shorter than the main memory DIMM and do not require as much force to install them.

Caution: Observe proper ESD practices (such as using a ground strap and an antistatic mat) when installing DIMMs.

1. Place each DIMM directly into a single DIMM location with the notch at the ejector tab end. The DIMMs are keyed to prevent improper installation.
2. Install the DIMMs a bank at a time.
3. Push down firmly with your thumbs to seat the DIMMs (see Figure 8-7). When the DIMM is properly installed, the ejector tab should be *up* (not down).

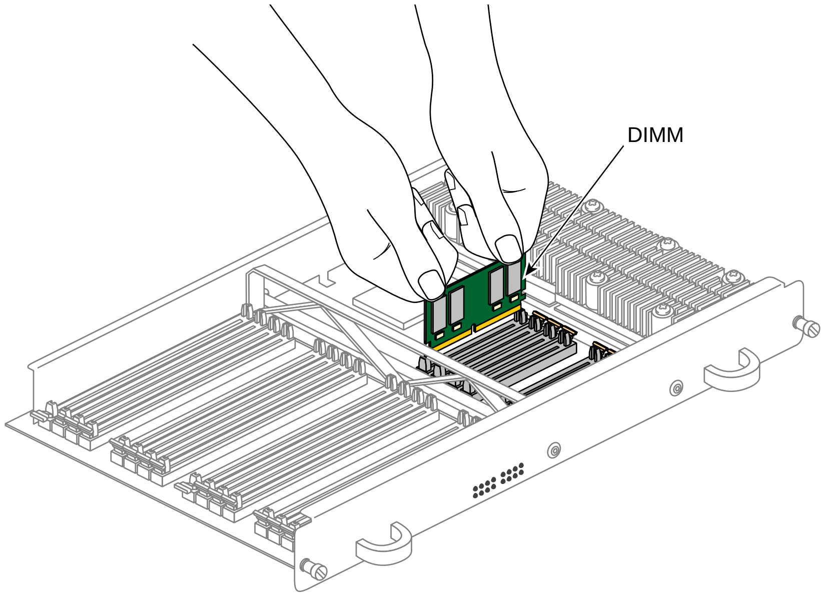

Figure 8-7 Installing Premium Directory DIMMs

8.2.3 Removing DIMMs

Follow these instructions to remove the DIMMs.

Caution: Observe proper ESD practices (such as using a ground strap and an antistatic mat) when removing DIMMs.

1. Push *down* on the DIMM ejector tab (see Figure 8-8).
2. Carefully remove the DIMMs and set them aside.

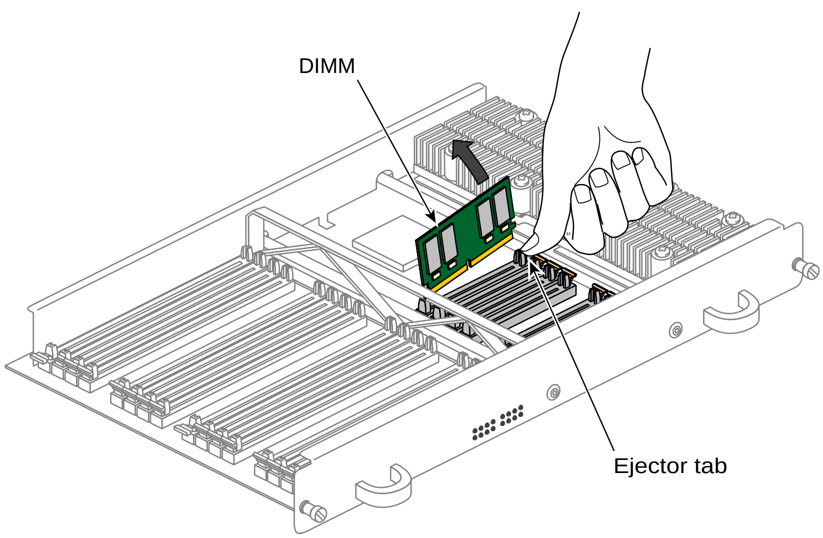

Figure 8-8 Removing Premium Directory DIMMs

8.2.4 HIMM for Secondary Caches

The horizontal inline memory module (HIMMs) provide a secondary cache for the IP27 Node board. The HIMMs are mounted along with the R10000 and heat sink on a single assembly (see Figure 8-9).

A single IP27 board supports the following configurations:

- minimum Scache: 1 MB (using one 1-MB HIMM)
- maximum Scache: 8 MB (using two 4-MB HIMMs)

To increase the Scache size from 1 MB to 4 MB, you must replace the entire R10000 and heatsink module. You cannot take the R10000 / HIMM assembly apart.

Note: Mixing 1-MB and 4-MB HIMMs in one IP27 is not allowed.

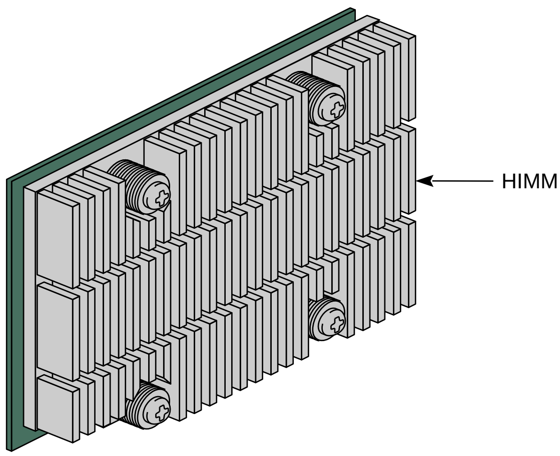

Figure 8-9 R10000 / HIMM Assembly

Chapter 9

Installing a Deskside or Graphics Module Into a Rack

This chapter describes how to install an Origin2000 deskside or an Onyx2 graphics expansion module in the new Silicon Graphics 19-inch rack or any standard 19-inch rack. This enables you to expand system configuration and provide greater grounding capability.

Caution: It is recommended that two people perform this procedure. Installation requires lifting the module chassis into a rack.

You can install up to two deskside modules in a rack (see Figure 9-1). If you are installing only one module, it needs to be installed in the lower half of the rack for greater balance and stability.

9.1 Reviewing Site Preparation Requirements

Consult the *Site Preparation Guide for Origin Family and Onyx2* (P / N 007-3452-xxx) for site preparation requirements for a rackmount system. For example, a rack system requires 220-volt power. Therefore, you must ensure that the site has the adequate power support. In addition to power requirements, here are some other important site preparation factors to consider:

- Floor preparation. Be sure that the floor is rated to support the additional weight.
- Maintenance space around cabinet. Be sure that you have sufficient space to install a rack system.
- Environmental factors such as ambient air temperature and humidity. Be sure that the air conditioning can handle the additional cooling requirements of a rackmount system.
- Be sure to have a single grounding source for the rack setup.

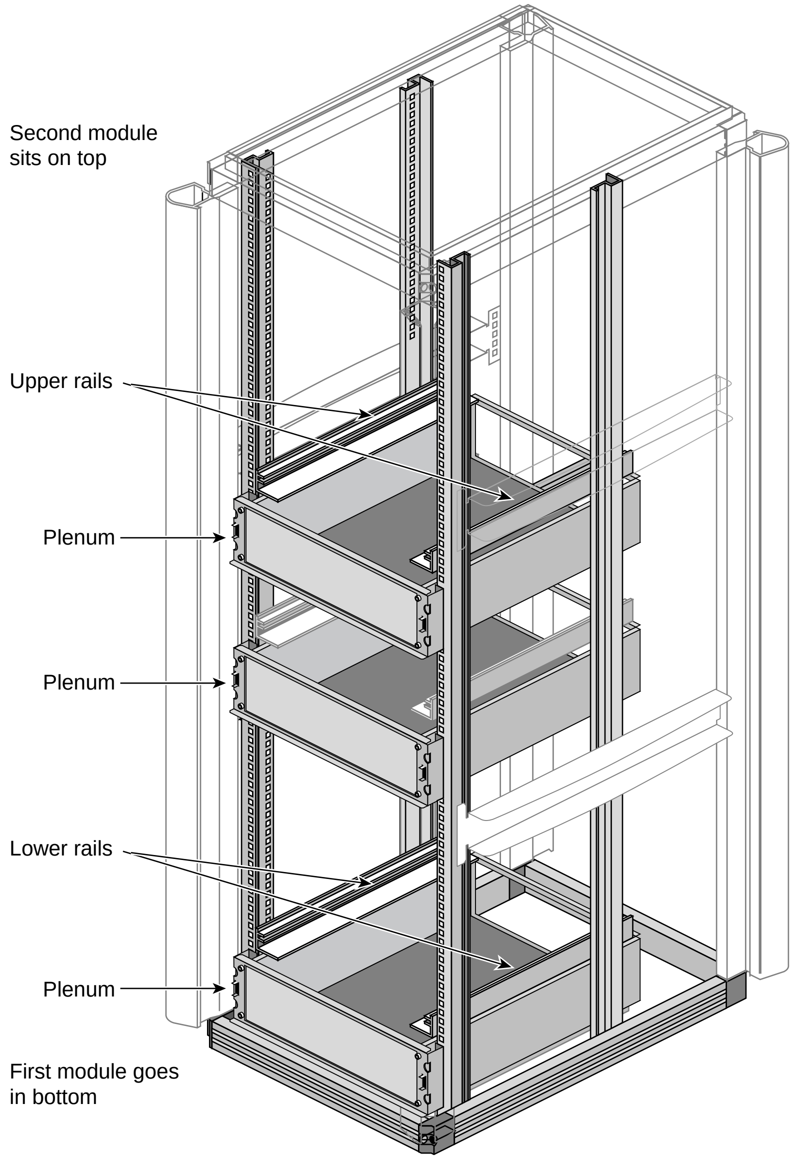

Figure 9-1 Origin2000 Rackmount Mounting Rails for Modules

9.2 Preparing the Module

Follow these guidelines to prepare a deskside or graphics module for installation into a rack system.

9.2.1 Deskside Module

Be sure that the deskside module is light enough to lift it into a rack by removing the following parts. Make sure that the system has been safely powered off and disconnected (see Section 10.2.1 for additional information on shutting off the system).

- Remove the facade (see Section 10.2.7).
- Remove the tophat (see Section 10.2.8).
- Remove the side skins (see Section 10.2.9).
- Remove the bottom caster assembly (see Section 10.2.10).
- Remove the drives (see Section 10.3.1).
- Remove the power supply (see Section 10.3.9).
- Remove the fan tray (see Section 10.3.10).

9.2.2 Graphics Module

Ensure that the graphics module is light enough to fit it into a rack by removing the following parts. Be sure that the system has been safely powered off and disconnected (see Section 10.2.1 for additional information on shutting off the system).

- Remove the power supply (see Section 10.4.2).
- Remove the fan tray (see Section 10.4.3).
- Remove the graphics boards, if required to lighten the module even further (see Section 10.4.1).

9.3 Installing a Deskside Module Into an Origin2000 Rack

After the site preparation requirements have been met (see Section 9.1 and the *Site Preparation Guide for Origin Family and Onyx2*, P / N 007-3452-xxx) and the deskside module has been stripped of essential components for weight reduction (see Section 9.2), you can install the deskside module into an Origin2000 rack.

The Origin2000 rack has a MMSC and display along with built-in cable management hardware (see Figure 9-2).

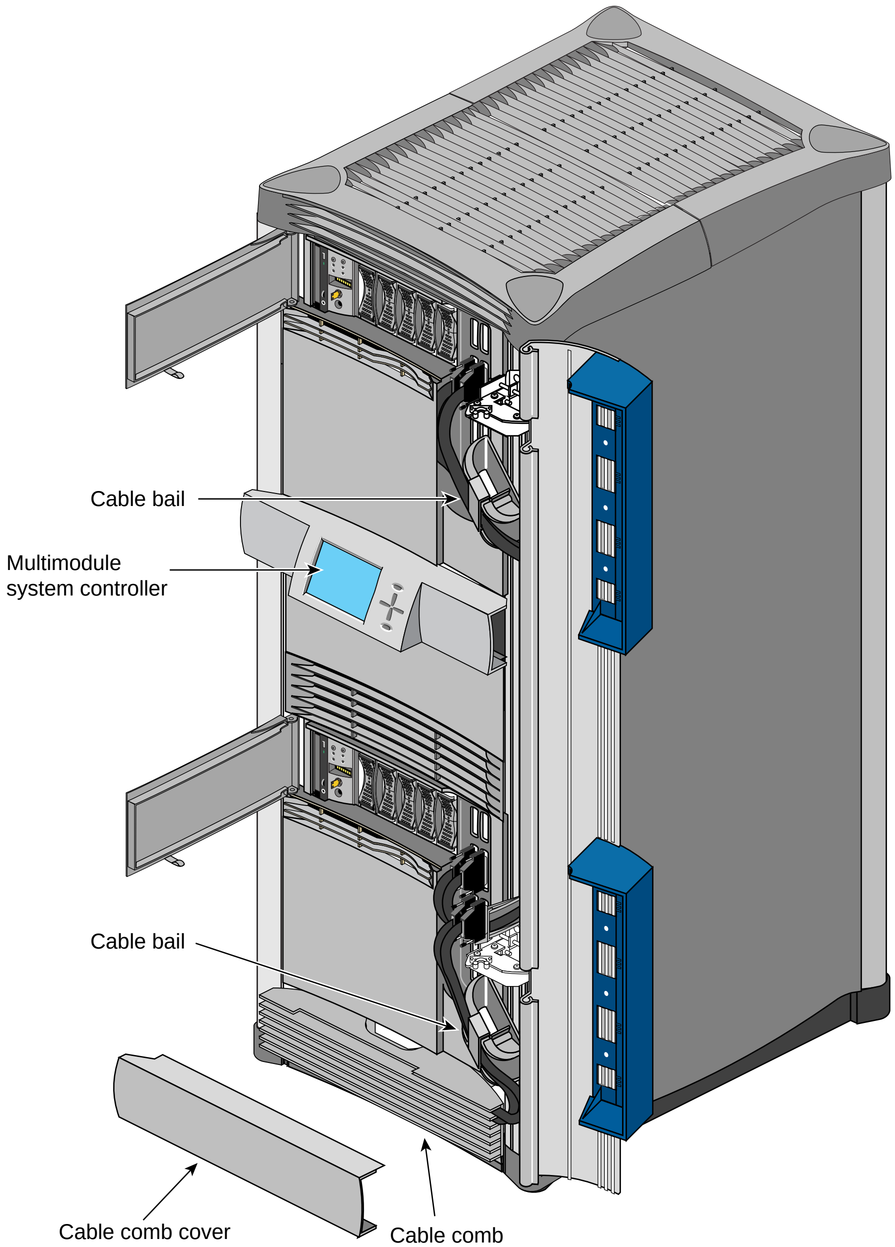

Figure 9-2 Origin2000 Rackmount Chassis with Cable Management Hardware and Multimodule System Controller

Follow these instructions to install a module in a rackmount system.

1. Add the front U-brackets to the sides of module in the front (see Figure 9-3).
2. Carefully lift and lower the stripped deskside module into the lowest set of rackmount rails, if no existing module is present (see Figure 9-3).

   Caution: If you are installing only one module, it needs to be installed in the lower half of the rack for greater balance and stability.

3. Secure the front of the module to the chassis using the U-brackets.
4. Attach the Z-brackets to secure the rear of the module to the back of the chassis (see Figure 9-4).
5. If you are installing the module in the upper set of rails, carefully raise the module into the upper set of rails (see Figure 9-1 for placement).
6. Secure the upper module to the chassis as described in steps 1 through 4.
7. After installing the module(s), go to Section 9.6 to complete installation.

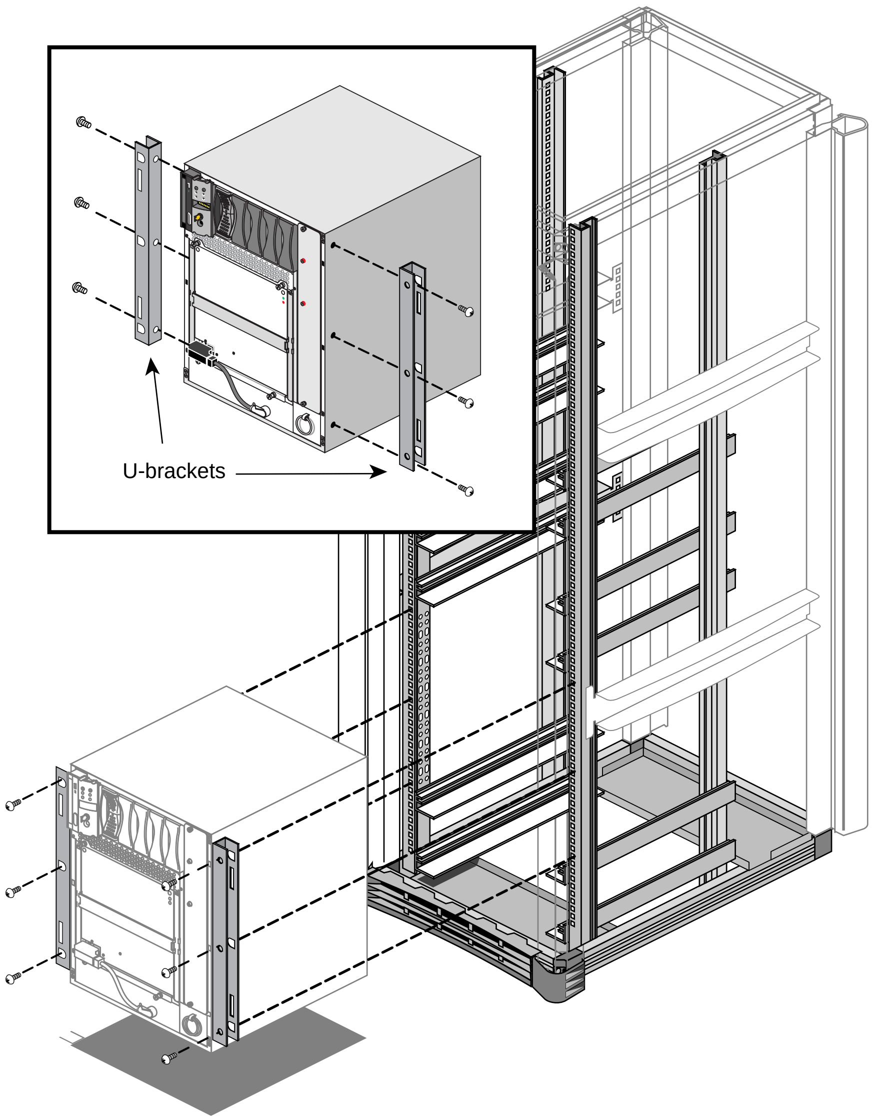

Figure 9-3 Installing an Origin2000 Deskside Module Into a Silicon Graphics 19-Inch Rack (Front View)

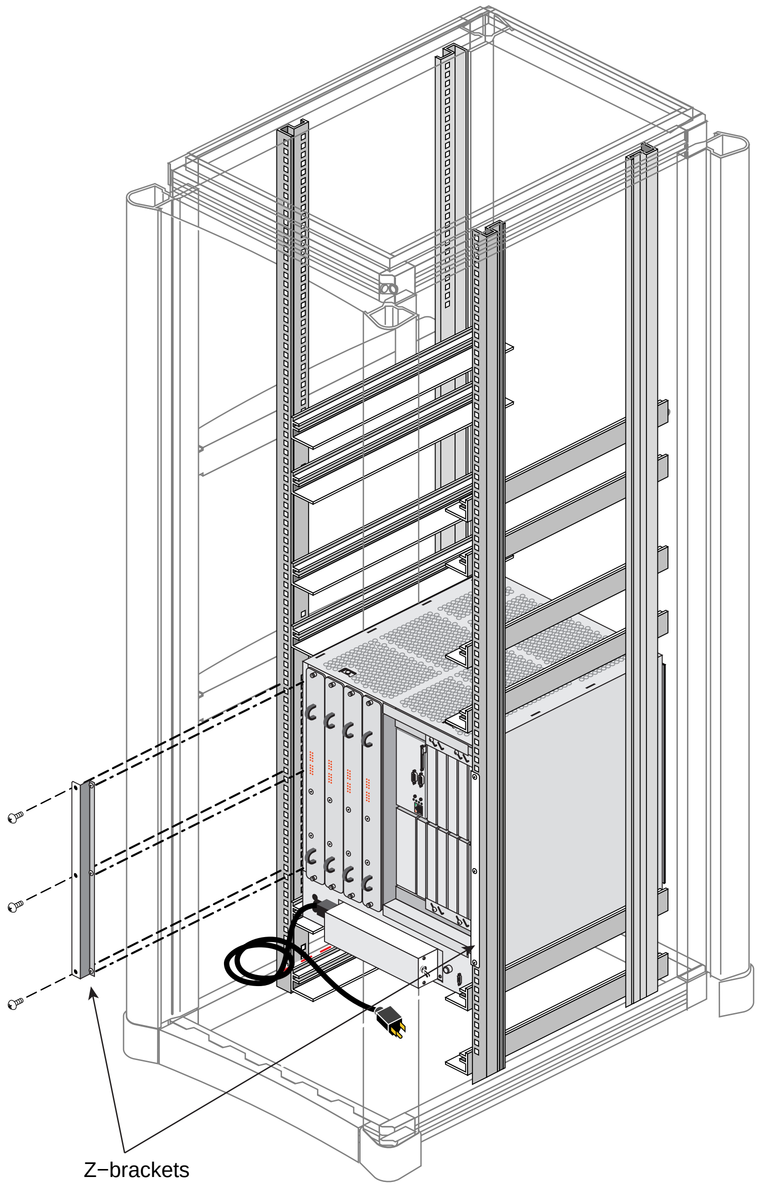

Figure 9-4 Installing an Origin2000 Deskside Module Into a Silicon Graphics 19-Inch Rack (Rear View)

9.4 Installing a Graphics Module In a Rack

Follow these instructions to install a module in a rackmount system.

1. Add the front U-brackets to the sides of module in the front (see Figure 9-5).
2. Carefully lift and lower the stripped deskside module into the lowest set of rackmount rails, if no existing module is present.

Caution: If you are installing only one module, it needs to be installed in the lower half of the rack for greater balance and stability.

3. Secure the front of the module to the chassis using the U-brackets.
4. Attach the Z-brackets to secure the rear of the module to the back of the chassis (see Figure 9-5).
5. If you are installing the module into the upper set of rails, carefully raise the module into the upper set of rails (see Figure 9-1 for placement).
6. Secure the upper module to the chassis as described in steps 1 through 4.
7. After installing the module(s), go to Section 9.6 to complete installation.

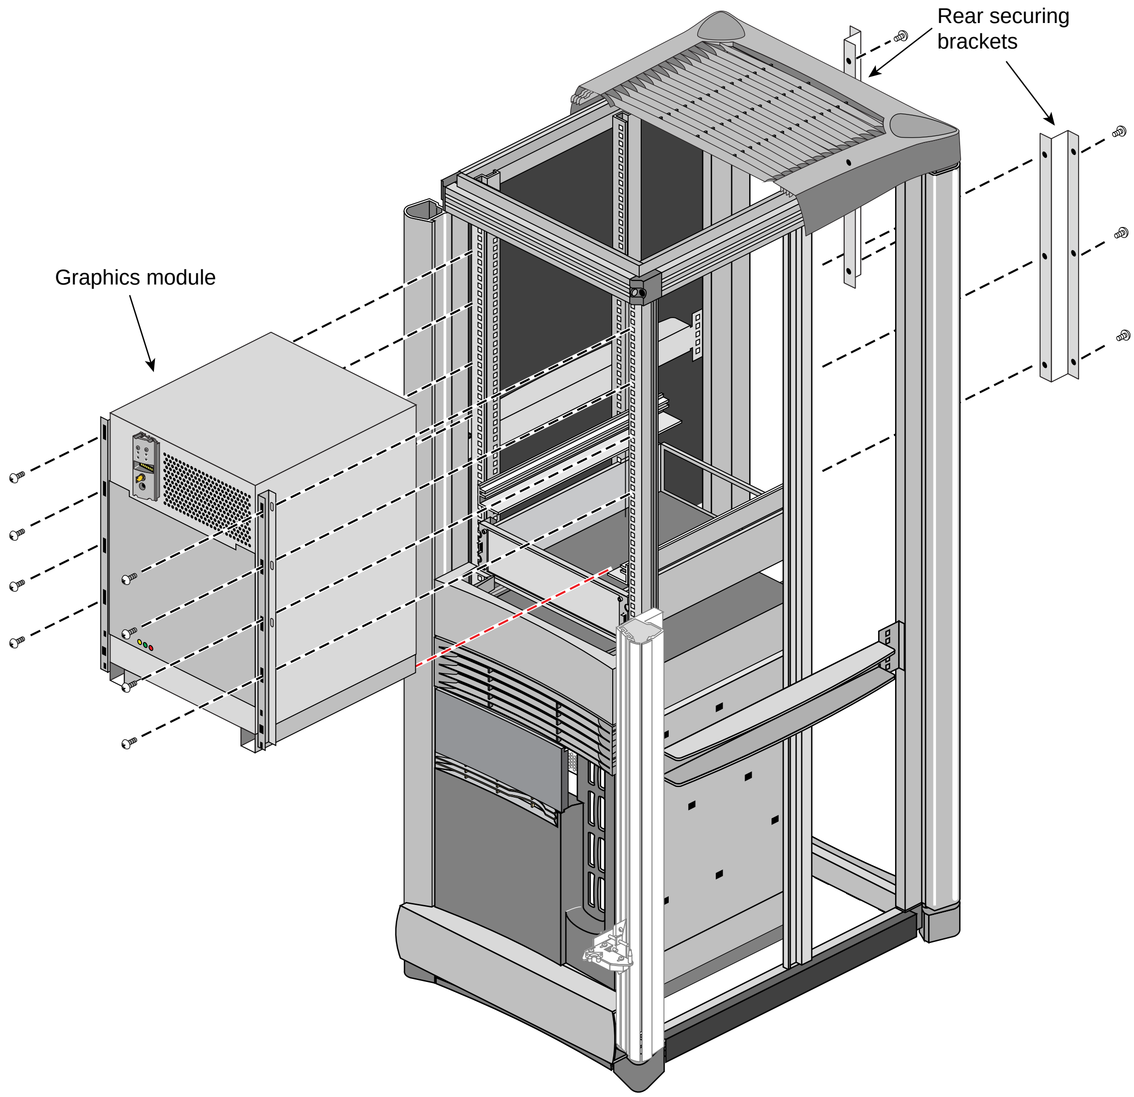

Figure 9-5 Graphics Module Installation

9.5 Installing a Deskside Module In a Standard 19-inch (Non-Origin2000) Rack

The Origin2000 deskside module can be installed into any standard 19-inch rack, however, the cable management hardware or multimodule System Controller is not provided with the Origin2000 rack. Nevertheless, it is possible to use CrayLink Interconnect cabling to expand these system configuration types.

To install a deskside module into a 19-inch rack, make sure that the site preparation requirements have been met (see Section 9.1 and the *Site Preparation Guide for Origin Family and Onyx2*, P / N 007-3452-xxx) and that the deskside module has been stripped of essential components for weight reduction (see Section 9.2).

CrayLink Interconnect capability across multiple modules installed in non-Origin2000 racks requires the *Crosstown* cable (see Figure 9-6). The Crosstown cable has a conduit wrapping that provides additional protection for the CrayLink Interconnect cable wiring.

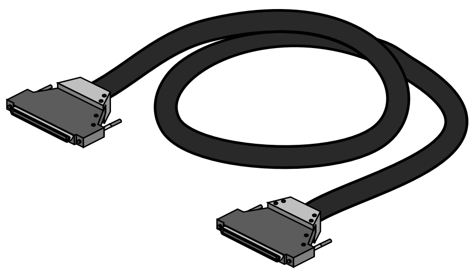

Figure 9-6 Crosstown Cable (With Conduit)

The modification kit for non-Origin2000 racks also includes the following major parts:

- exhaust and intake plenums
- non–Silicon Graphics cosmetic facade

Caution: Make sure that the railings on the racks are sufficiently spaced to allow enough room to hold the deskside module.

9.6 Completing Installation

After installing the module in the rack, follow the general procedures outlined in Table 6-1 or Table 6-2 in Chapter 6, “Installation” to complete the module installation.

Chapter 10

General Replacement Procedures

This chapter describes the installation and removal procedures for the field-replaceable units (FRUs) in the Origin2000 and Onyx2 deskside and rackmount systems.

10.1 General Information

Read the following subsections for information about safety and required tools.

10.1.1 Safety Information

Before you begin the replacement procedures, observe these precautions.

Warning: **This equipment uses electrical power internally that is hazardous if the equipment is improperly disassembled. The procedures in this manual require specific training and technical knowledge and are for use only by Silicon Graphics system support engineers or other Silicon Graphics-trained personnel.**

Warning: **Wait at least 2 minutes after you power off a rackmount system before you work on any part of the power supply or midplane/backplane. Because of the large amount of capacitance in the system, a significant amount of the operating voltage remains on the midplane/backplane for up to 2 minutes after the system is powered off.**

Caution: This equipment is extremely sensitive and susceptible to damage by electrostatic discharge (ESD). The buildup of electrical static potential on clothing and other materials can cause ESD.

Use proper ESD preventive measures (see Figure 10-1 and Figure 10-2 for built-in grounding points on a rack system). In addition, observe these precautions:

- Wear a properly grounded wrist strap when you connect and disconnect peripherals.
- Be sure that you and all the electrical equipment that you handle are at ground potential to avoid damage from ESD. Do not rely on the power source ground: the ground is lost when the system is disconnected from the power source.

- Keep boards in their antistatic bags until you are properly grounded to the chassis ground with a ground strap.
- Do not use an ohmmeter or digital voltmeter on a board.

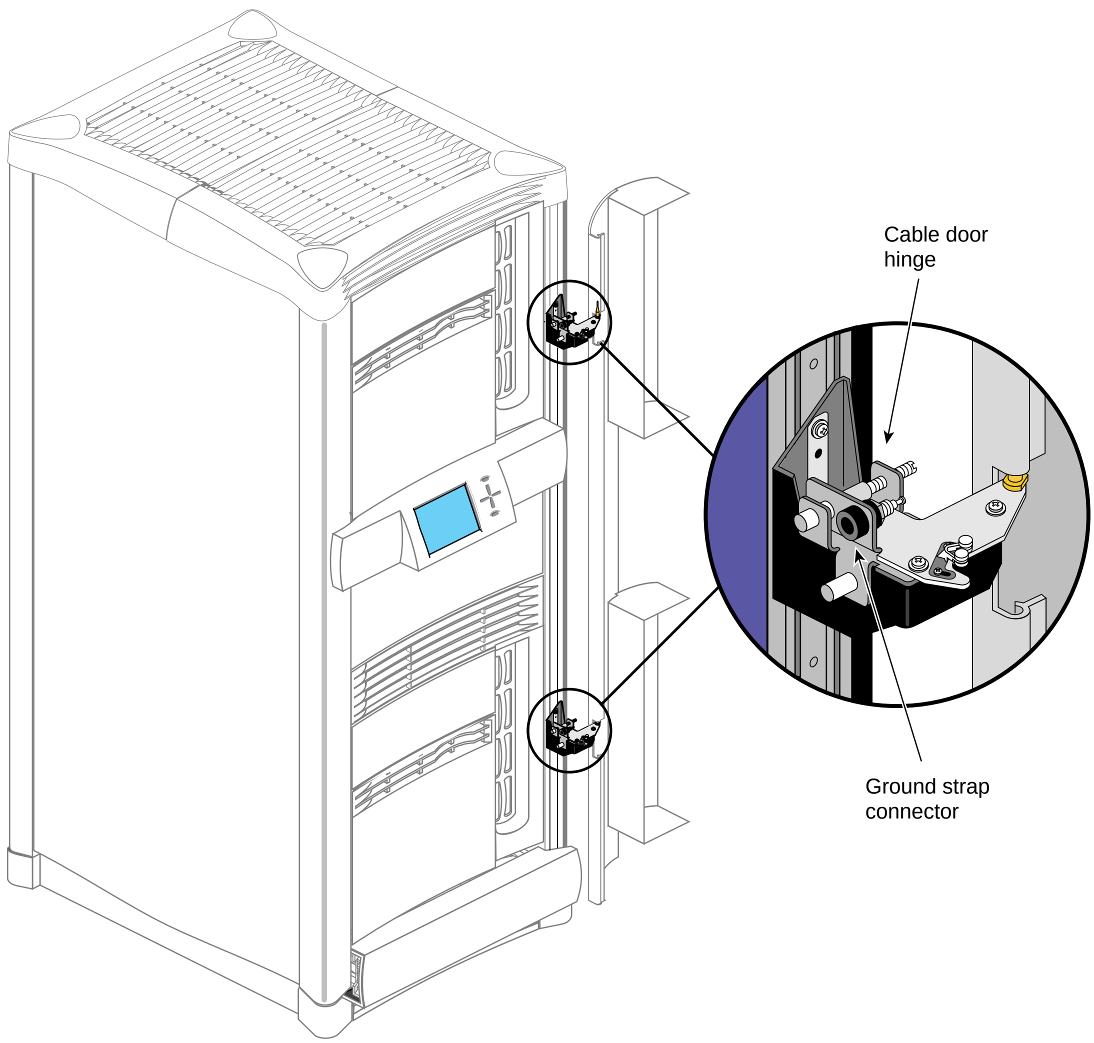

Figure 10-1 Grounding Points in Front of Rack Chassis

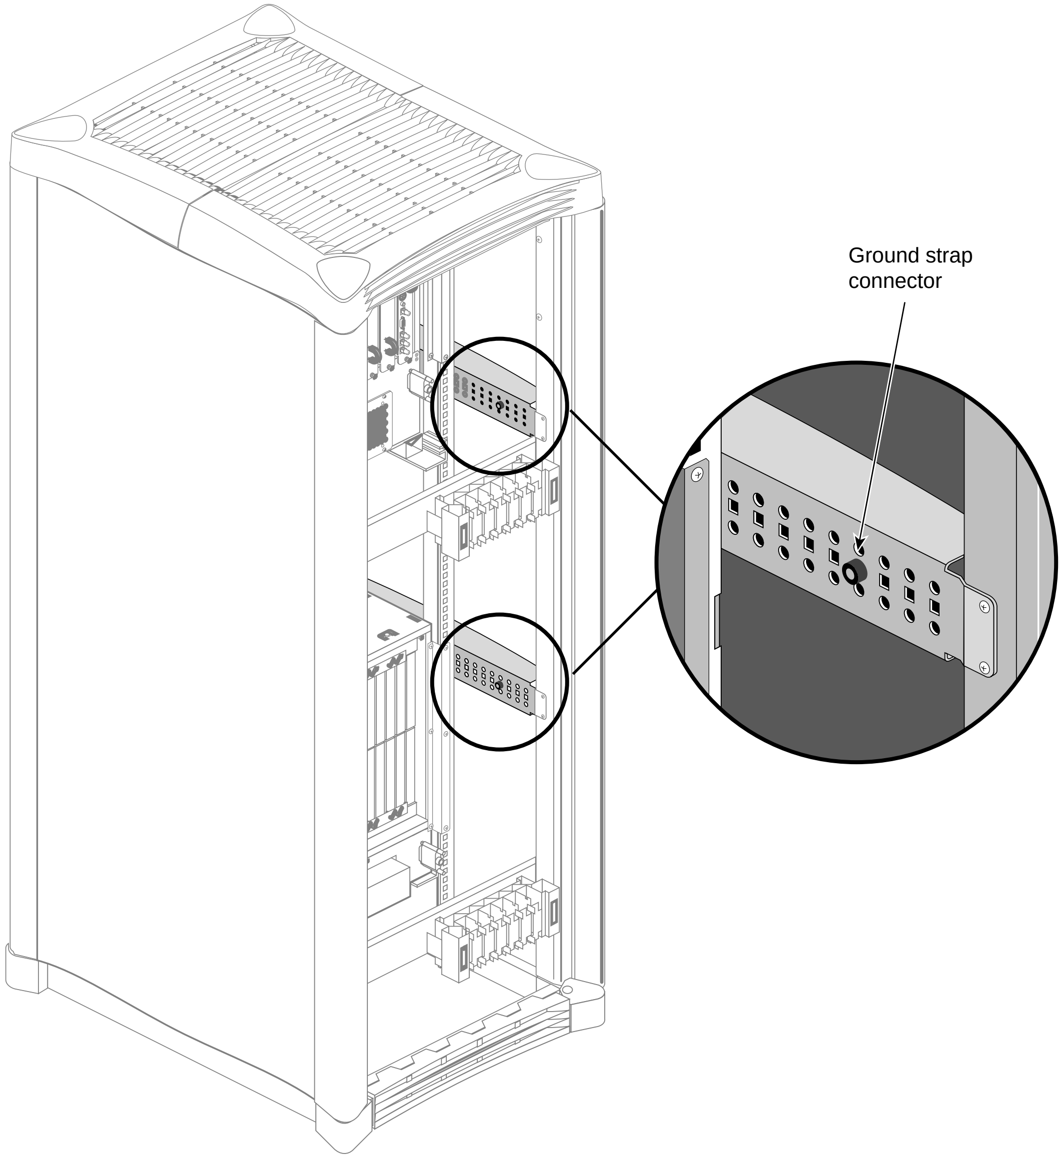

Figure 10-2 Grounding Points in Rear of Rack Chassis

10.1.2 Required Tools

Here is a minimum list of required tools

- #2 Phillips screwdriver (14-in. long)
- #1 Phillips screwdriver
- 7/64-in. hex driver (to loosen compression connectors)
- flatblade screwdriver

10.2 General Procedures

This section describes procedures that are common to many maintenance activities. Figure 10-3 and Figure 10-4 provide an exploded view of the doors, panels, and major hardware components of the Origin2000 server systems. Figure 10-6 and Figure 10-7 provide exploded views of the Onyx2 graphics systems. Table 10-1 lists common Origin2000 and Onyx2 FRUs and their part numbers.

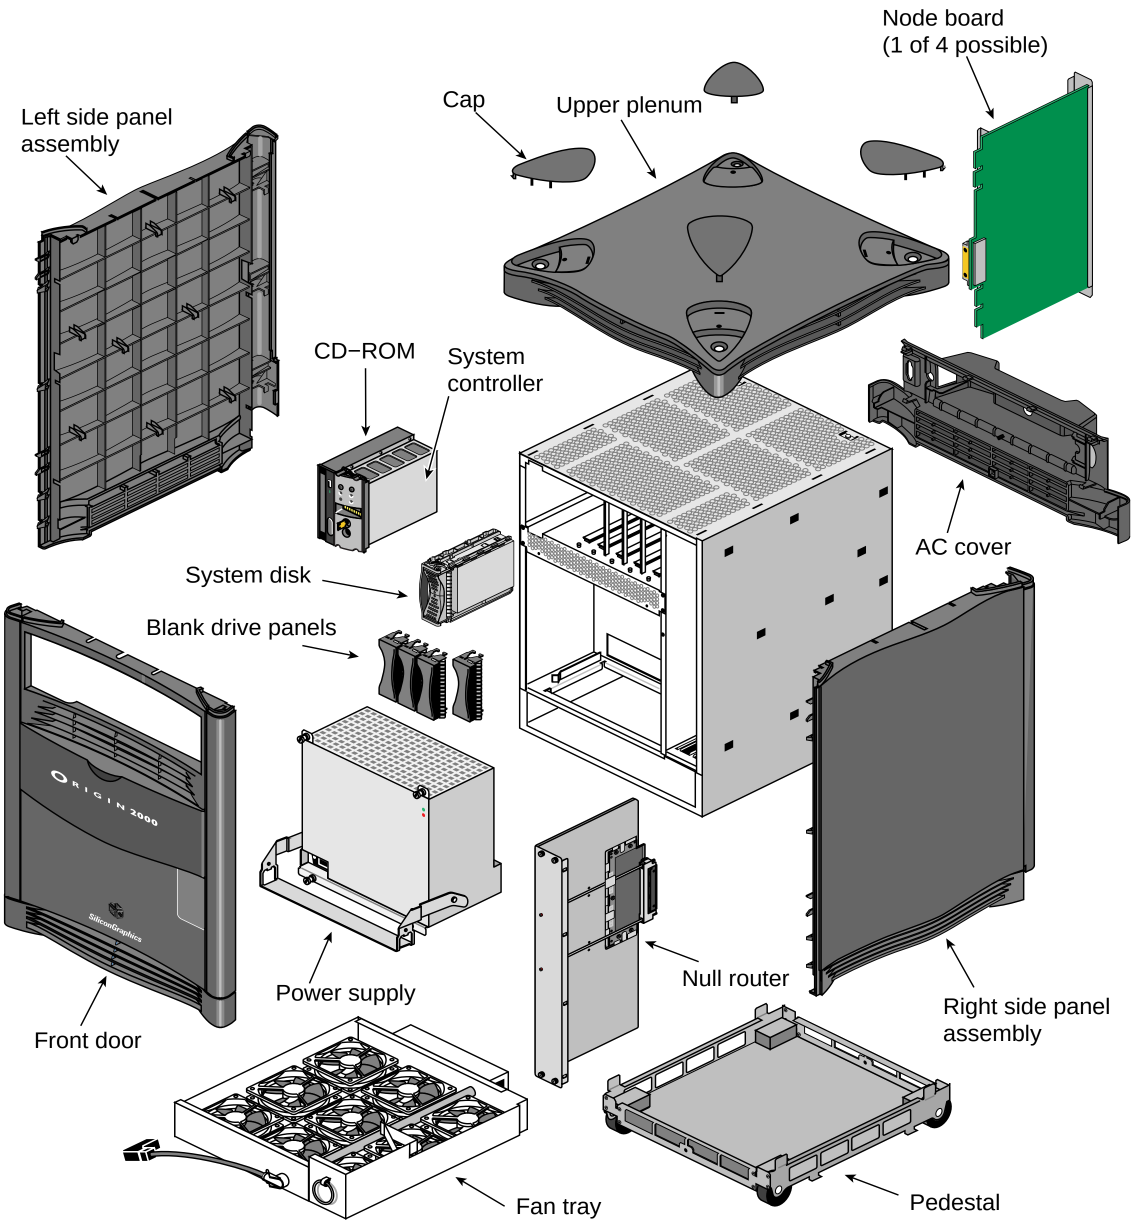

Figure 10-3 Origin2000 Deskside System Components (Exploded View)

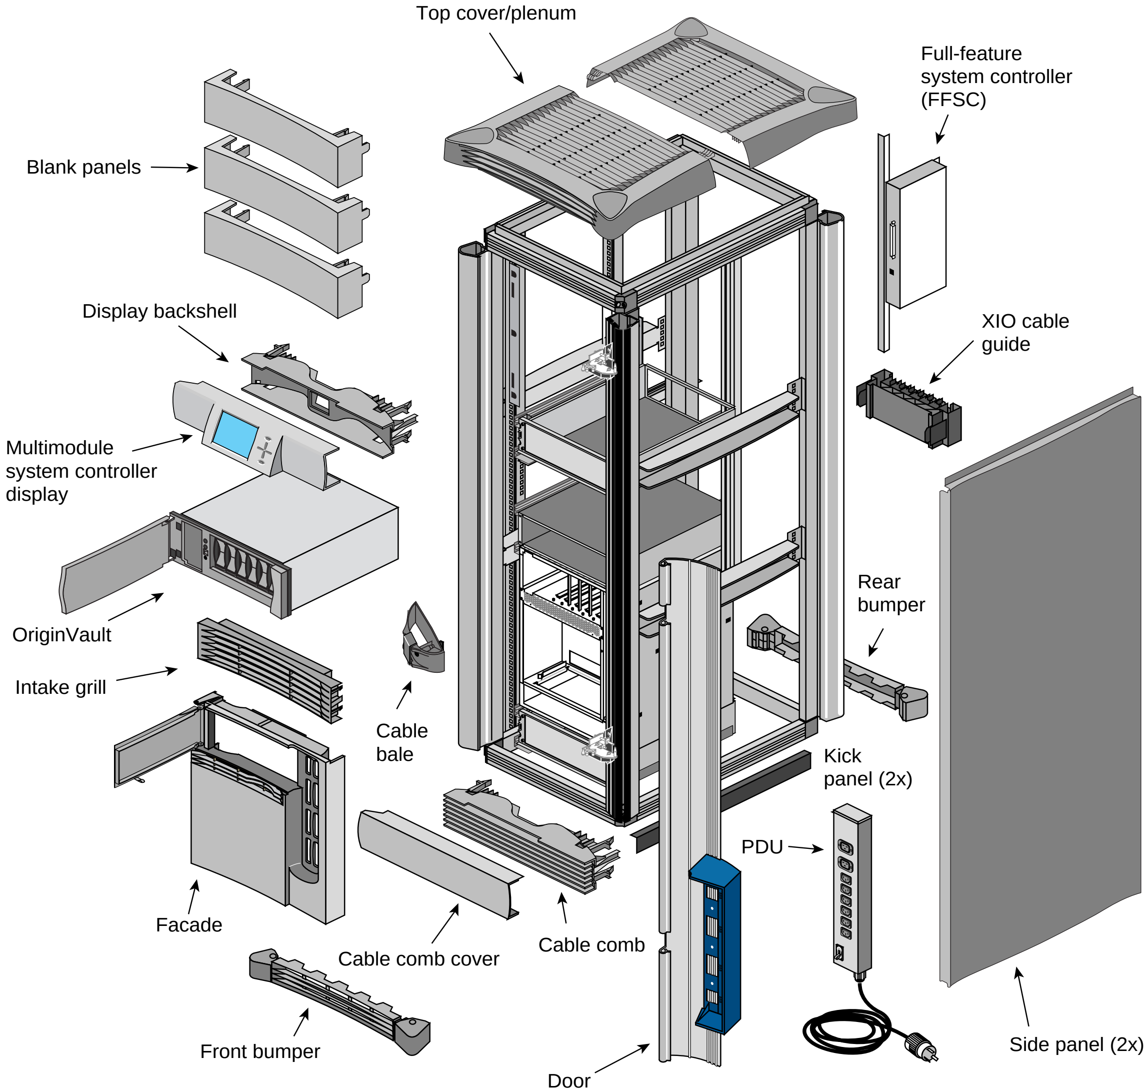

Figure 10-4 Origin2000 Rackmount Major Components Exploded View (Part 1)

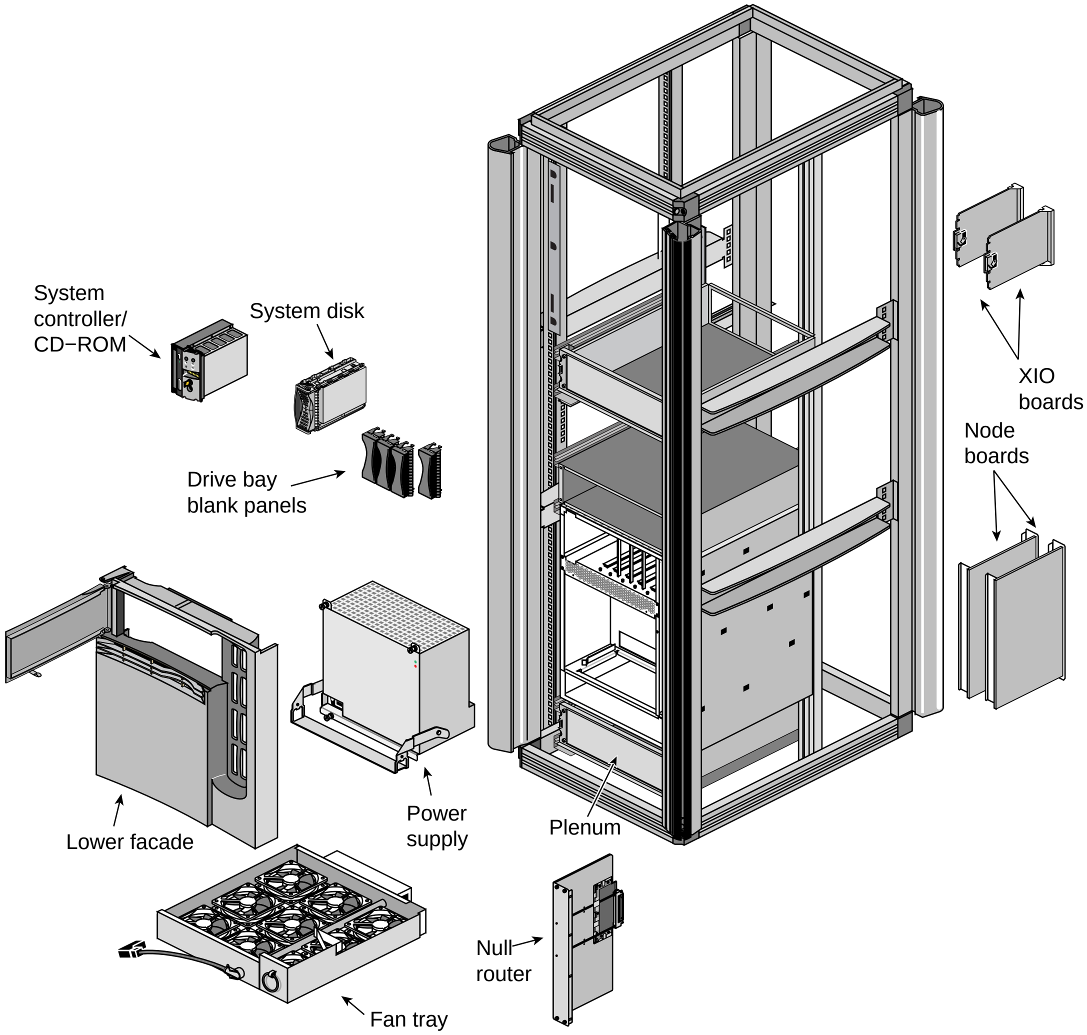

Figure 10-5 Rackmount Major Components Exploded View (Part II)

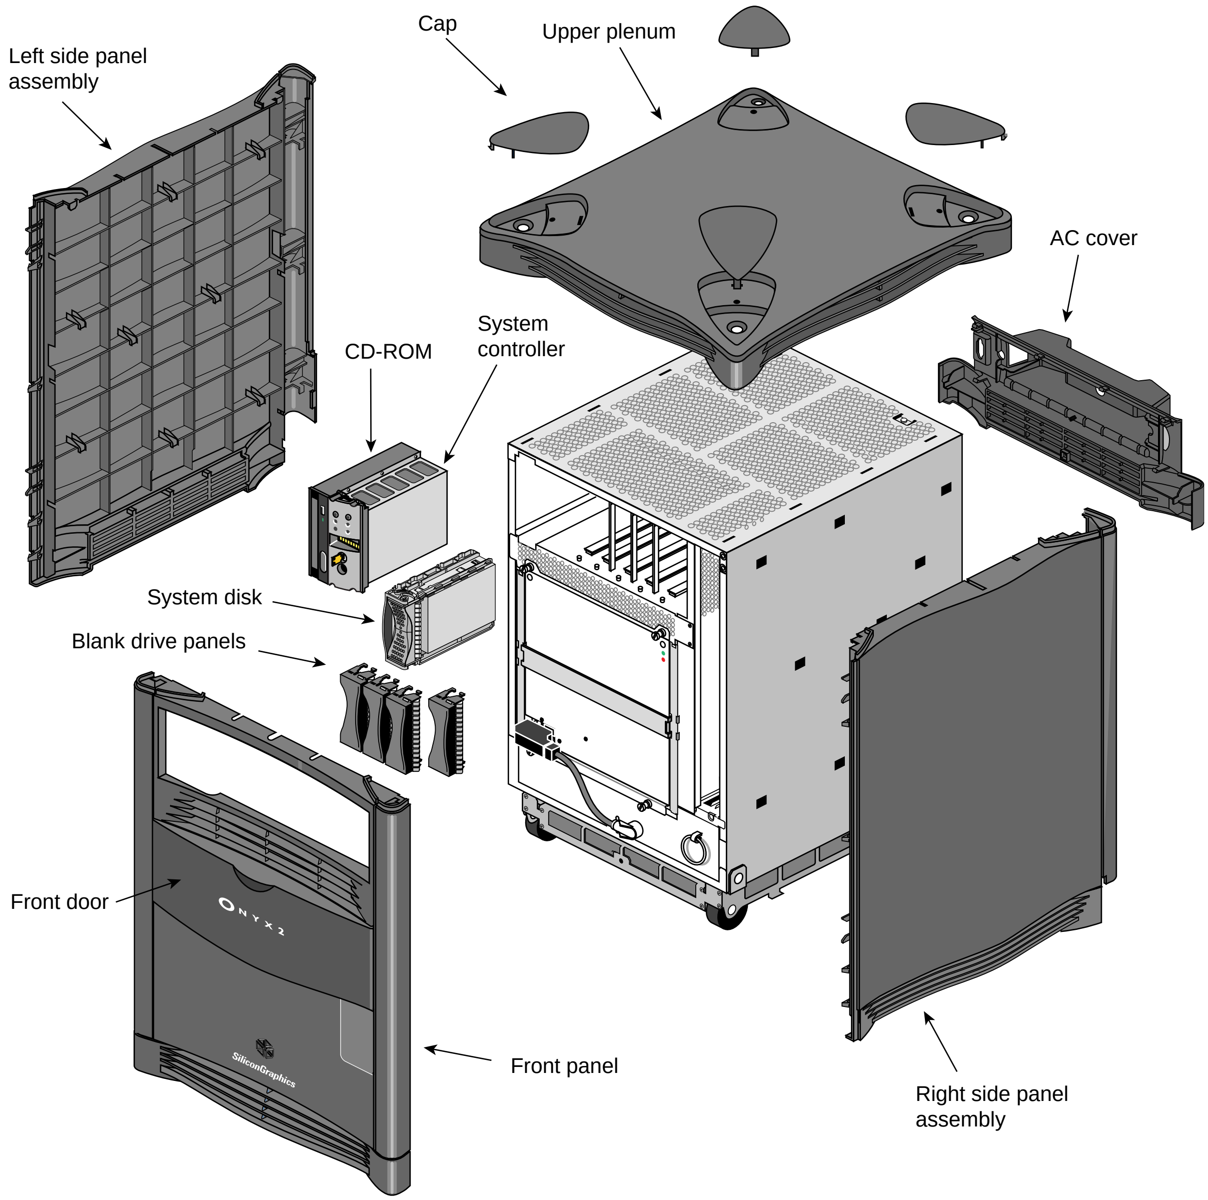

Figure 10-6 Onyx2 Deskside Exploded View

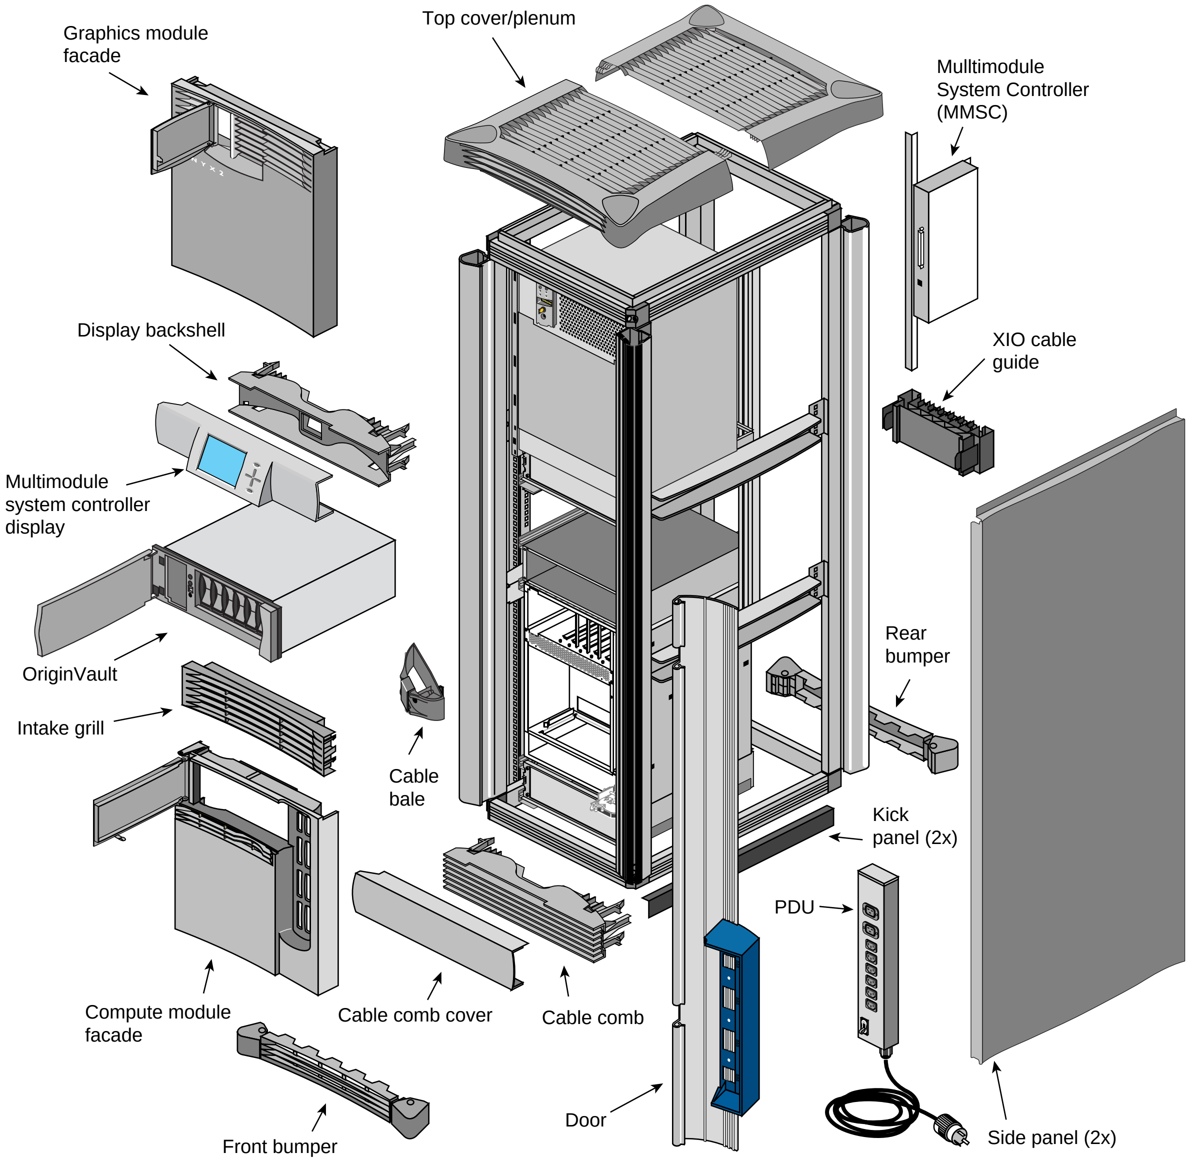

Figure 10-7 Onyx2 Rackmount Exploded View

Table 10-1 Selected Origin2000 and Onyx2 FRU Part Numbers

Description	**Part number**
Module System Controller	030-0859-xxx
Multimodule System Controller (“Single Board Computer”)	041-0229-xxx
Multimodule System Controller display	013-1714-xxx
Multimodule System Controller display cable	018-0606-xxx

Table 10-1 (continued) Selected Origin2000 and Onyx2 FRU Part Numbers

Description	Part number
IP27 Node Board, 1 CPU, 1MB SCache	013-1346-xxx
IP27 Node Board, 2 CPU, 1MB SCache	013-1347-xxx
IP27 Node Board, 1 CPU, 4MB SCache	013-1830-xxx
IP27 Node Board, 2 CPU, 4MB SCache	013-1831-xxx
32 MB DIMM	030-0759-xxx
64 MB DIMM	013-1372-xxx
256 MB DIMM	030-0775-xxx
16 MB directory DIMM	030-0783-xxx
64 MB directory DIMM	030-0781-xxx
Midplane	013-1547-xxx
BaseIO (IO6)	013-1631-xxx
Fan tray assembly	013-1589-xxx
Power supply	060-0005-xxx
Null Router	030-0842-xxx
Star Router (8p router)	030-0841-xxx
Standard router (128p router)	030-0845-xxx
CrayLink cable, 58”	018-0564-xxx
CrayLink cable 84”	018-0693-xxx
CrayLink cable 95”	018-0694-xxx
CrayLink cable 108”	018-0568-xxx

10.2.1 Powering Off a Deskside System

Use this procedure to power off the system:

1. Verify that the desired filesystems are backed up and make sure that all users are off the target system.
2. Become superuser; then shut off the system software as follows:

   # shutdown -y g0

3. Turn the Module System Controller (MSC) key switch to the Off/Standby position (see Figure 10-8).
4. Power off the module. The switch is located in the rear of the chassis, near the AC power cord receptacle (see Figure 10-9). Disconnect the system from the power source.

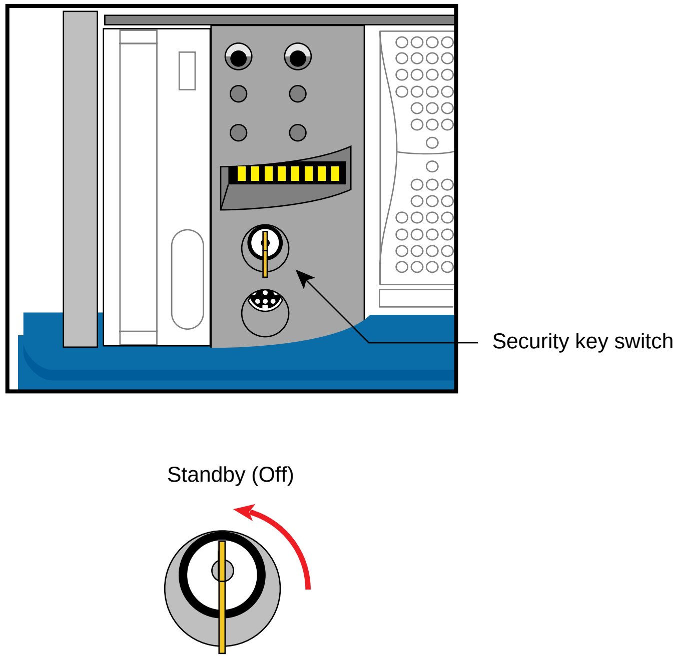

Figure 10-8 Turning Off the MSC

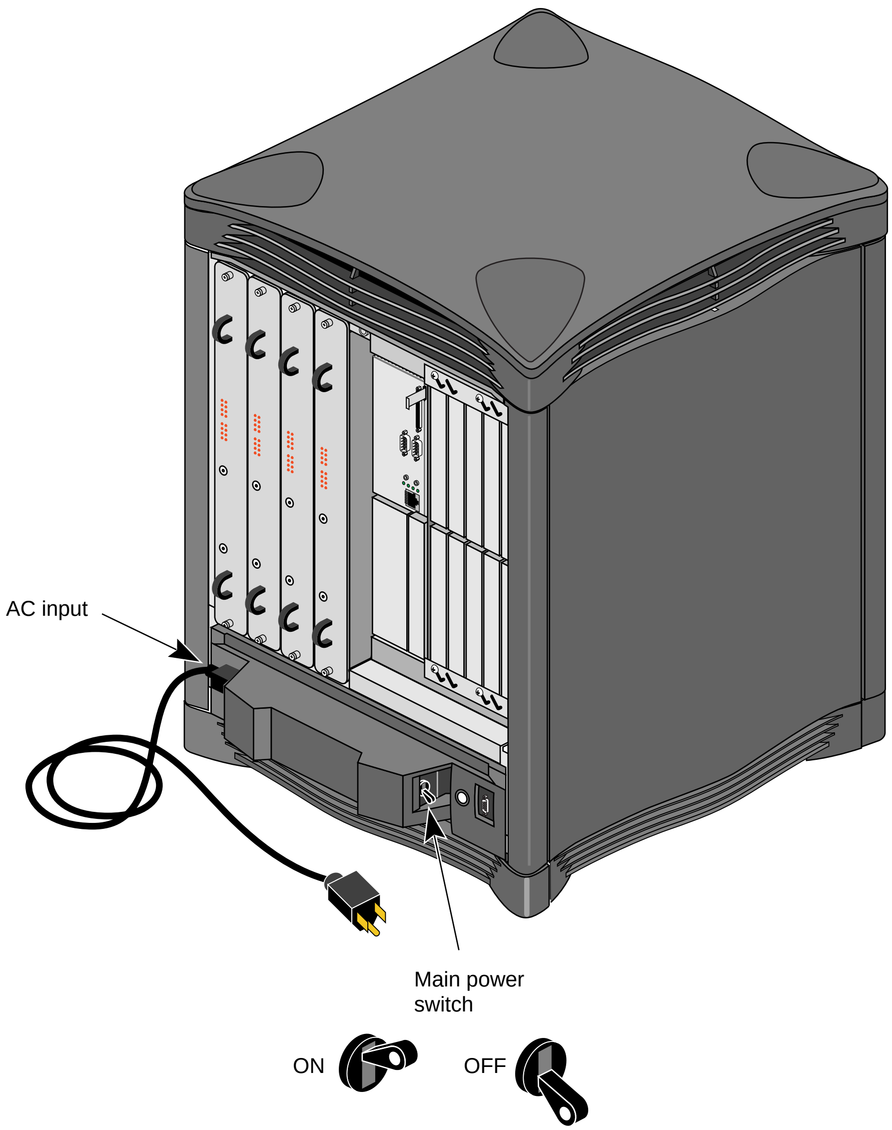

Figure 10-9 Powering Off a Deskside System

10.2.2 Powering Off a Rackmount System

You can either power off an individual module or power off the entire rack configuration using the multimodule System Controller (MMSC). When you power off an individual module, the other module(s) in the configuration continue to function normally.

10.2.2.1 Powering Off an Individual Module System

Use this procedures to power off a module in a rack:

1. Verify that the desired filesystems are backed up and make sure that all users are off the target system.
2. Become superuser; then shut down the system software as follows:

   ```
   # shutdown -y g0
   ```
3. Turn the MSC key switch to the Off/Standby position (see Figure 10-8).
4. Power off the module. The switch is located in the rear of the chassis, near the AC power cord receptacle (see Figure 10-10). Disconnect the system from the power source.

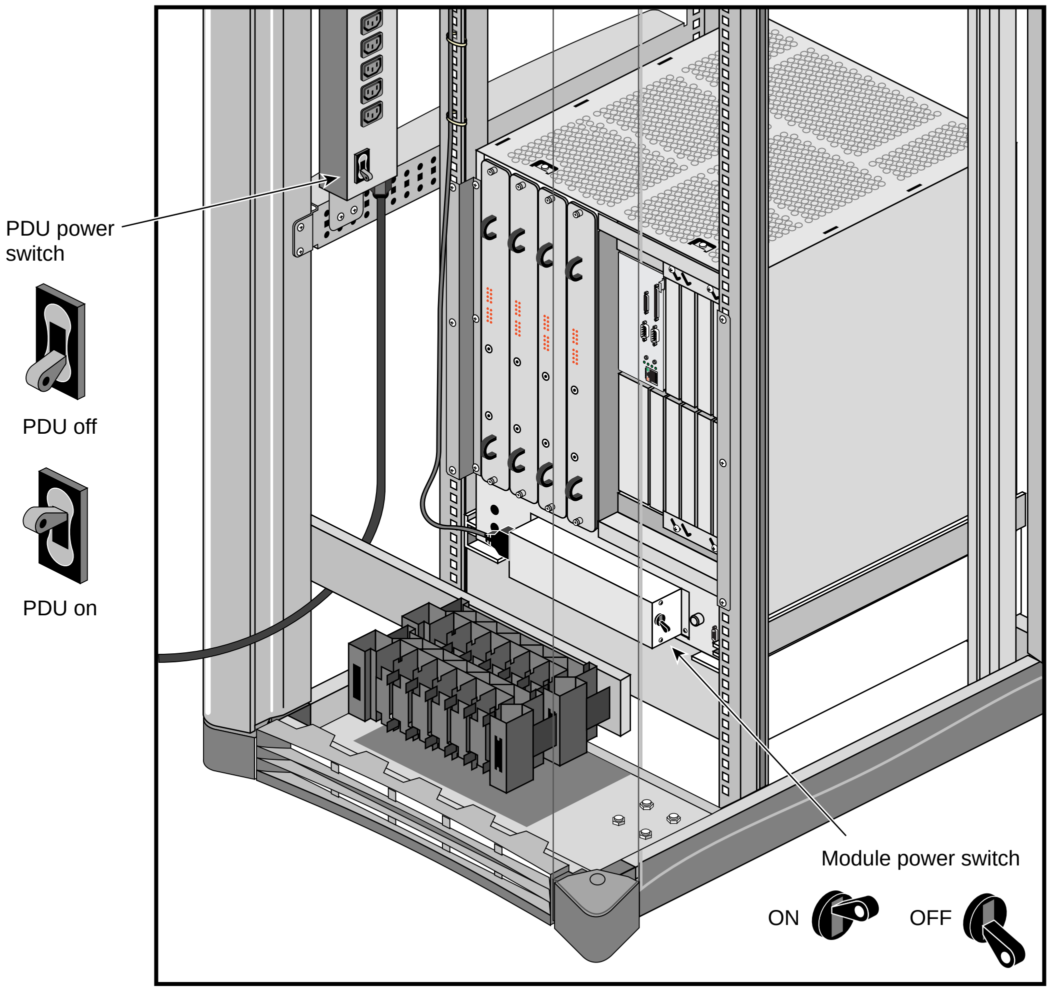

Figure 10-10 Powering Off a Rackmount Module

10.2.2.2 Powering Off a Single-rack Configuration

Follow these instructions to power off an entire rack configuration through the multimodule System Controller display (see Figure 10-11).

1. Verify that the desired filesystems are backed up and make sure that all users are off the target system.
2. Become superuser; then shut down the system software as follows:

   ```
   # shutdown -y g0
   ```
3. Turn the MSC keyswitch(es) to Standby.
4. At the MMSC display, select the Focus menu using the direction keys; then press the Execute key.

   Make sure that the All option is highlighted.
5. Using the direction keys, go to the “Action” menu on the display.
6. Select the Power Off option and press the Execute key.

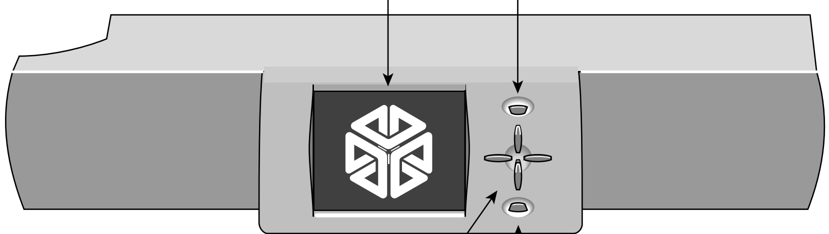

Figure 10-11 Powering Off the Multimodule System Controller and Display

7. Turn off power to the power distribution unit (PDU); see Figure 10-10.

10.2.2.3 Powering Off a Multimodule Rack Configuration

Use this procedure to power off an entire multirack configuration.

1. Power off the configuration using the multimodule display (see Section 10.2.2.2).
2. Power off the PDUs on the individual racks. It is recommended that you start from the rightmost rack and then continue with the rack to the left. This helps to ensure a cleaner and more efficient shutdown process.
3. Turn off the PDU on the last rack (see Figure 10-12).

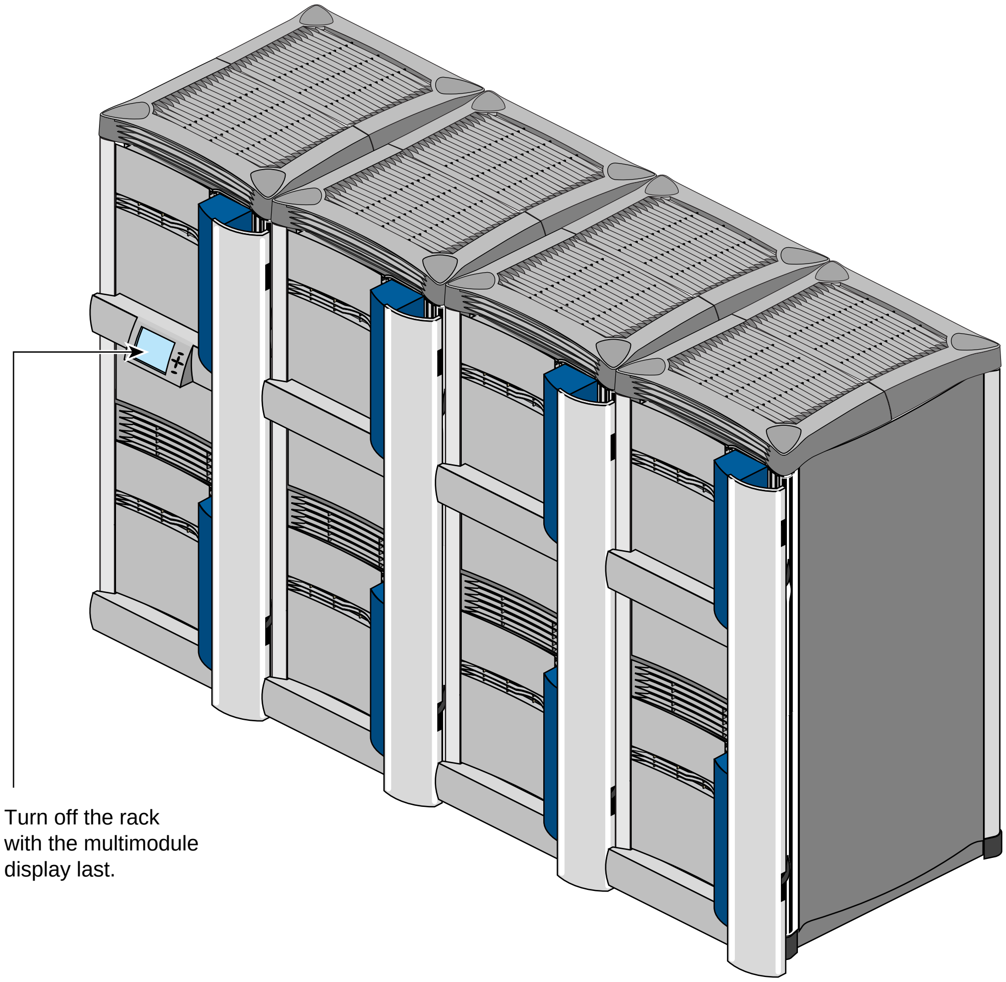

Figure 10-12 Turning Off a Multirack Configuration

10.2.3 Opening and Closing the Drive Door (Deskside)

To open the drive door on a deskside system:

1. Open the drive door to expose the drive by grasping the top of the door and sliding it down (see Figure 10-13).

   Note: The drive door is spring-loaded.

2. To close the drive door, push down on the drive door. The door should pop up automatically.

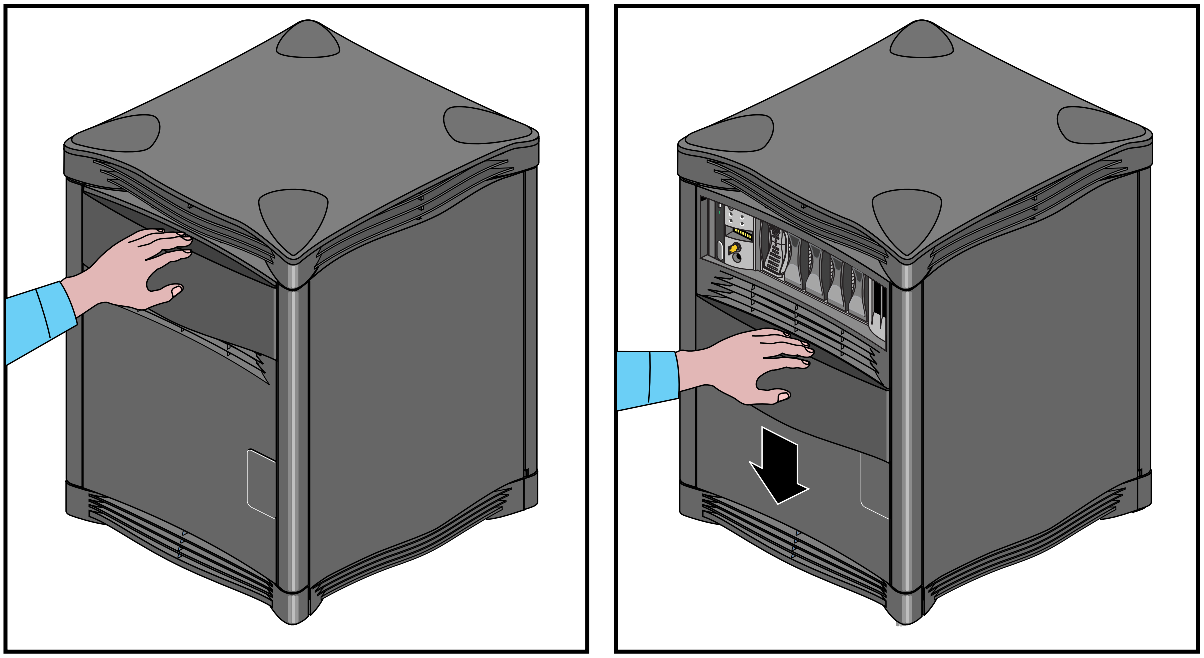

Figure 10-13 Opening the Drive Door on a Deskside System

10.2.4 Opening the Drive Door (Rackmount)

To open the drive door on a rackmount system, swing open the drive door as shown in Figure 10-14). When you close the door, make sure to push the door all the way in to engage the plastic tab on the bottom of the drive door.

Note: The door should normally be in the closed position to prevent dust and other contaminants from entering the drives and System Controller.

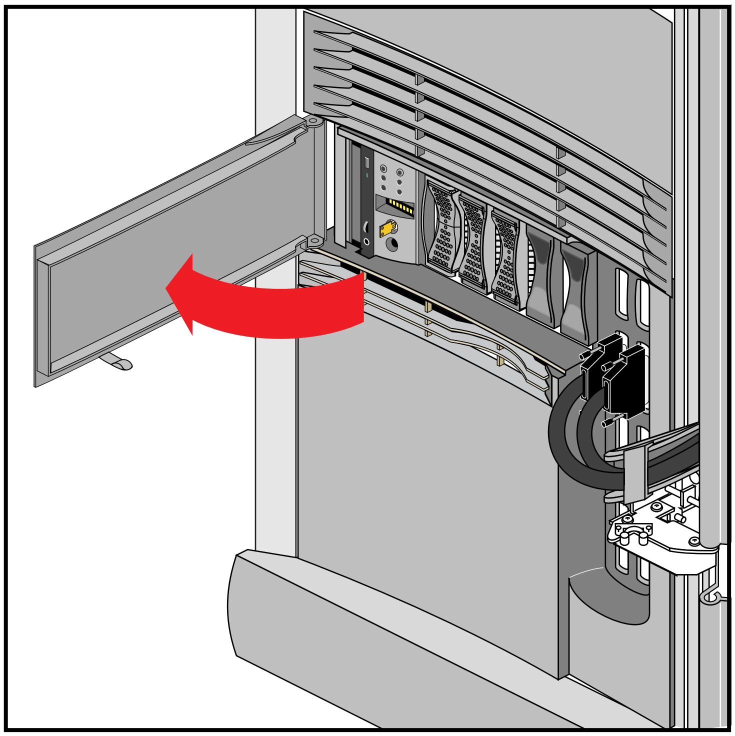

Figure 10-14 Opening the Drive Door on a Rackmount System

10.2.5 Opening the Cable Cover Door (Rackmount Systems Only)

The cable cover door (see Figure 10-15) provides shielding for the interconnection fabric cabling on the rackmount chassis and between side-by-side rackmount systems. The door can also be removed when you install CrayLink cabling.

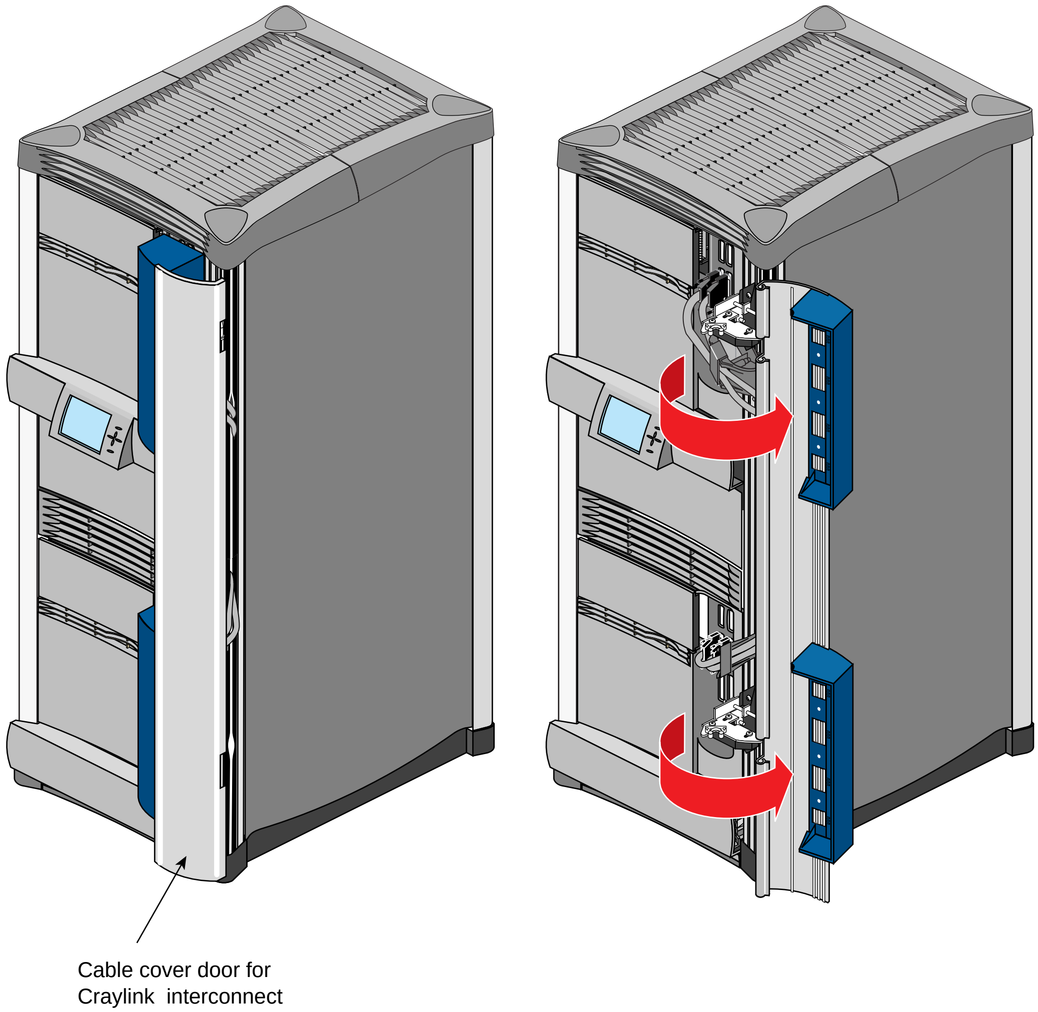

Figure 10-15 Opening the Cable Cover Door

10.2.6 Removing the Cable Cover Door

Use this procedure to remove the door.

1. Swing the door open.
2. Remove the ground strap near the top hinge (see Figure 10-16)
3. Lift it up and pull it away.

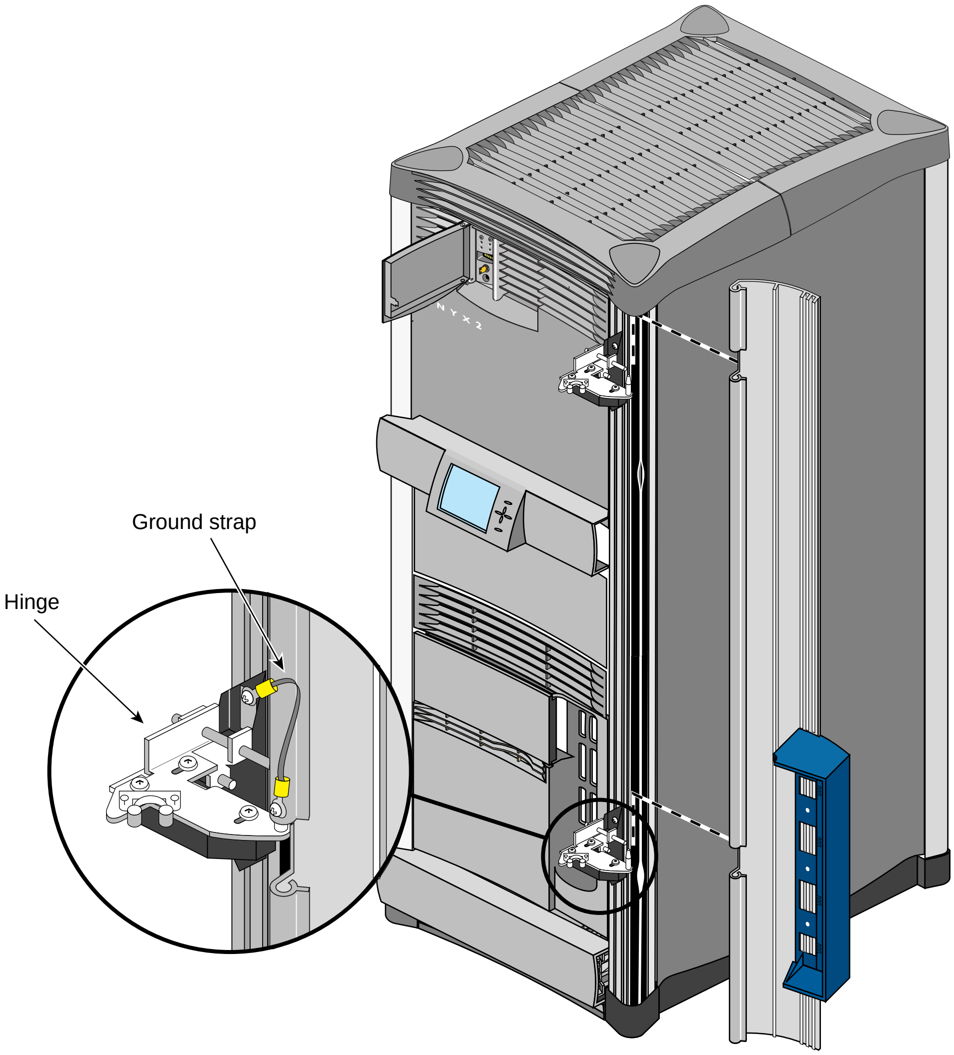

Figure 10-16 Removing the Cable Cover Door

10.2.7 Removing the Facade

Use this procedure to remove the facade (or front door) from a deskside (see Figure 10-17 or rackmount system (see Figure 10-18).

1. Remove or loosen the screw that holds the facade in place (see Figure 10-17 or Figure 10-18).
2. Push down on top of the door frame to release the door and pull it out.
3. Reverse these steps to install the front door.

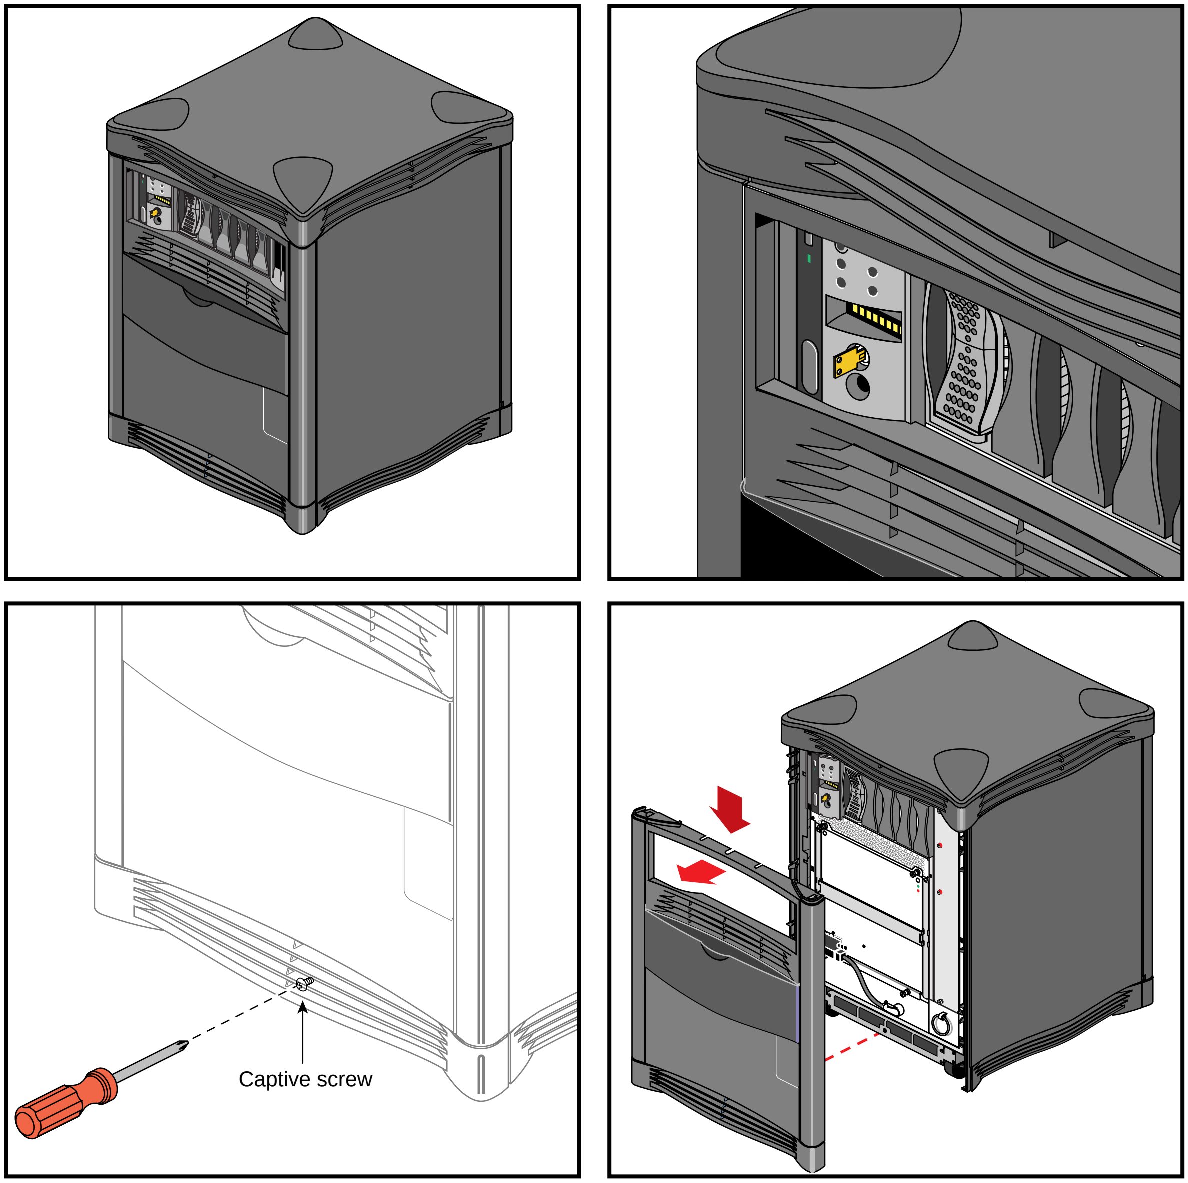

Figure 10-17 Removing the Facade (Deskside)

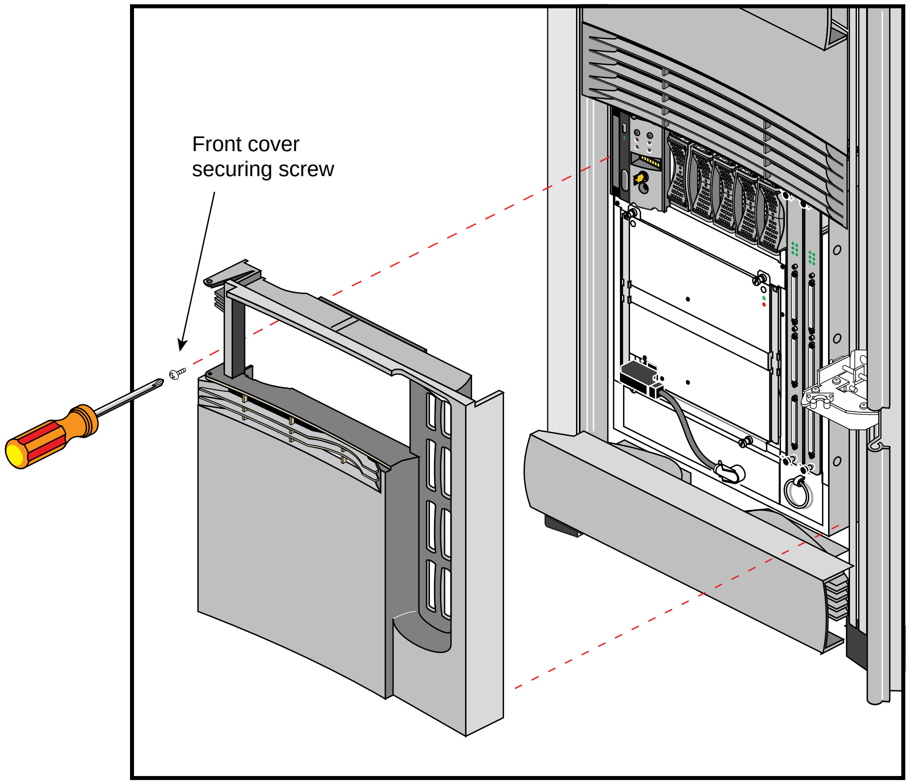

Figure 10-18 Removing the Facade (Rackmount)

10.2.8 Removing the Deskside Top Panel (Tophat)

Use this procedure to remove the deskside top panel (tophat):

1. Insert a flathead screwdriver through the plenum and underneath one of the four caps on the top panel.
2. Locate the securing latch that is located near the rear of the cap and push the latch forward (see Figure 10-19).
3. Lift up the cap.
4. Repeat Steps 1 through 3 for each of the caps.
5. Use a #2 Phillips screwdriver to remove the four screws that secure the top panel to the chassis (see Figure 10-19).
6. Lift off the top panel.

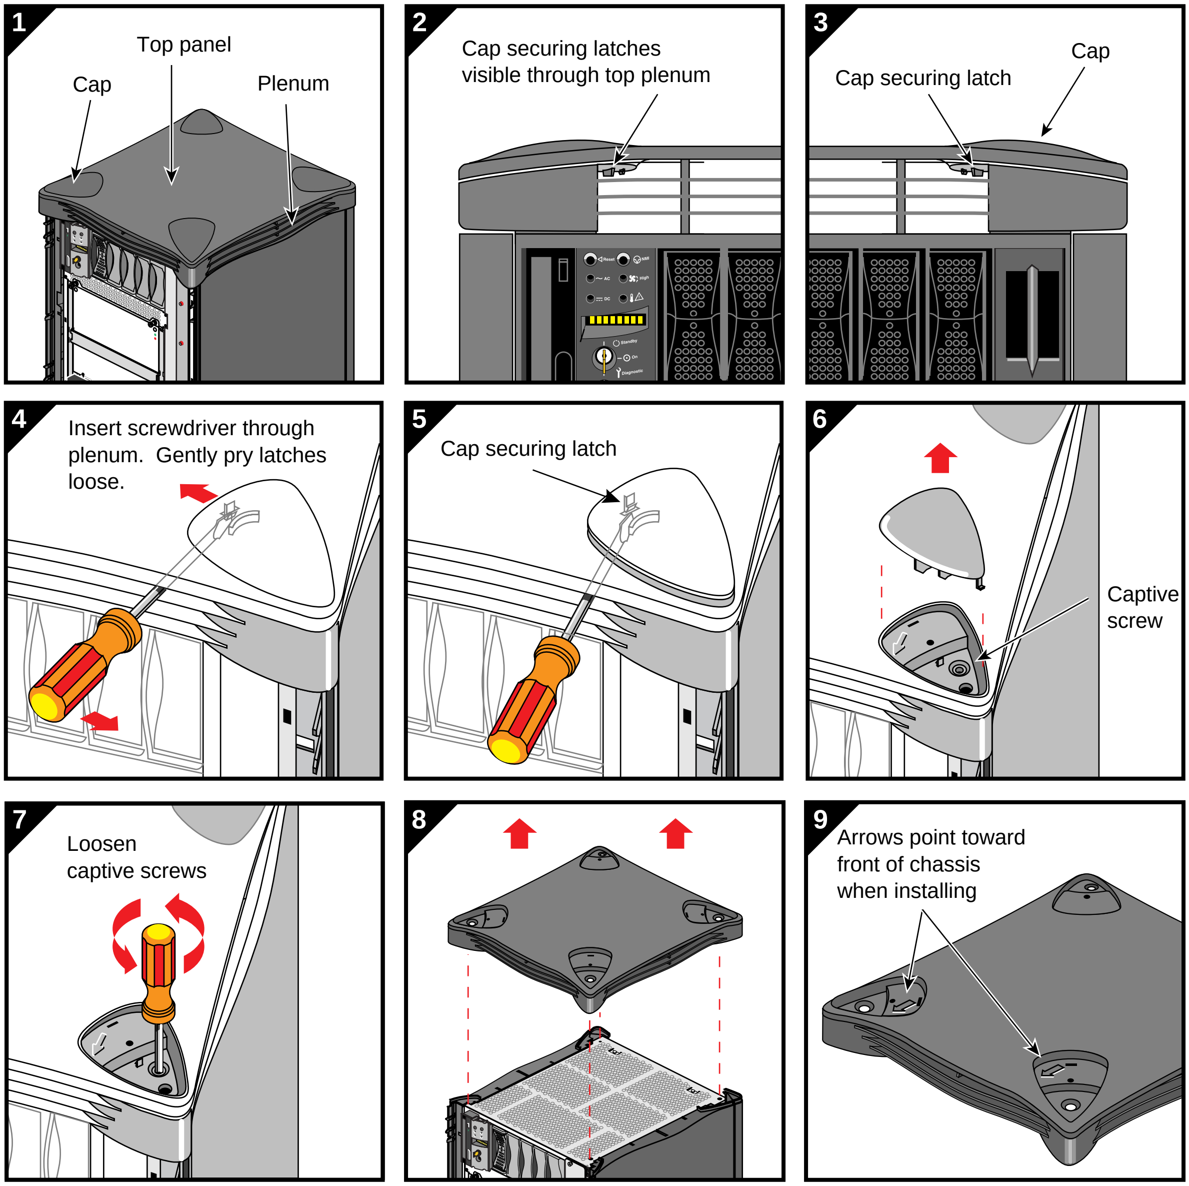

Figure 10-19 Removing the Top Panel

10.2.9 Removing the Deskside Side Panels

Use this procedure to remove the side panels from a deskside system:

1. Remove the top panel as described in Section 10.2.8.
2. Grasp the side panel by the bottom, pull it up, and then lift it out (see Figure 10-20).
3. Reverse the steps to reinstall the panel.

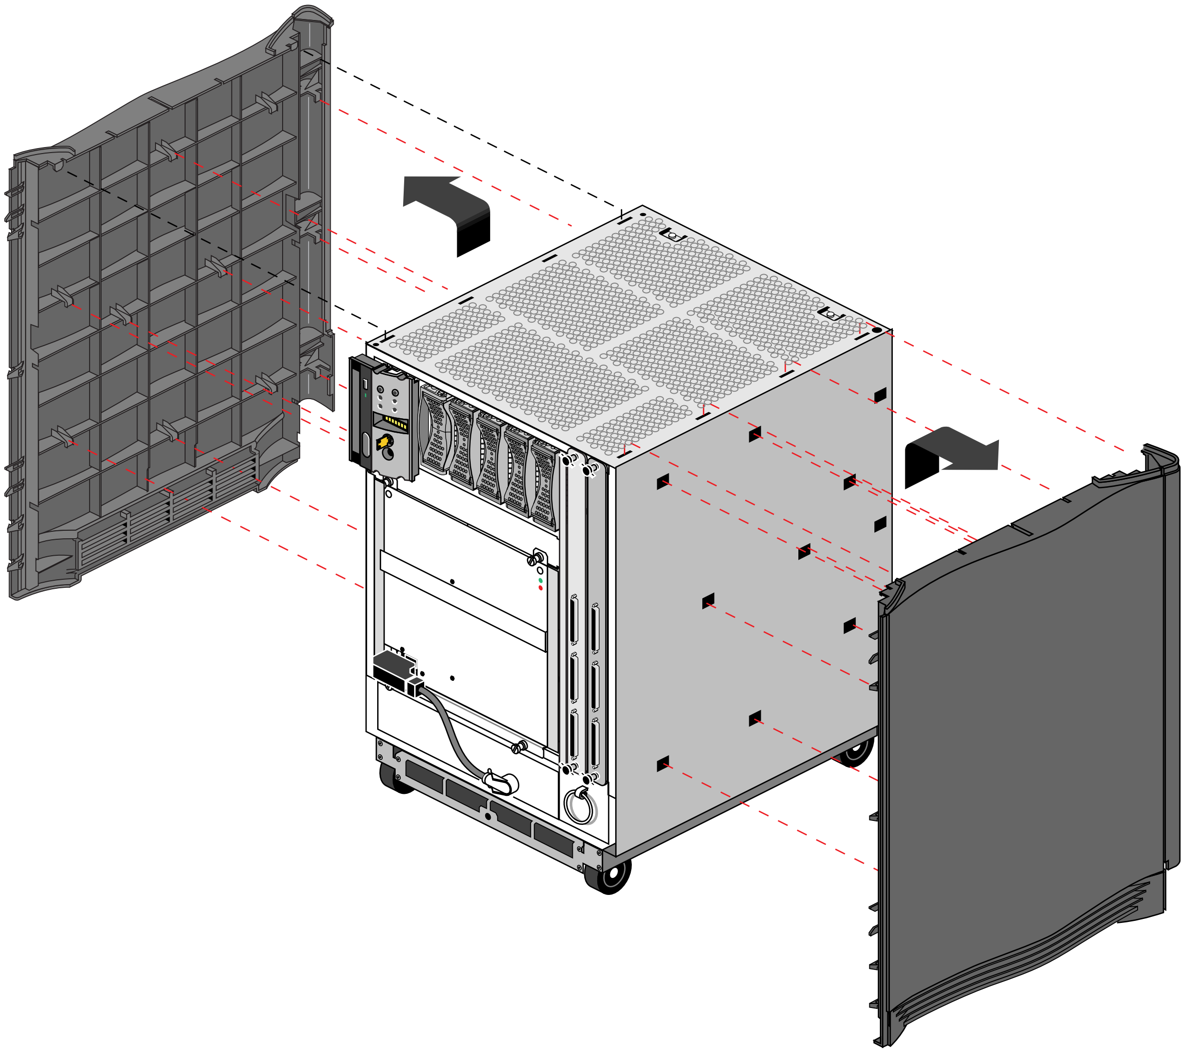

Figure 10-20 Removing the Side Panels

10.2.10 Removing the Deskside Bottom Plenum and Pedestal Assembly

Use this procedure to remove the deskside pedestal assembly:

1. Remove the top panel as described in Section 10.2.8.
2. Remove the side panels as described in Section 10.2.9.
3. Remove the four screws that secure the pedestal assembly at the sides (see Figure 10-21).
4. Reverse these steps to reinstall the pedestal assembly.

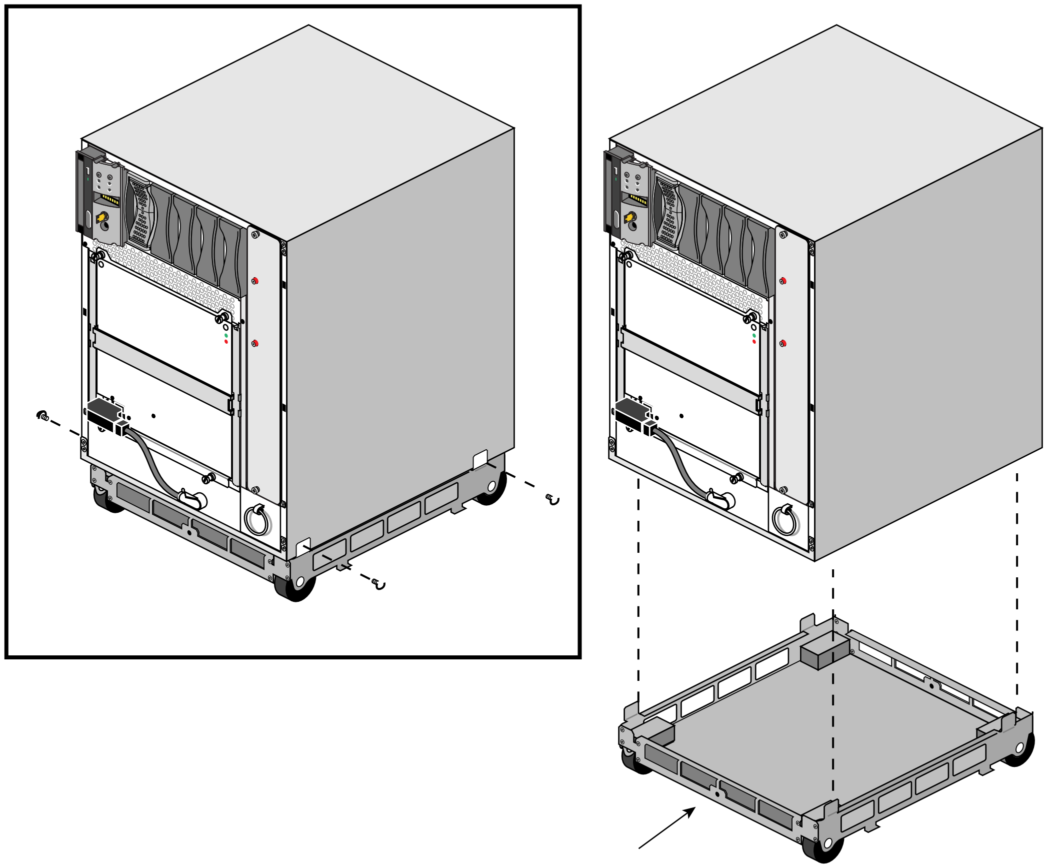

Figure 10-21 Removing the Bottom Plenum and Pedestal Assembly

10.2.11 Removing the Rackmount Side Panels

Use this procedure to remove the side panels from a rackmount chassis:

1. Remove one-half of top cover / plenum from the top of the chassis by pulling it forward.
2. Remove the other half of the top cover / plenum.
3. Grasp the top edge of the side panel and pull it down. (The panel is friction-fitted into the chassis.)
4. Reverse these steps to reinstall the side panel.

   Note: To replace the side panel, align it and then push it into place. If necessary, gently flex the panel to ensure alignment.

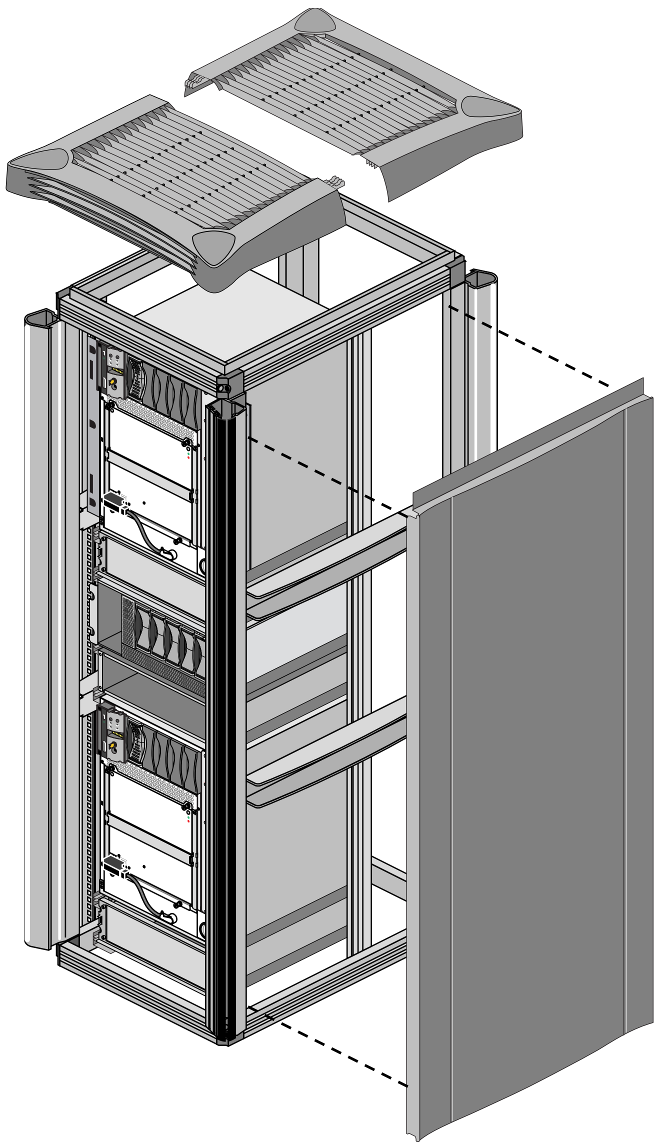

Figure 10-22 Removing the Rackmount Side Panel

10.3 Specific Procedures

The following section provides instructions for replacing the field-replaceable units (FRUs). To replace an FRU, first use Figure 10-3 and Figure 10-4 to identify the appropriate unit and its position in the chassis. Then proceed to the appropriate section and perform the steps.

10.3.1 Removing and Installing a Disk Drive Module

Disk drive modules are aligned vertically at the front of the chassis, as shown in Figure 10-23. In the *server module and systems*, note that the leftmost disk drive—the system drive—is oriented upside down relative to the other drives. All the drives in an Onyx2 deskside system have the same orientation.

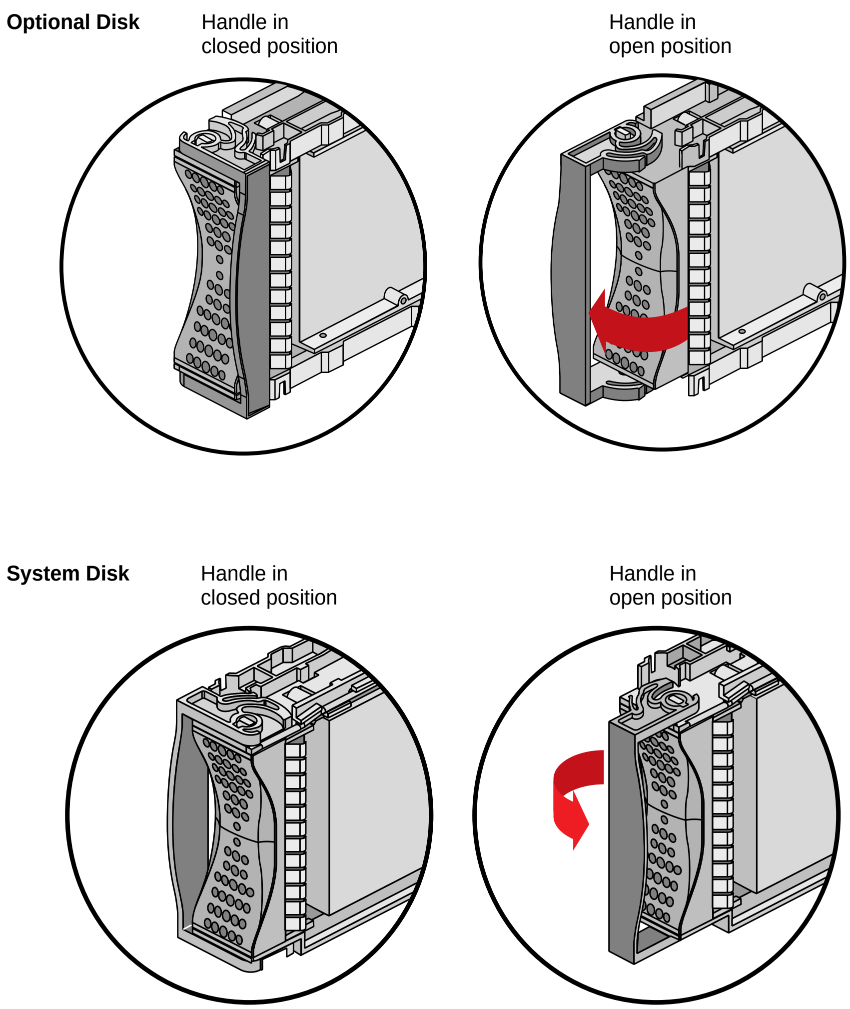

Figure 10-23 Disk Drive Unit Module

To *remove* a disk drive module, follow these steps:

1. To remove a disk drive module, snap the handle to the right or left (depending on the drive orientation) to the open position. The handle is centered, as shown in Figure 10-23.
2. Pull the disk drive module straight out (see Figure 10-24).

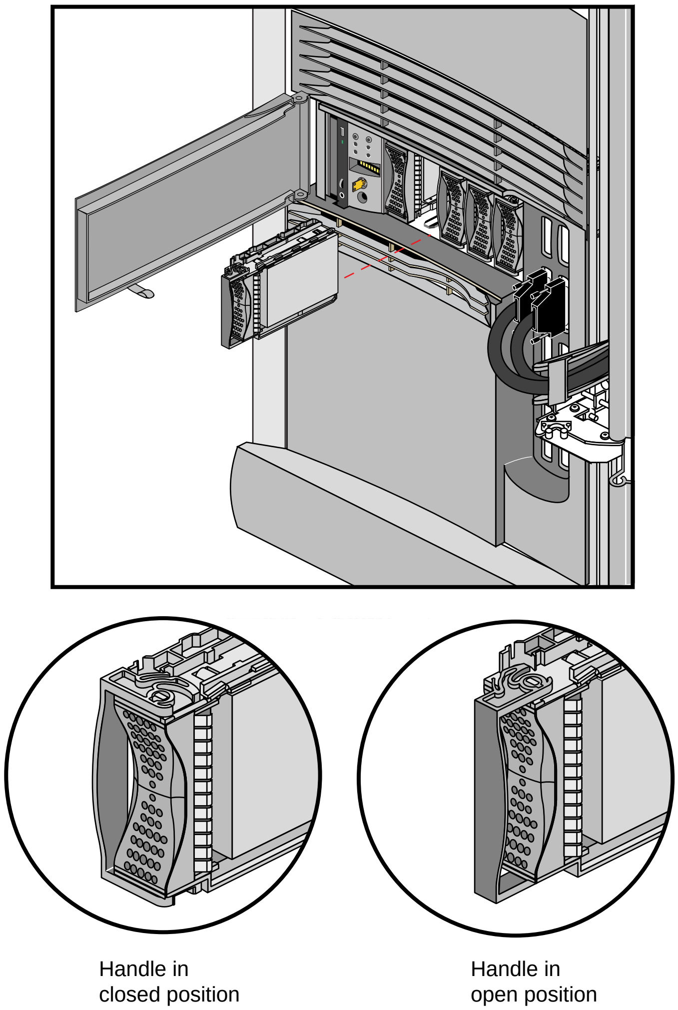

Figure 10-24 Removing the Drive

To *install* a disk drive module, follow these steps:

1. If necessary, snap the handle to the open position so that it is centered, as shown in Figure 10-24.
2. If you are adding a disk drive module, remove the drive filler plate that covers the drive slot that you want to use.
3. Align the new disk drive module with the drive guide (see Figure 10-24).
4. Gently but firmly slide the disk drive module on the guides over the pin.

5. When the disk drive module is fully inserted, snap the handle to the right (closed) position, as shown in Figure 10-24. In the case of the system disk module, which is upside down relative to the other drives, snap the handle to the left to close it.
6. Use the packaging from the new disk module to repackage the old disk module.

10.3.2 Removing the Module System Controller (MSC) and CD-ROM

The module system controller (MSC) and CD-ROM drive are packaged in one assembly. To replace either component, you must remove the old assembly and insert the new one.

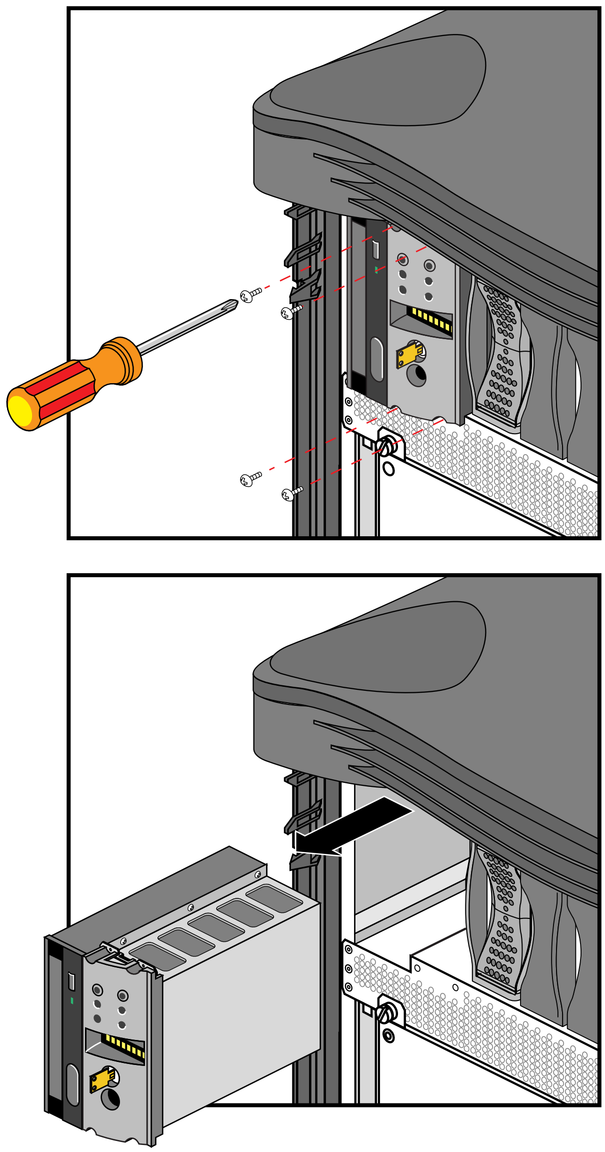

Figure 10-25 Removing the Module System Controller

1. Power off the system.
2. Remove the screws that hold the assembly in place, as shown in Figure 10-25, and save them for reuse.

   Note: Ensure that you do not damage the flex circuit on the CD-ROM.
3. Insert the new assembly and replace the screws.

10.3.3 Removing and Installing a Router Board

To *remove* a router board, follow these steps:

1. Prepare an antistatic surface where you can place the old router board.
2. Remove the facade to expose the Router boards (see Section 10.2.7).
3. Using a 7/64-inch hex driver, loosen the router board's compression connector (CPOP) screws. To avoid damaging the CPOP connector, alternate between the hex screws as you unscrew them. (The hex screws are spring-loaded and captive.)
4. Loosen the two captive Phillips screws at the top and bottom of the router board.
5. Pull the router board straight out and place it on an antistatic surface. Do not touch the CPOP connector or let it contact anything.

To *install* a router board, follow these steps:

Caution: The end of the CPOP connector is extremely sensitive and damages easily. To prevent damage, do not touch the connector, do not let it contact anything other than the midplane, and install the router board with even pressure at the top and the bottom, square with the midplane.

1. Inspect the CPOP connector screws on the new router board. If they have debris or hard gray material in their threads, return the router board and replace it with a clean one, as noted in FIB 100006.
2. Align the router board with the card guide at the bottom of the board slot.

3. Gently but firmly slide the router board all the way into the slot.

   Note: When you insert the board, ensure that it does not contact the EMI wipes (EMI gasket) of an adjacent board and pull the wipes loose. To prevent this, apply slight pressure away from the wipes as you install the board.

4. To help prevent cross-threading, start the hex screws by hand. Insert each screw until it contacts the midplane, turn the screw *counterclockwise* until you hear the threads engage; then partially tighten the screw.

   Caution: You can damage the midplane if you cross-thread the hex screws. If you feel resistance, remove the screws and the router board, examine the hardware, and start over.

5. Using a 7/64-inch hex driver and alternating periodically between the top and bottom screws, gently tighten the hex screws until they are finger tight. Do not overtighten the screws (the torque specification is 6 in.-oz.). Note that some boards have colored rings on the hex screws and black rings on the cutouts for the screws; on these boards, tighten the screws until the colored rings align with the black rings.

6. Using a Phillips screwdriver, secure the retaining screws at the top and bottom of the router board.

7. Use the packing material from the new router board to repackage the old router board.

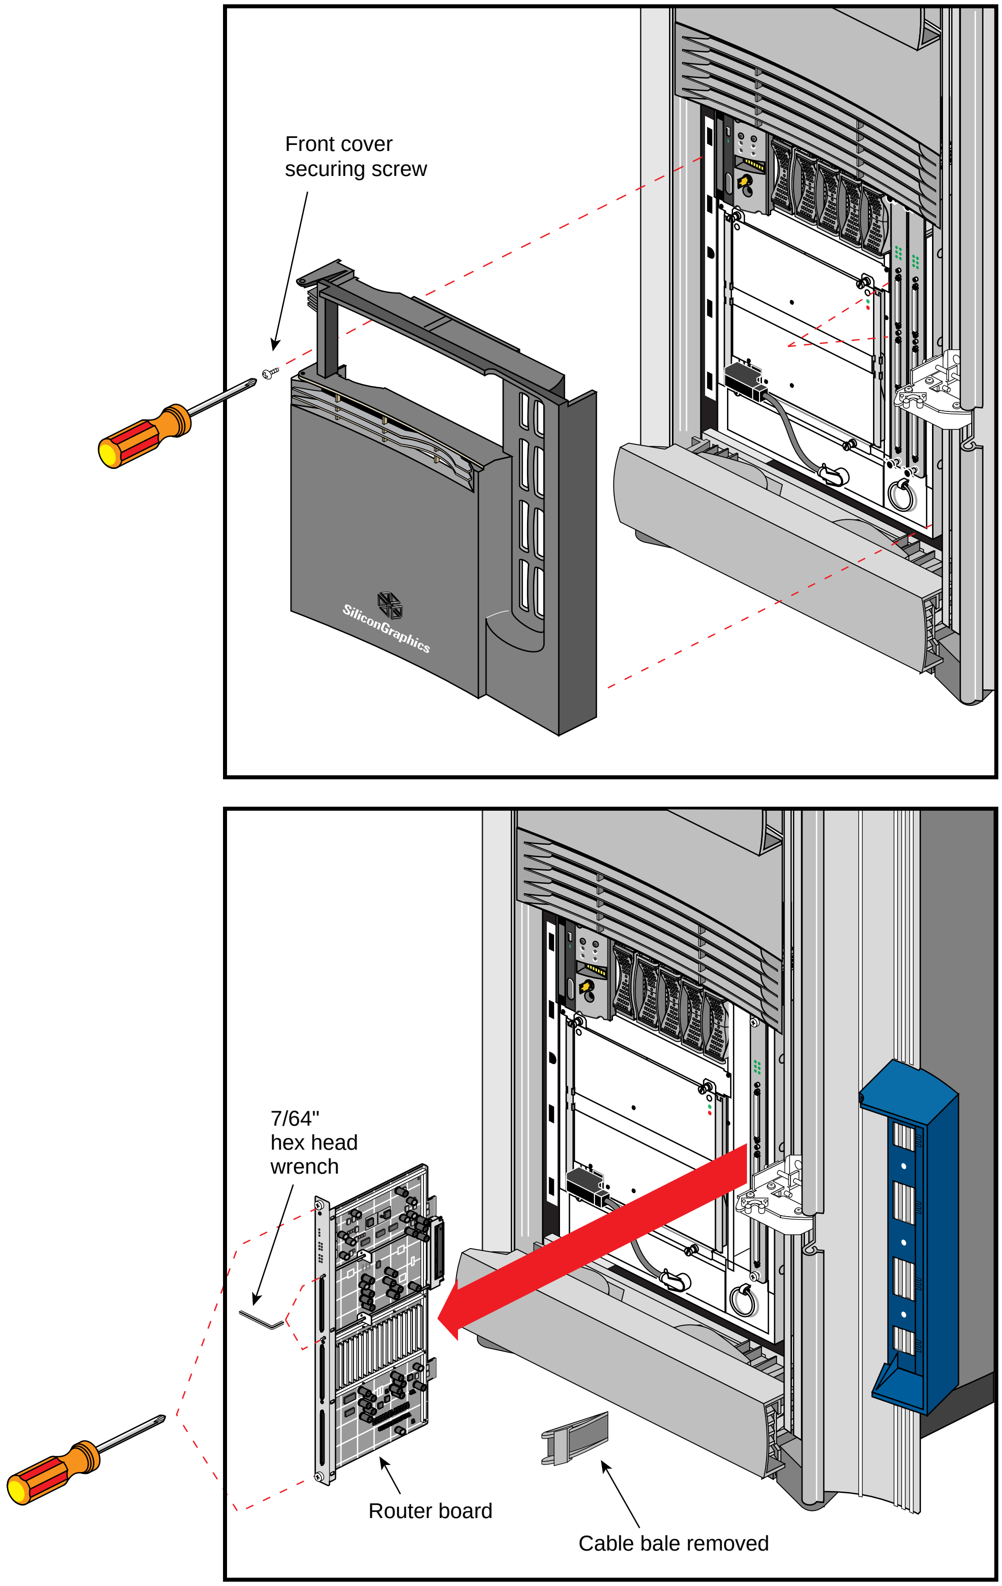

Figure 10-26 Removing a Router Board

10.3.4 Removing and Installing a Node Board

To *remove* a node board, follow these steps:

1. Prepare an antistatic surface where you can place the node board.
2. Using a 7/64-inch hex driver, loosen the node board's compression (CPOP) connector screws. To avoid damaging the CPOP connector, alternate between the hex screws as you unscrew them. (The screws are captive.)
3. Loosen the two captive Phillips screws at the top and bottom of the node board.
4. Slide out the node board, as shown in Figure 10-27, and set it on an antistatic surface. Do not touch the CPOP connector or let it contact anything.

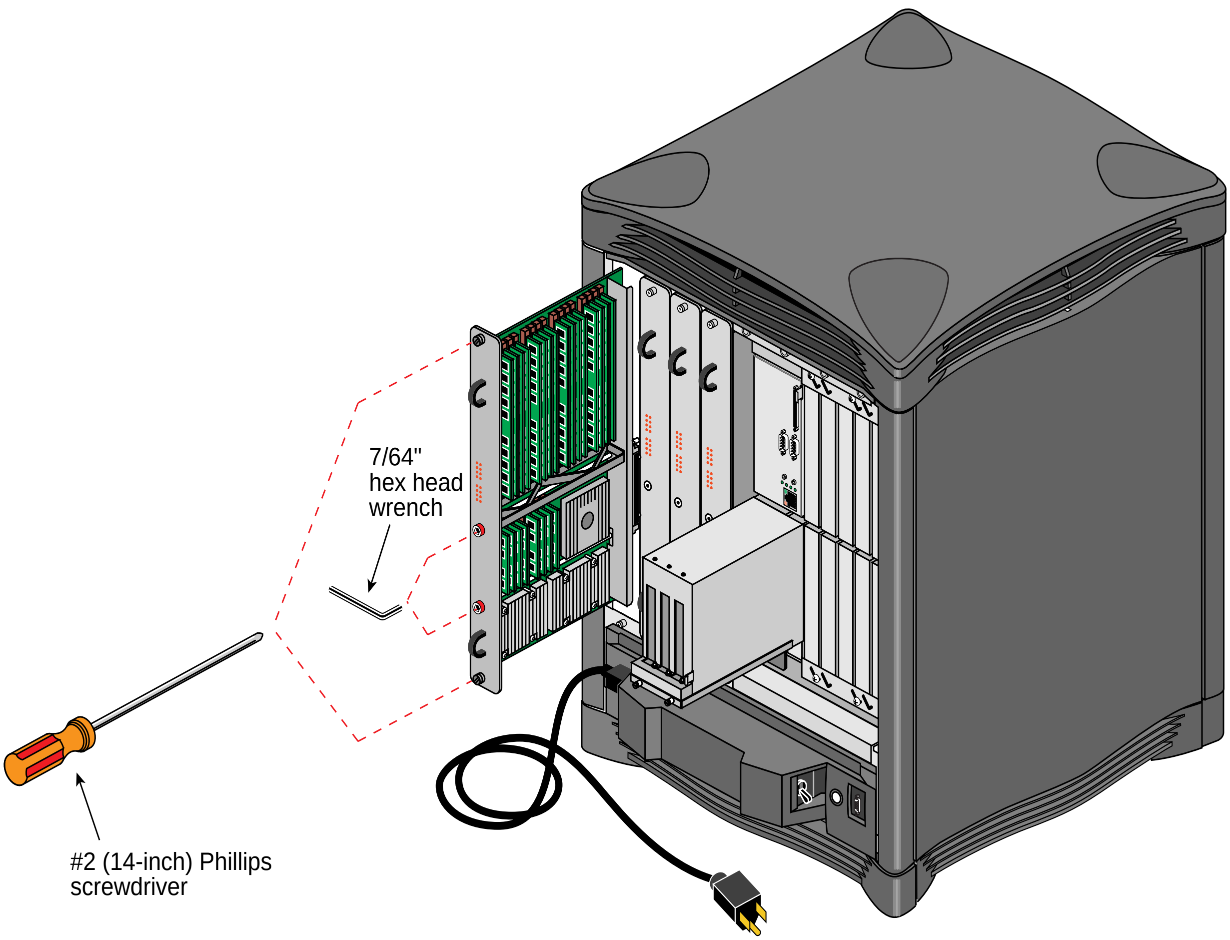

Figure 10-27 Removing a Node Board (Fourth Board Shown)

To *install* a node board, follow these steps:

Caution: The end of the CPOP connector is extremely sensitive and damages easily. To prevent damage, do not touch the connector, do not let it contact anything other than the midplane, and install the router board with even pressure at the top and the bottom, square with the midplane.

1. Inspect the CPOP connector screws on the new node board. If they have debris or hard gray material in their threads, return the board and replace it with clean a one, as noted in FIB 100006.
2. Align the node board with the card guide at the bottom of the board slot.
   - If you are *replacing* a board, align the new board with the card guide at the bottom of the original board's slot.
   - If you are *adding* a board, align it with the card guide at the bottom of the slot closest to a slot that is already occupied. That is, fill the slots from right to left.
3. Gently but firmly slide the node board all the way into the slot.

   Note: When you insert a board, ensure that it does not contact the EMI wipes (finger stock) of an adjacent board and pull the wipes loose. To prevent this, apply slight pressure away from the wipes as you install the board.
4. To help prevent cross-threading, start the hex screws by hand. Insert each screw until it contacts the midplane, turn the screw *counterclockwise* until you hear the threads engage, and then partially tighten the screw.

   Caution: You can damage the midplane if you cross-thread the hex screws. If you feel resistance, remove the screws and the node board, examine the hardware, and start over.
5. Using a 7/64-inch hex driver and alternating periodically between the top and bottom screws, gently tighten the hex screws until they are finger tight. Do not overtighten the screws (the torque specification is 6 in.-oz.). Note that some boards have colored rings on the hex screws and black rings on the cutouts for the screws; on these boards, tighten the screws until the colored rings align with the black rings.
6. Using a Phillips screwdriver, secure the retaining screws at the top and bottom of the node board.
7. Use the packing material from the new node board to repackage the old node board.

10.3.5 Removing and Installing an XIO Board

You must follow configuration guidelines when you install or remove XIO boards. Failure to follow these guidelines can result in system or peripheral malfunction.

Always

- keep the BaseIO board installed in XIO slot 1
- fill the top XIO slots first (XIO slots 3 and 5 should be filled first)
- ensure that the PCI adapter is installed in XIO slot 2
- handle compression connectors with care, as described in “Optional XIO Boards” in Chapter 3

Never

- move the BaseIO board to a slot other than XIO 1
- install a SCSI board in XIO slot 2

- install an XIO board in an unsupported slot (see Section 3.6, "Board Configuration and Layout")

This section explains

- selecting a slot for the XIO board
- installing and removing an XIO board

Refer to Section 10.3.6.1, "Correcting ConsolePath and Console-moving Problems," for details on correcting ConsolePath and console-moving problems that are related to BaseIO boards.

10.3.5.1 Selecting a Slot for the XIO Board

For the Origin2000 systems and the Onyx2 rackmount systems, the XIO boards have a different installation orientation from slot to slot.

- IO1 / IO2, IO5 / IO6, and IO9 / IO10 are installed with the component side facing left.
- IO3 / IO4, IO7 / IO8, and IO11 / IO12 are installed with the component side facing right.

Figure 10-28 shows the orientation of XIO boards for Onyx2 rackmount systems and for Origin2000 deskside and rackmount systems.

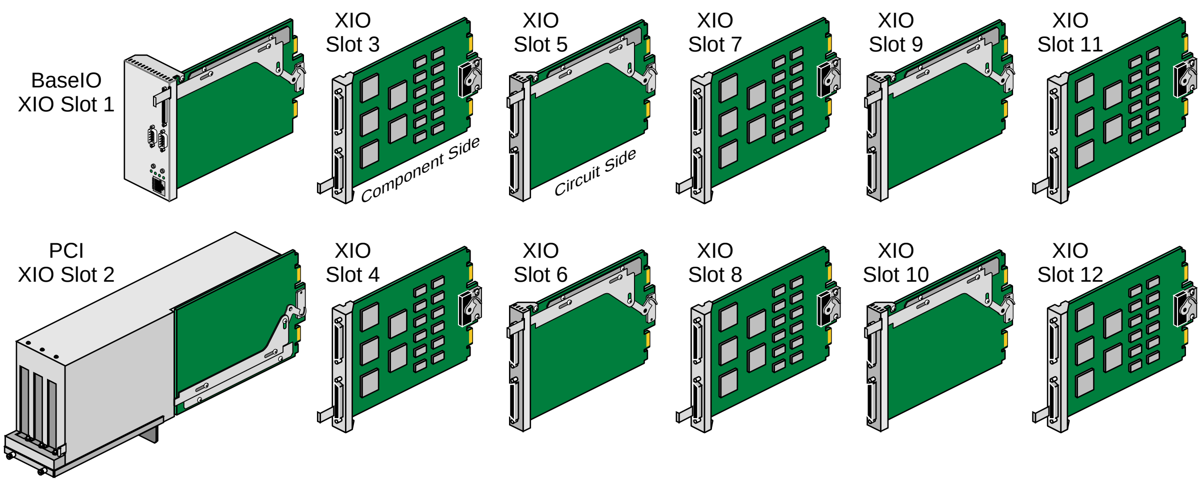

Figure 10-28 XIO Board Orientation (Origin2000 Systems and Onyx2 Rackmount)

Generally, fill all available odd-numbered slots before you fill any even-numbered slots, and fill the lower-numbered slots before you fill the higher-numbered slots. For example, fill slot 3 before you fill either slot 2 or slot 5, and fill slot 7 before slot 2.

In general, if a chassis (module) has a node board in slot *N1* or *N3*, then XIO slots 1 through 6 are available. If it has a node board in slot *N2* or *N4*, XIO slots 7 through 12 are available. If a chassis has two node boards, one in *N1* or *N3* and one in *N2* or *N4*, then all 12 XIO slots are available.

Note: Certain XIO option boards may have restrictions on the slot they can occupy. In addition, rackmount and multimode systems have restrictions on the total number of certain XIO boards per system.

Figure 10-29 shows I/O components in the Origin2000 deskside and rackmount chassis and Onyx2 rackmount chassis.

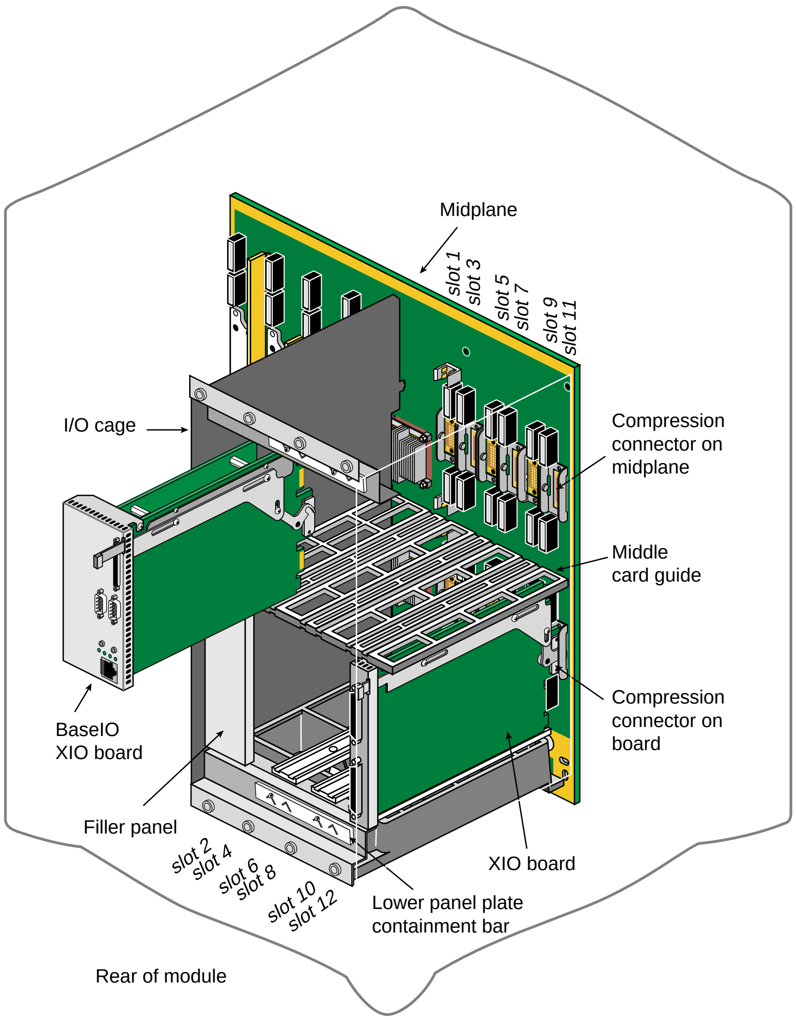

Figure 10-29 I/O Components in the Origin2000 Chassis and Onyx2 Rackmount Chassis

The Onyx2 deskside system has a different board orientation scheme (see Figure 10-31). You can install either three full-height boards or five half-height boards. Slot 4 combines with slot 3 to provide a full-height slot. You can install a half-height board into slot 3, but *you cannot install a half-height board into slot 4.*

Figure 10-30 shows I/O components in the Onyx2 deskside chassis.

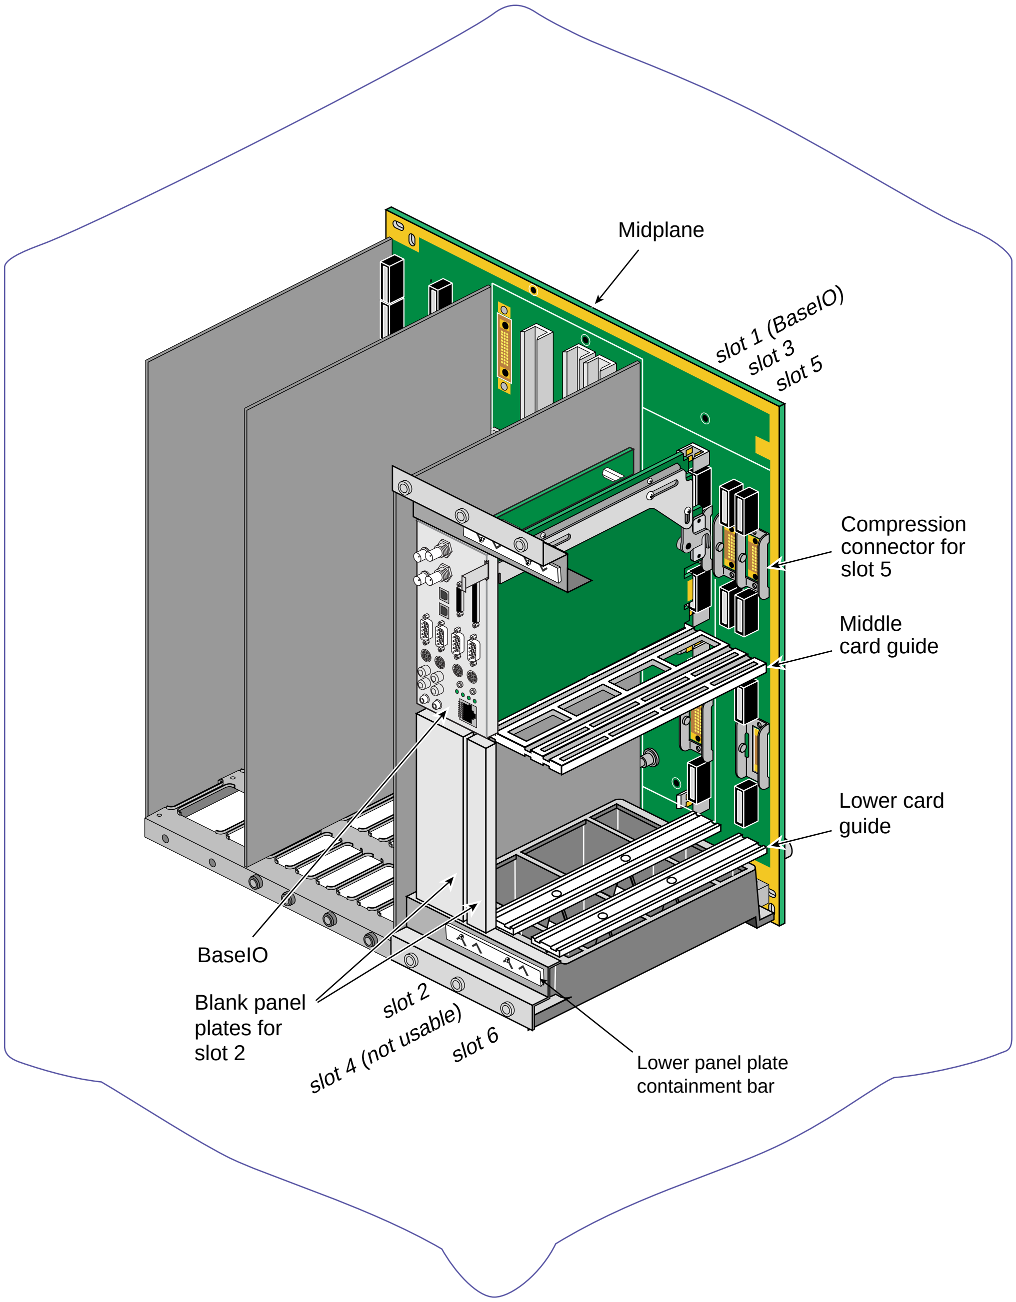

Figure 10-30 I/O Components in Onyx2 Deskside

Figure 10-31 shows the XIO board orientation for the Onyx2 deskside system.

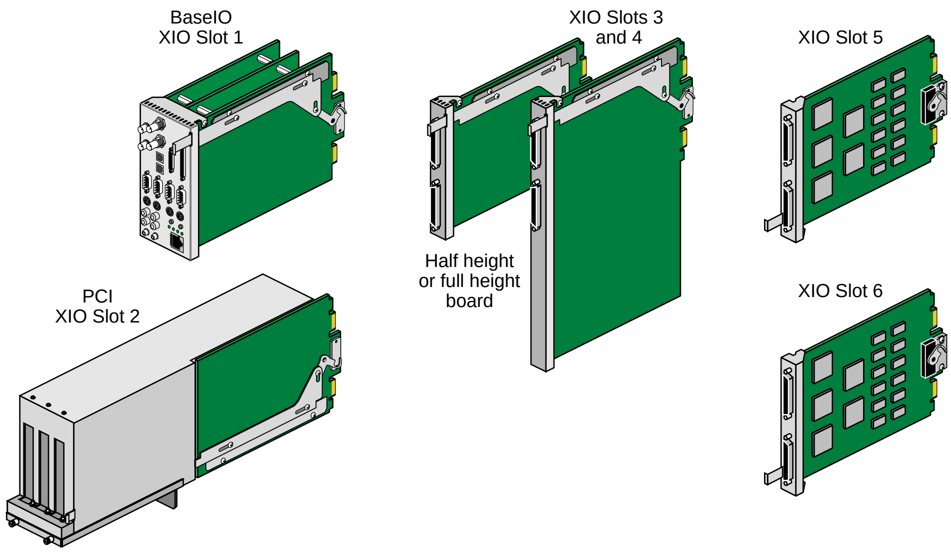

Figure 10-31 XIO Board Orientation (Onyx2 Deskside system)

10.3.5.2 Installing and Removing an XIO Board

To *install* an XIO board, see Figure 10-32 and Figure 10-33 and follow these steps:

Note: If you are installing a BaseIO board assembly, use the procedure in Section 10.3.6.

1. Be familiar with the care and handling of XIO boards, particularly the compression connectors, as explained in "Optional XIO Boards" in Chapter 3.
2. If you are installing the XIO board in a system of interconnected racks, determine the rack and chassis into which you will install the board.
3. Select the XIO slot into which the board will go, following guidelines in "Board Configuration and Layout" in Chapter 3.
4. Prepare an antistatic surface where you can place the XIO board when you remove it.
5. Loosen the holding brackets (containment plates) that secure the XIO boards, and slide the bracket diagonally up or down to fully expose the board slots (see Figure 10-32).

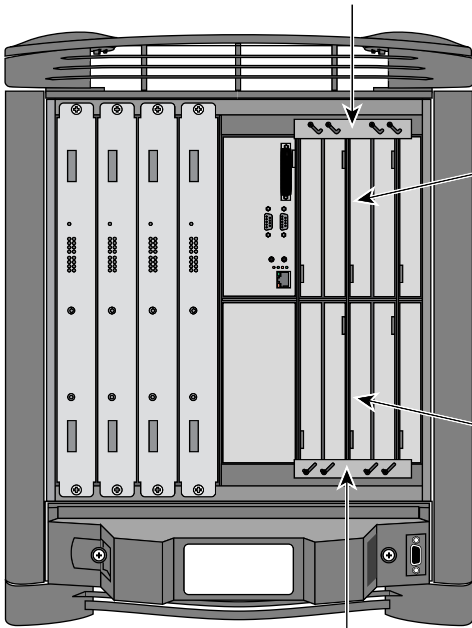

Figure 10-32 XIO Cardcage Holding Brackets and Cover Blanks

6. Remove the appropriate XIO slot cover (see Figure 10-32).
7. Remove the protective cap from the new board's compression connector. Save this cap. You will need it to cover the compression connector if you remove the board for any reason.
8. Orient the board appropriately for the slot; see Figure 10-28 or Figure 10-31.
9. Position the board between the card guides and slide it into the chassis.

   Caution: Take care that no board components are damaged as you slide the board past other XIO boards in the chassis.

10. Verify that the board's panel plate is flush with the other panel plates. If it is not flush, ensure that the board is properly positioned between the card guides; then press gently until it is flush.

    When the board is in place, push the hook actuator handle forward (toward the board) as shown by the large arrow in Figure 10-33 to engage the compression connector.

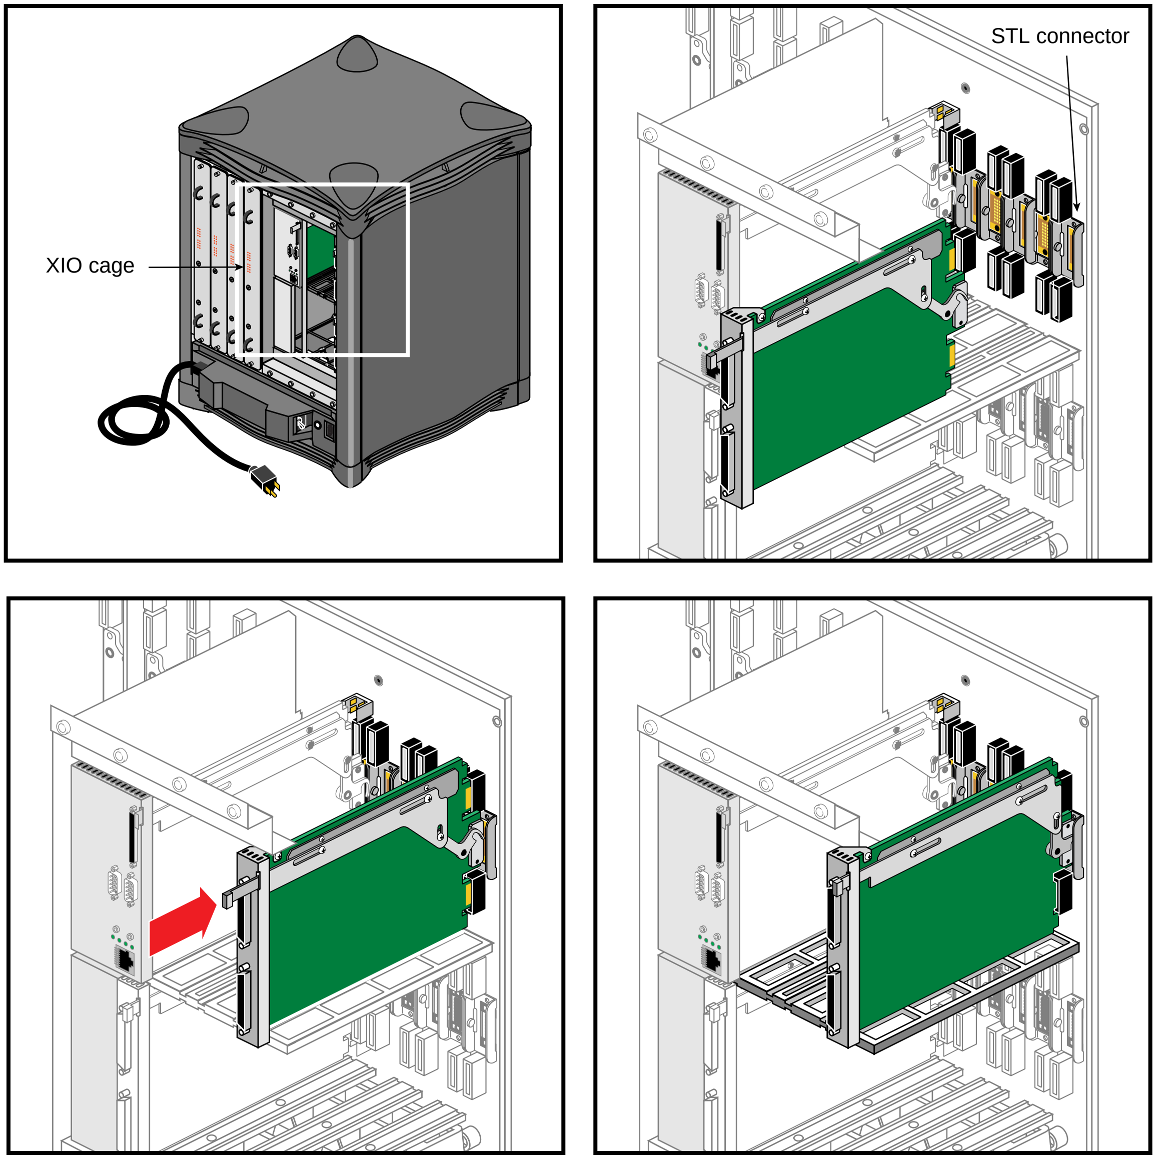

Figure 10-33 Installing an XIO Board

11. Secure the XIO board in place with the holding bracket (see Figure 10-32). Slide the bracket diagonally up or down to hold the board in place in the cardcage, and then tighten the bracket screws.

12. For a rackmount system, use the XIO cable management guide to route, track, and protect the associated board cables (see Figure 10-34).

- Remove the rubber securing band to open a guide slot.
- Flip open the slot tab.
- Insert the cable into the guide slot, close the slot tab, and reattach the securing band to keep the cable in place.

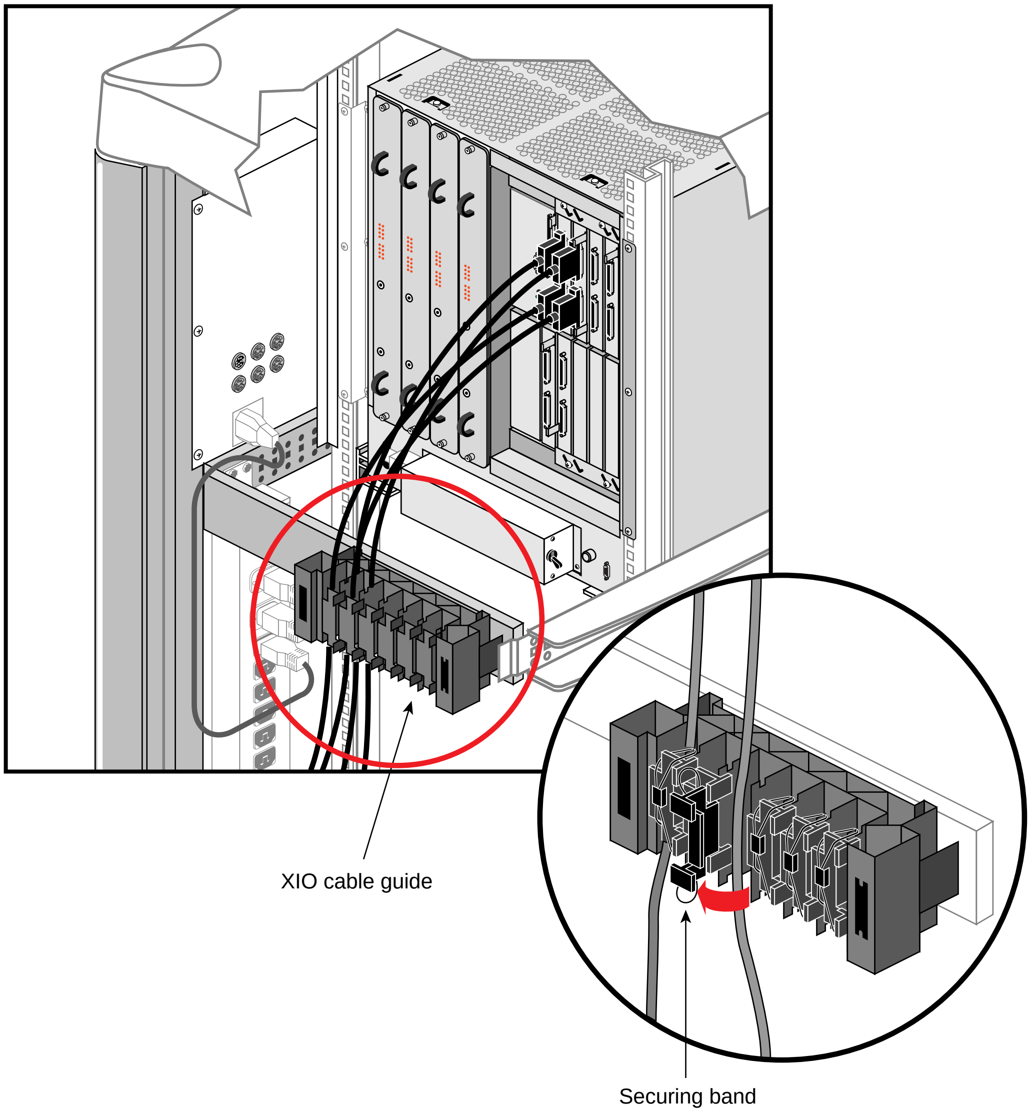

Figure 10-34 XIO Cable Guide (Rackmount System only)

13. Attach any I/O cables for the option.

To *remove* an XIO board, follow these steps:

1. At the back of the chassis, detach all I/O cables for the XIO boards that you are removing. If you have a rackmount system, remove the rubber securing bands that hold the cables in place (see Figure 10-34).

   Caution: Take care that no board components are damaged as you move the carrier past other XIO boards in the cage.

2. Loosen the holding brackets that secure the XIO boards, and slide the bracket diagonally up or down (see Figure 10-32).
3. Pull out the hook actuator handle (locking cam mechanism) to disengage the compression connector.
4. Gently pull the board out of the slot.

10.3.6 Removing and Installing a BaseIO Board Assembly

The BaseIO board is attached to the leftmost XIO board slot (IO1) in the chassis.

Note: If, after you install a BaseIO board, the ConsolePath changes and communication with the console is lost, you must connect the console to the new console connection, which is designated by the letter C in the MSC display. Refer to Section 10.3.6.1, "Correcting ConsolePath and Console-moving Problems," for more information.

To *remove* the board assembly, follow these steps:

1. Prepare an antistatic surface where you can place the board when you remove it.
2. Detach the mouse, keyboard, monitor, and Ethernet cables from the BaseIO panel at the back of the chassis.
3. Loosen the holding bracket that secures the BaseIO board in place, and slide the bracket diagonally up or down (see Figure 10-32).
4. Pull out the hook actuator handle (locking cam mechanism) to disengage the compression connector, as shown in Figure 10-35.
5. Gently pull the board assembly out of the slot.

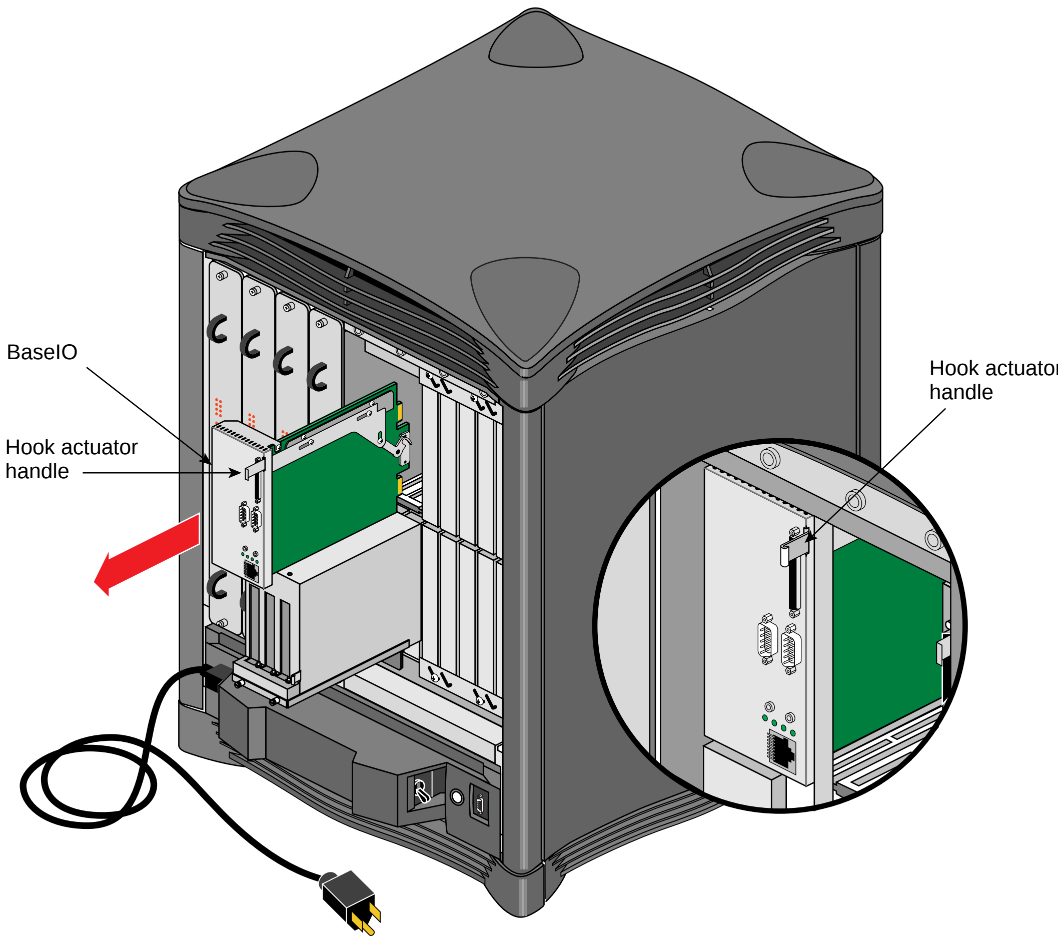

Figure 10-35 Removing the BaseIO Board

To *install* the board assembly, follow these steps:

1. If you are replacing a BaseIO board, remove the removable (coin-shaped) NIC from the old board and install it on the replacement board. Spring clips hold the NIC to the component side of the board.

   Note: If you do not install the old board's NIC on the new board, software will not recognize the new board.

2. Align the replacement BaseIO board with the top and bottom card guides in the IO1 slot, and insert the board all the way.
3. Push the hook actuator handle forward to engage the compression connector.
4. Slide the XIO holding bracket into place to secure the BaseIO board (see Figure 10-32). Tighten the bracket.

10.3.6.1 Correcting ConsolePath and Console-moving Problems

At power-up, the system nodes compare data to determine whether they share a consistent "picture" of the system topology. If they do, the system selects a BaseIO and module to serve as the console connection. It makes the selection based on the ConsolePath variable that is stored in all of the BaseIO NVRAMs in the system. (All BaseIO NVRAMs are in sync; therefore, they all point to the same BaseIO board.)

If the ConsolePath points to a path that is unavailable, software designates the BaseIO in the lowest slot in the lowest module as the console connection. If two nodes in the machine have a valid ConsolePath, then software chooses the master on the basis of the HUB NIC.

You may also modify the ConsolePath. To do so, set the ConsolePath variable, then reset the system. For example:

```
>> setenv ConsolePath /hw/module/2/slot/io1
>> reset
```

When the ConsolePath is modified (by the system or by hand) and the system is reset, the designated module becomes the console connection, the designated module's BaseIO board becomes the master BaseIO, and the attached CPU becomes the global master. You must then ensure that the console physically connects to the console connection, which is designated by the letter C in the MSC display.

Note: Software now tracks the old ConsolePath and prints a warning in the IO6PROM if the ConsolePath moved and does not match the old ConsolePath. If you see the warning, you must set the ConsolePath manually.

Causes of Console-moving Problems

If one of the following conditions exists, the ConsolePath changes and the system loses communication with the console:

- The Hub IO link to the console BaseIO board goes down.

  In this case, the system cannot find the console BaseIO. As a result, the system searches for another module with a BaseIO (as described previously); it then designates that module as the console connection. In other words, the system modifies the ConsolePath variable to point to another module and BaseIO.

  As a result of the change in ConsolePath, communication with the console is lost until you connect it to the new console connection, which is designated by the letter C in the MSC display.

  If the system does not contain another BaseIO board, the system defaults to POD mode in the IP27PROM. When this occurs, communication with the console is lost, and you must communicate via the debug port on the MSC. As a result, you must add a BaseIO board to a functional IO slot to restore the system to operation.

- Module IDs change.
  Because the ConsolePath variable is in the hwgraph format, any change in the module IDs might point the variable to the wrong place. In this case, you must change either the module IDs or the ConsolePath variable to set the console BaseIO appropriately.

  The module IDs can change if one of the following events occurs:

– The user alters the module IDs with a command.

– The system automatically renumbers the modules in the IO6PROM.
The system prompts the user to enable automatic renumbering when it detects an inconsistency in module numbering that is caused by one of the following problems:

A new module, with the same module ID as an existing module, is added to the system. This should not occur if you add a new module from Manufacturing; however, it could occur if you add a module from another system.

The system cannot read the module ID from the MSC, and the backup of the module IDs in the IP27log is not consistent.

• The BaseIO NIC does not contain valid information.
In this case, software does not recognize the board as a BaseIO, and it attempts to select a new BaseIO as the master (as noted previously).

Corrections

The method that you use to correct the console-moving problem depends on the cause of the problem. The following list contains three methods for correcting console-moving problems:

• If the console connection transferred to another module but you wish to continue running without changing the ConsolePath, you may simply reconnect the console and continue running:
Unplug the console and connect it to the module that is designated as the new console connection (C). This will restore communication between the system and the console.

• If the console connection transferred to another module and you wish to transfer it (and the global master) back to the original module, you may reset the ConsolePath and continue running:

  1. Unplug the console and connect it to the module that is designated as the new console connection (C). This will restore communication between the system and the console.
  2. In the Command Monitor, change the ConsolePath to point to the preferred console connection. For example:
     >> **setenv ConsolePath /hw/module/1/slot/io1**

     >> **reset**
  3. Unplug the console and connect it to the module that you specified in Step 2. In other words, connect the console to the module that is now designated by a C in the MSC display.

• If the system is running and you need to replace the master BaseIO, you may replace it while preserving old variables:

Note: All NVRAMs are in sync with the master BaseIO board. If you replace the master BaseIO board, all other NVRAMs are overwritten with the information in the new board. If you replace a BaseIO board that is not the master, the new board will have all of the values of the master. A change to the NVRAM is copied to the NVRAM in all BaseIOs in the system; therefore, you may choose any BaseIO board to be the console connection by setting the ConsolePath variable.

1. Change the ConsolePath to point to another module and BaseIO, but *do not* reset the system. For example:

   ```
   >> setenv ConsolePath /hw/module/2/slot/io1
   ```

2. Power down the system and replace the old master BaseIO. (For details, refer to “Removing and Installing a BaseIO Board Assembly” on page 41.)
3. Unplug the console and connect it to the new console connection, which you set in Step 1.
4. Power up the system. This should bring the new BaseIO board up to date.
5. Change the ConsolePath to point to the new BaseIO. For example:

   ```
   >> setenv ConsolePath /hw/module/1/slot/io1
   >> reset
   ```

6. Unplug the console and connect it to the console connection that you set in Step 5.

10.3.7 Removing and Installing the PCI Adapter

The PCI (peripheral connector interface) adapter allows you to install third-party PCI boards into the Origin2000 and Onyx2 chassis (see Figure 10-36). The PCI signals are converted to crosstalk protocol for communication with the Origin2000 and Onyx2 I/O subsystem through a bridge ASIC. The bridge ASIC also changes crosstalk protocol back to PCI levels. The PCI adapter is attached to the carrier board in the bottom, leftmost slot.

Follow this procedure to install a PCI adapter that already contains the required PCI boards. If you need to install PCI boards into the unit, see Section 10.3.8.

1. Loosen the bottom holding bracket (if it is covering the PCI slot) and move it to the right.
2. Install the PCI adapter in the IO2 slot (the double-wide, bottom, leftmost slot).
3. Push the hook actuator (locking cam mechanism) in to engage the compression connection.
4. Secure the PCI adapter to the chassis by attaching screws to the PCI mounting bracket beneath the unit.
5. Reverse these steps to remove the PCI adapter.

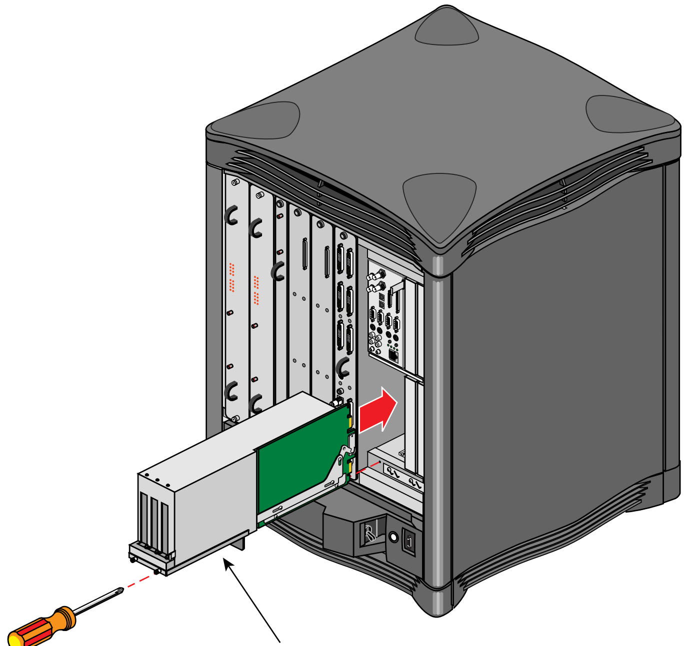

Figure 10-36 Location of PCI Adapter (From Rear of System)

10.3.8 Replacing or Adding a PCI Adapter Board

After the PCI adapter has been removed from the system chassis, follow this procedure to remove or install a PCI board. The PCI carrier can hold up to three half-size PCI boards or up to two full-size boards plus one half-size board (see Figure 10-37).

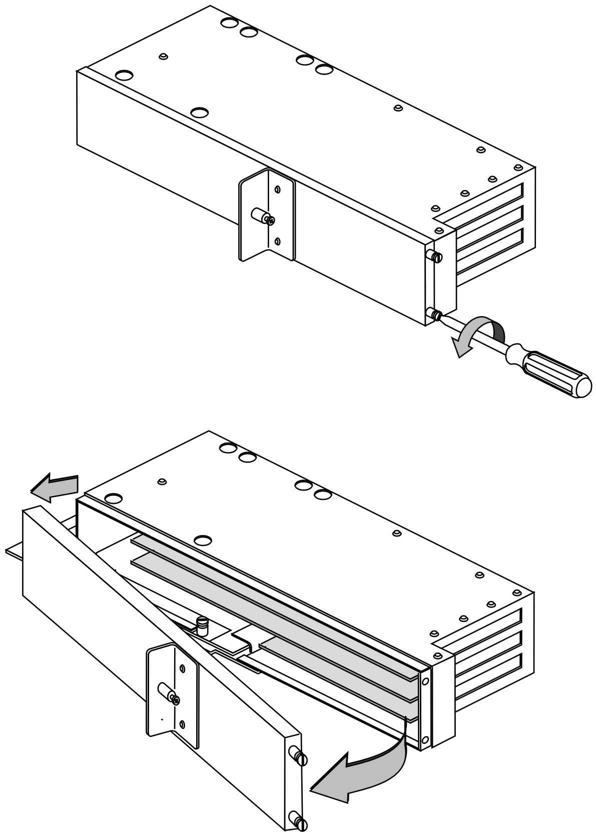

Figure 10-37 Opening the PCI Adapter

Follow these steps to install or remove a PCI board.

1. Set the PCI adapter on its side on a flat surface.
2. Loosen the two screws that secure the top of the PCI carrier (see Figure 10-37) and lift up the top.
3. There are three board slots. Two can hold full-size boards; the third can hold only half-size boards (see Figure 10-37). Install or remove the PCI boards, as required.
4. If you are installing a board, make sure the edge connector makes full contact with the socket inside the PCI adapter.
5. After installing or removing the PCI board(s), resecure the top of the PCI carrier.
6. Reinstall the PCI adapter as described in Section 10.3.7.

10.3.9 Replacing the Power Supply (Origin2000 Systems and Onyx2 Deskside System)

The power supply, at 13 kg (about 28 pounds), is the heaviest component in the Origin2000 and Onyx2 systems.

To replace the power supply, follow these steps:

1. Power off the system, following the usual procedures.
2. At the back of the chassis, unplug the power cord from the power source.
3. At the front of the chassis, detach the fan tray cable, as shown in Figure 10-38.
4. Remove an air baffle at one side of the power supply module by pushing down the tab at the top to release it and pulling the air baffle out from its tab slot at the bottom, as shown in Figure 10-38. Repeat the process with the other baffle plate.

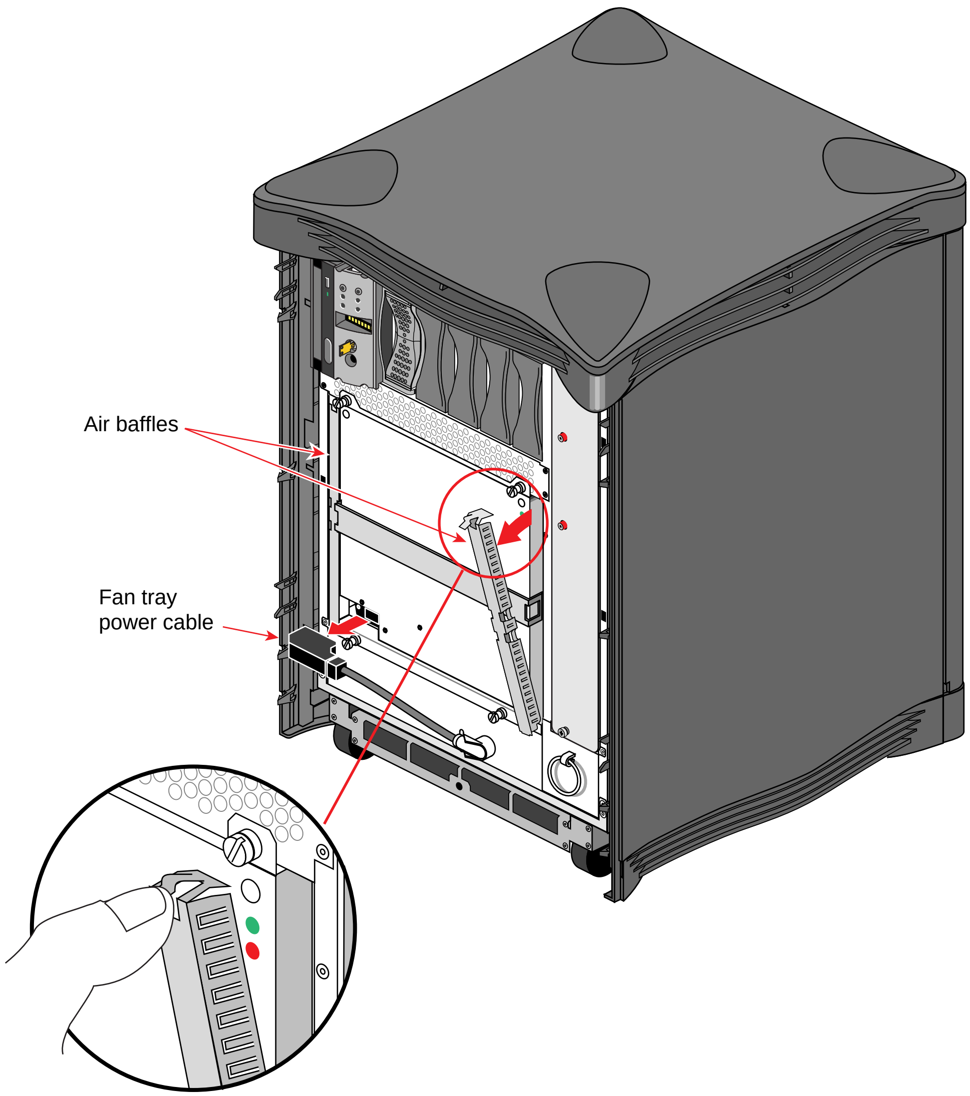

Figure 10-38 Removing the Side Air Baffles

5. Remove the power supply screws, as shown in Figure 10-39, and save them for reuse.

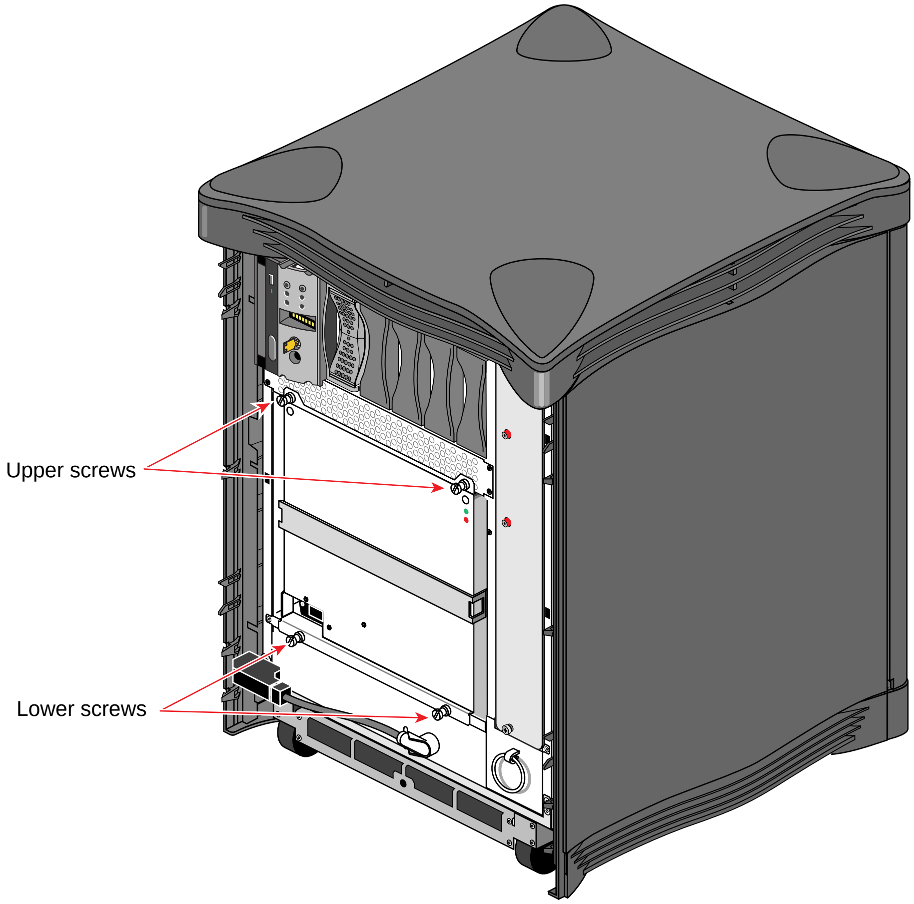

Figure 10-39 Removing Power Supply Screws

6. Pull the power supply handle down and out, as shown in Figure 10-40.

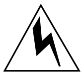

Warning: Carefully ease the power supply out while you support it at the bottom. Be careful not to hurt your back or drop the power supply on your feet.

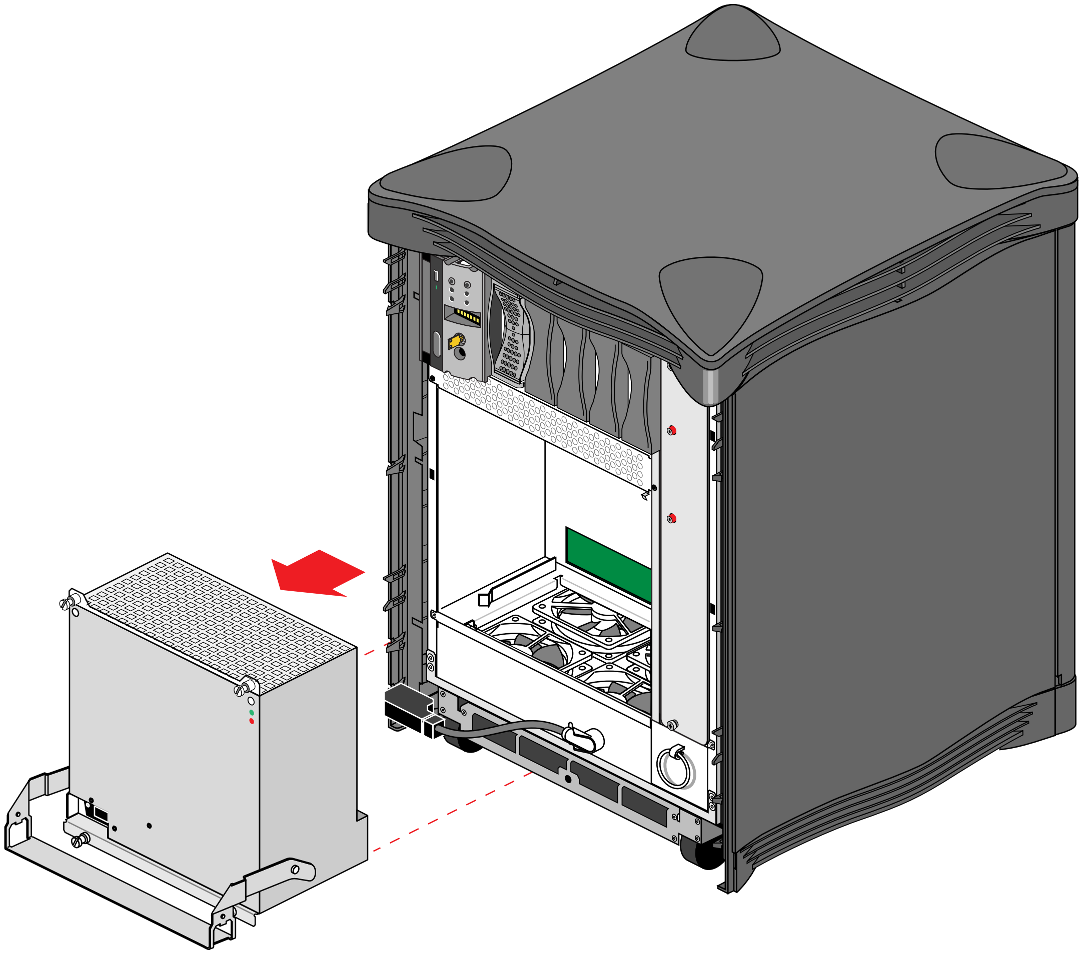

Figure 10-40 Removing the Power Supply

7. Align the replacement power supply in the middle of the bay while you support it with one hand. Slide it into the bay as far as it will go.
8. To secure the power supply in the bay, replace the screws that you removed earlier in diagonal order; for example, top right, then bottom left, then top left, then bottom right.
9. Reinsert the air baffles into the tab slots at the left and right of the power supply module. Replace the screws at the top.

10.3.10 Removing the Fan Tray

Follow these steps to remove the fan tray:

1. Remove the facade if applicable (see Section 10.2.7).
2. Disconnect the AC plug that attaches the power supply to the fan tray (see Figure 10-41).

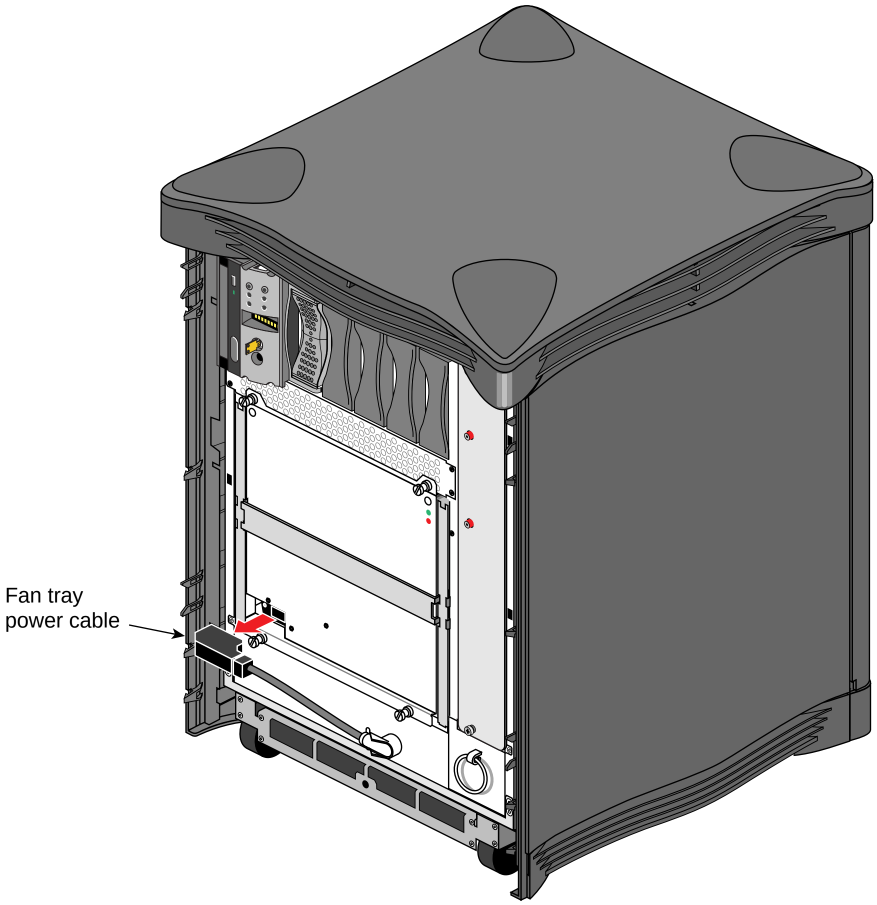

Figure 10-41 Disconnecting the Fan Tray AC Cord (From Front of Chassis)

3. If you have a deskside chassis, pull out the AC and fan cover (see Figure 10-42).

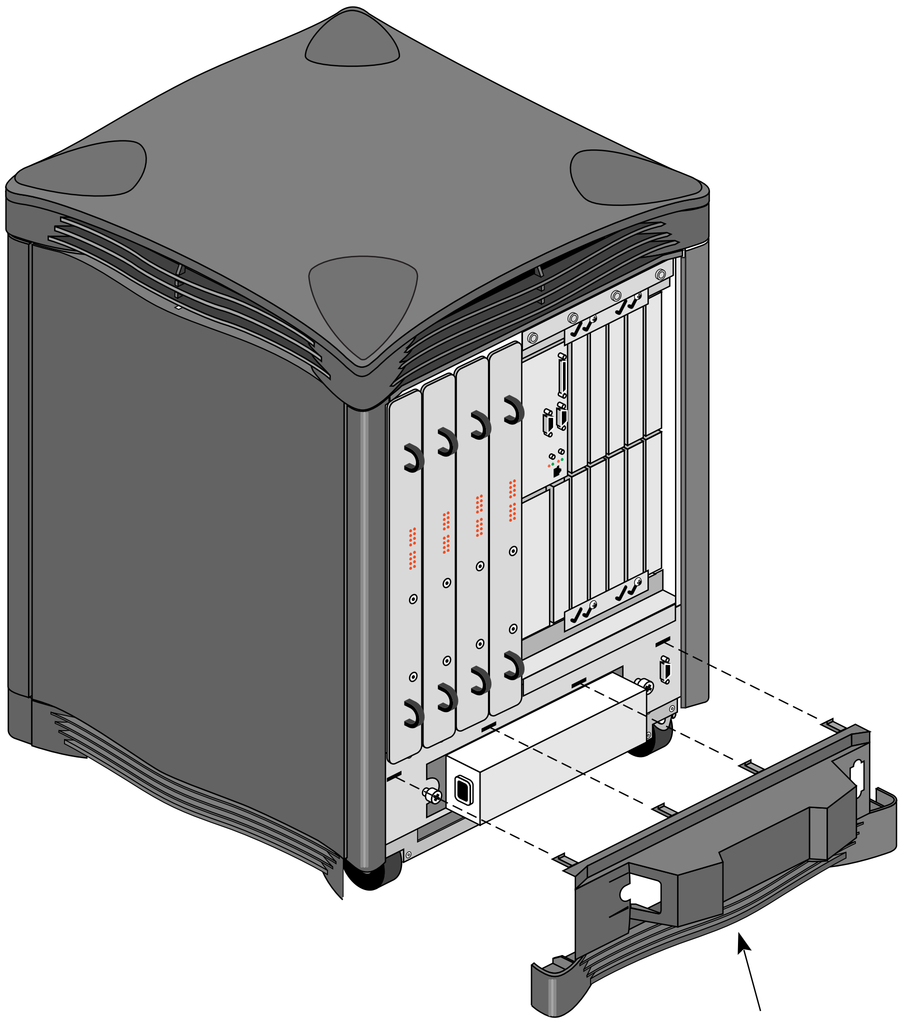

Figure 10-42 Remove the Deskside AC

4. From the rear of the chassis, loosen the two captive screws that secure the fan tray (see Figure 10-43).

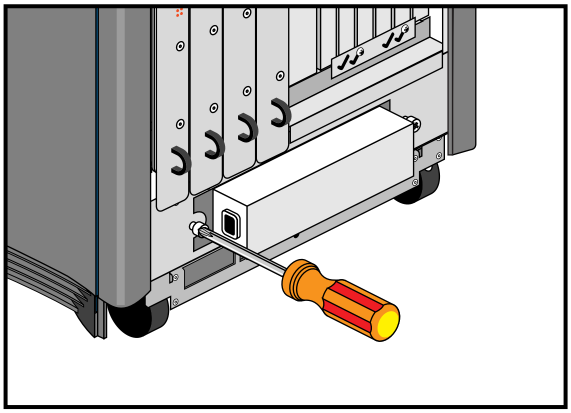

Figure 10-43 Removing / Installing the Fan Tray Screw From Rear of Chassis

5. From the rear of the chassis, partially push the fan tray straight out through the front of the chassis (see Figure 10-44).

   Note: Be careful not to damage the switch or plug.

6. Go to the front of the chassis and pull the fan tray straight out of the chassis. Do not pull the ring on the right side of the fan tray; pulling the ring can cause the fan tray to twist sideways and become stuck in the slot.

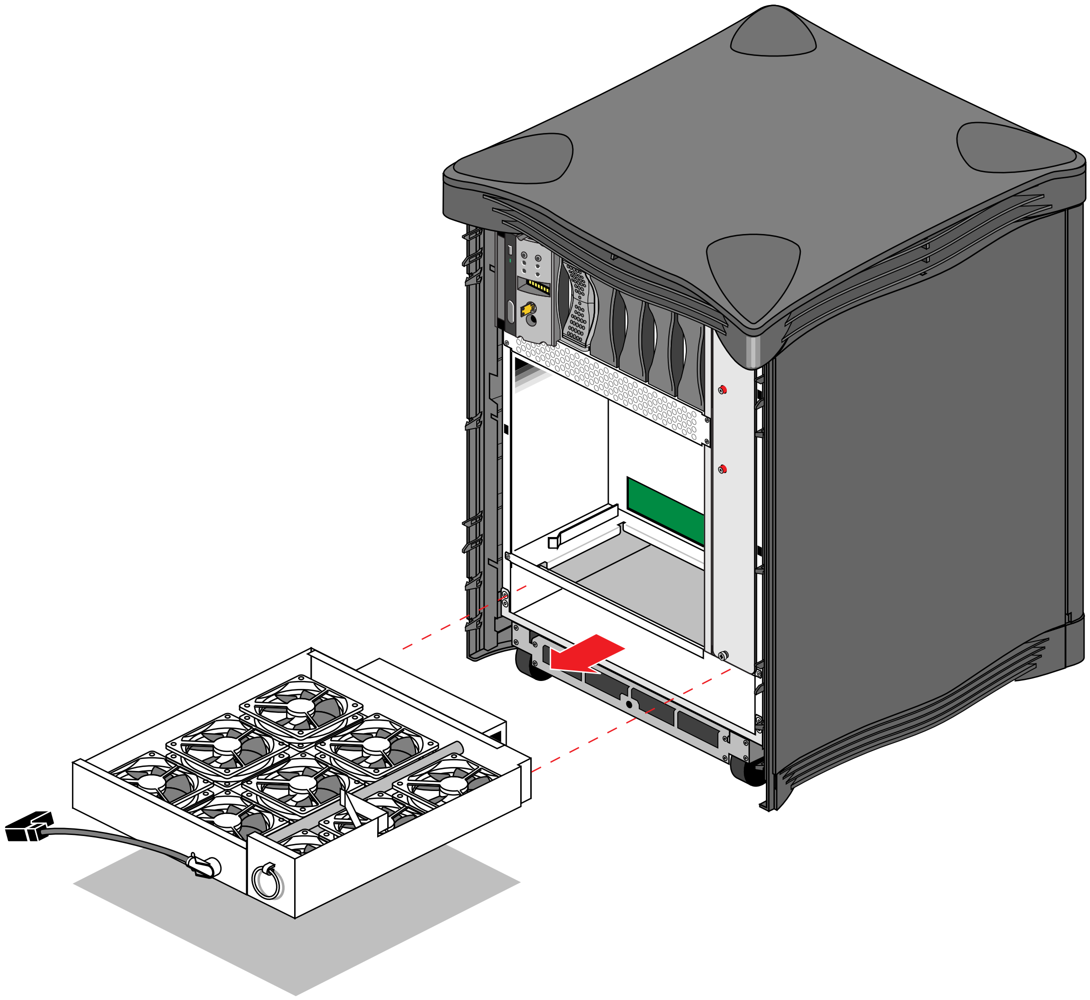

Figure 10-44 Pulling out the Fan Tray

7. To install the fan tray, insert the tray in the front of the chassis.
8. Carefully push the tray back all the way into the rear of the chassis (see Figure 10-45).

   Caution: When you insert the fan tray, the AC connector on the tray may bump against the left captive screw (see Figure 10-45) and prevent you from pushing it all the way into place. If this occurs, go to rear of the chassis and gently push the tray to the right to clear the left captive screw.

9. Resecure the two captive screws to hold the tray in place.
10. Reattach the AC cord to the power supply in the front of the chassis.

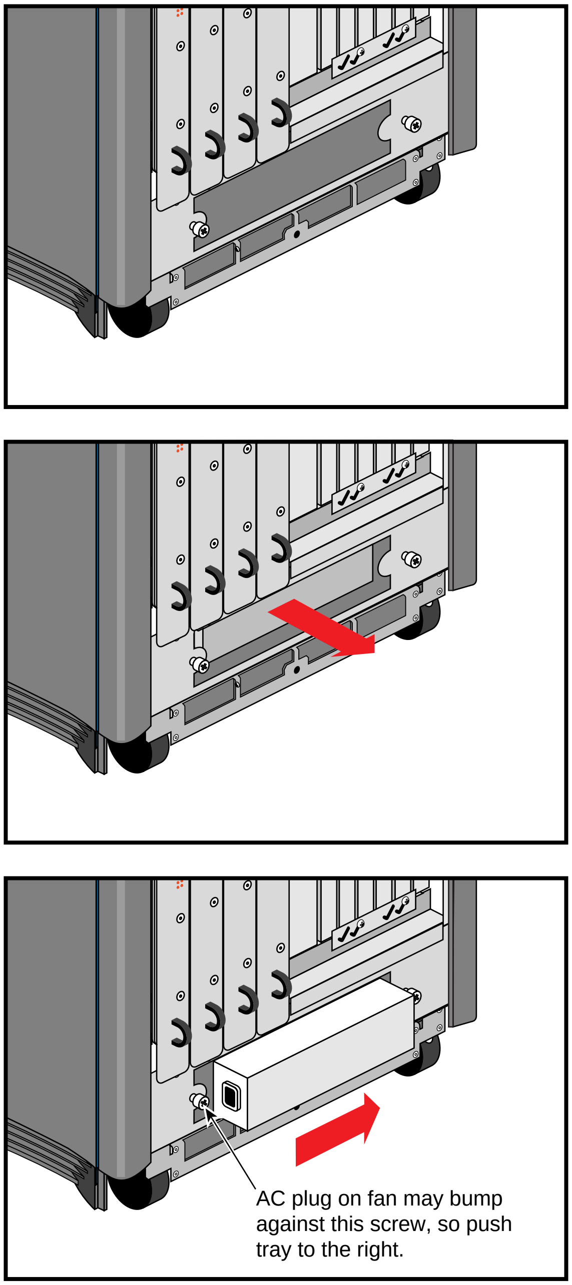

Figure 10-45 Installing the Fan Tray

10.3.11 Removing the Upper and Lower Card Guides

To remove the upper and lower card guides, see Figure 10-46 for Origin2000 and Onyx2 deskside systems or Figure 10-47 for graphics modules, and follow these steps:

1. Remove the retaining screws the secure the guides to the upper and lower part of the chassis.
2. Pull the guides away from the chassis.

   Note: The upper card guide is secured by a locking tab (see Figure 10-46), which you must release when you remove the card guide.

3. To reinstall the card guides, make sure that the guides are properly aligned with the midplane by slightly bowing the card guide inward using a flat-blade screwdriver as shown in Figure 10-46.

   Caution: If you do not perform this step, you could damage the midplane and other boards when you attempt to install node or XIO boards into the rear of the chassis.

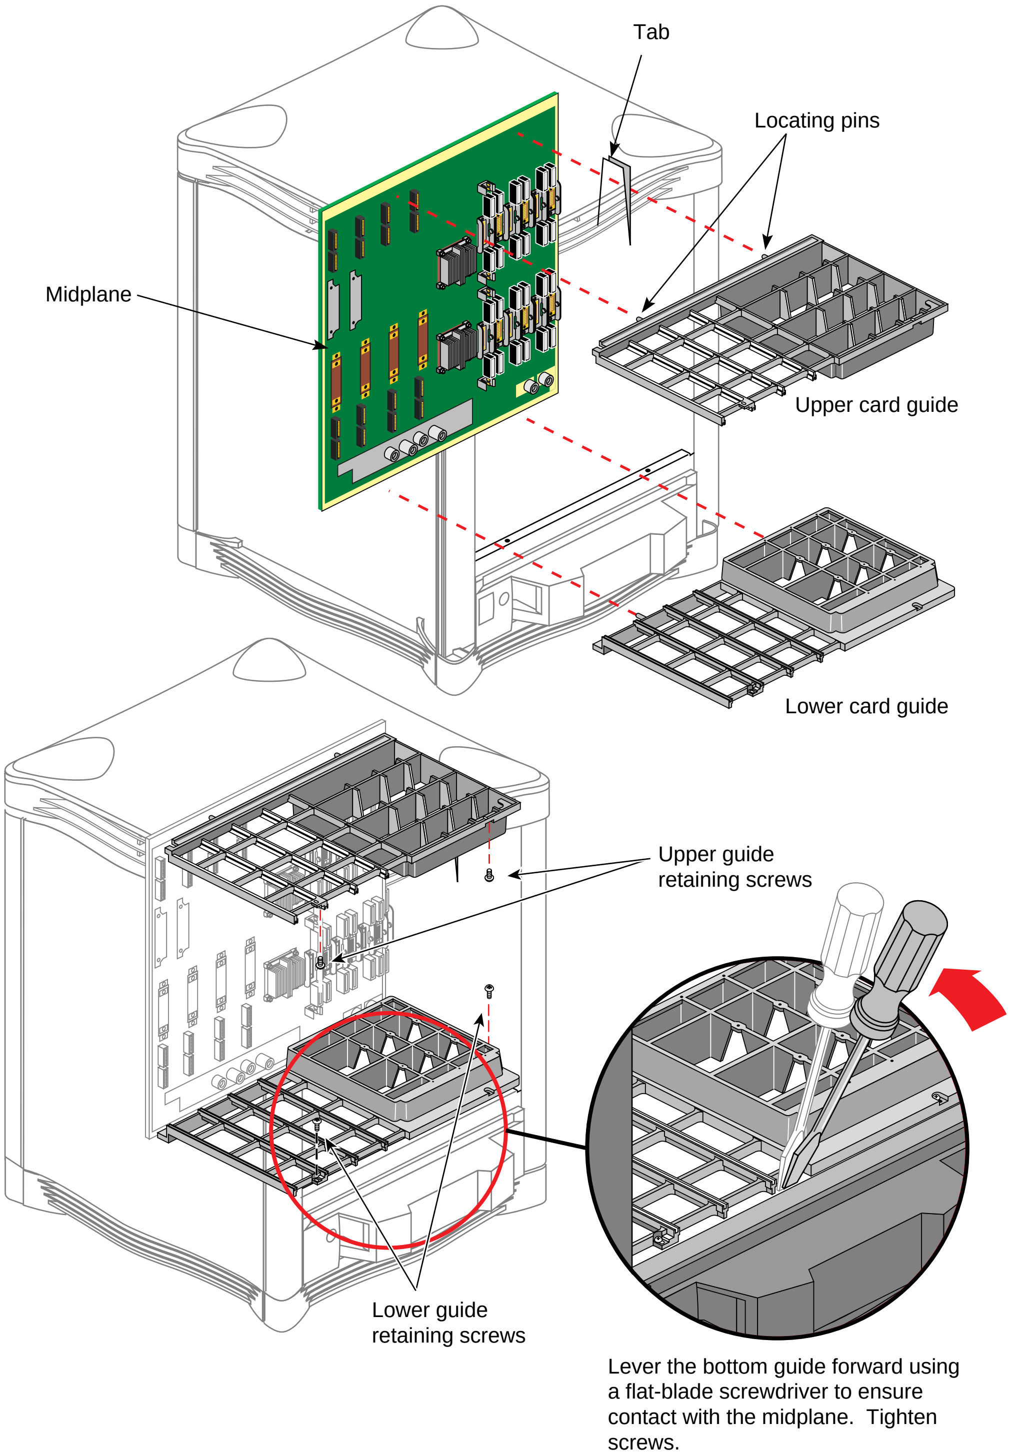

Figure 10-46 Removing / Replacing the Upper and Lower Card Guides

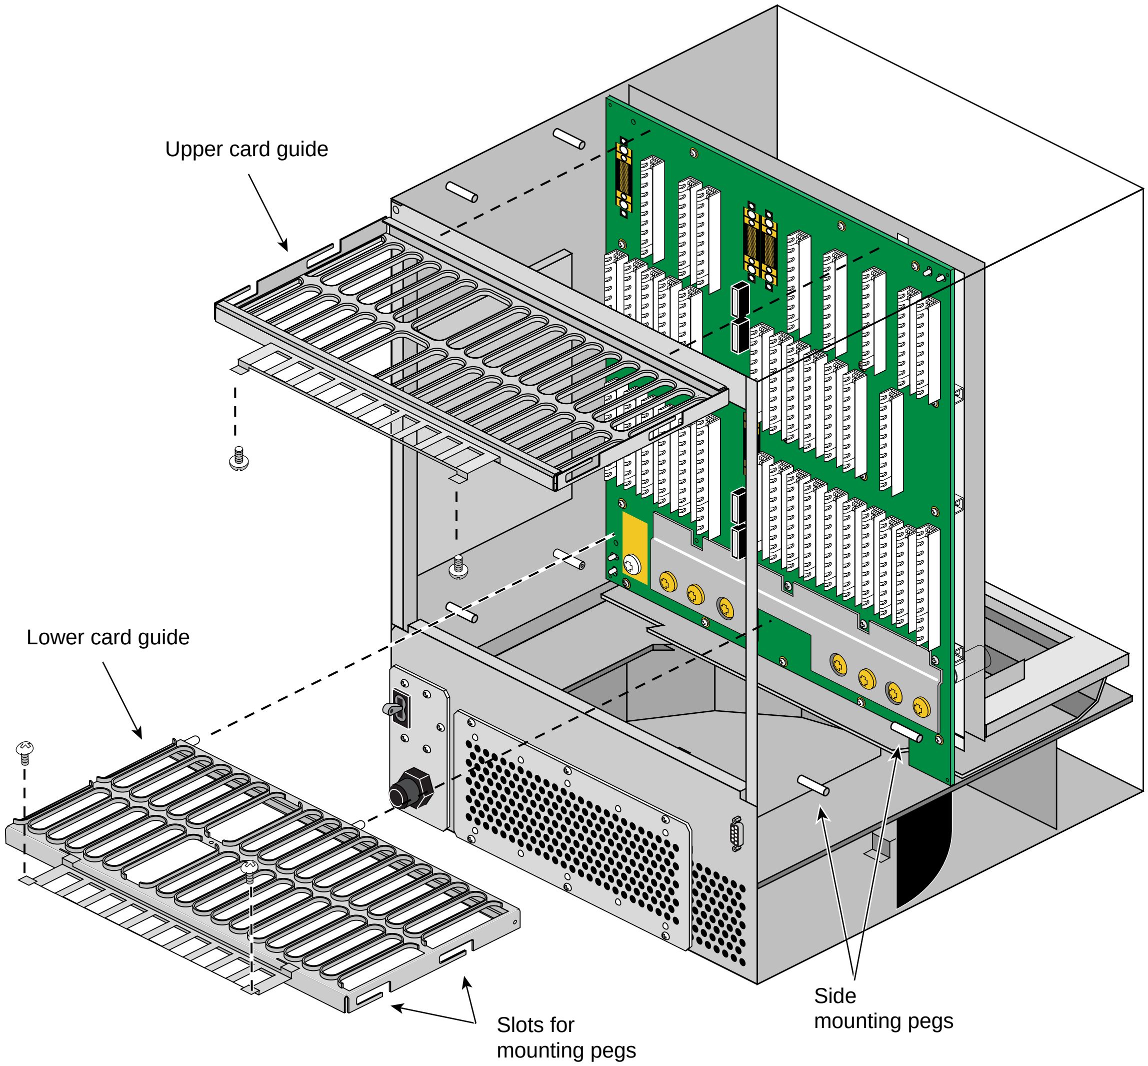

Figure 10-47 Removing the Upper and Lower Card Guides in an Onyx2 Rack Graphics Module

10.3.12 Removing the Card Cage from an Onyx2 Deskside System

To remove the card cage, see Figure 10-48 and follow these steps.

1. Remove the screws that secure the card cage to the chassis.
2. Pull the card cage from the rear of the chassis.

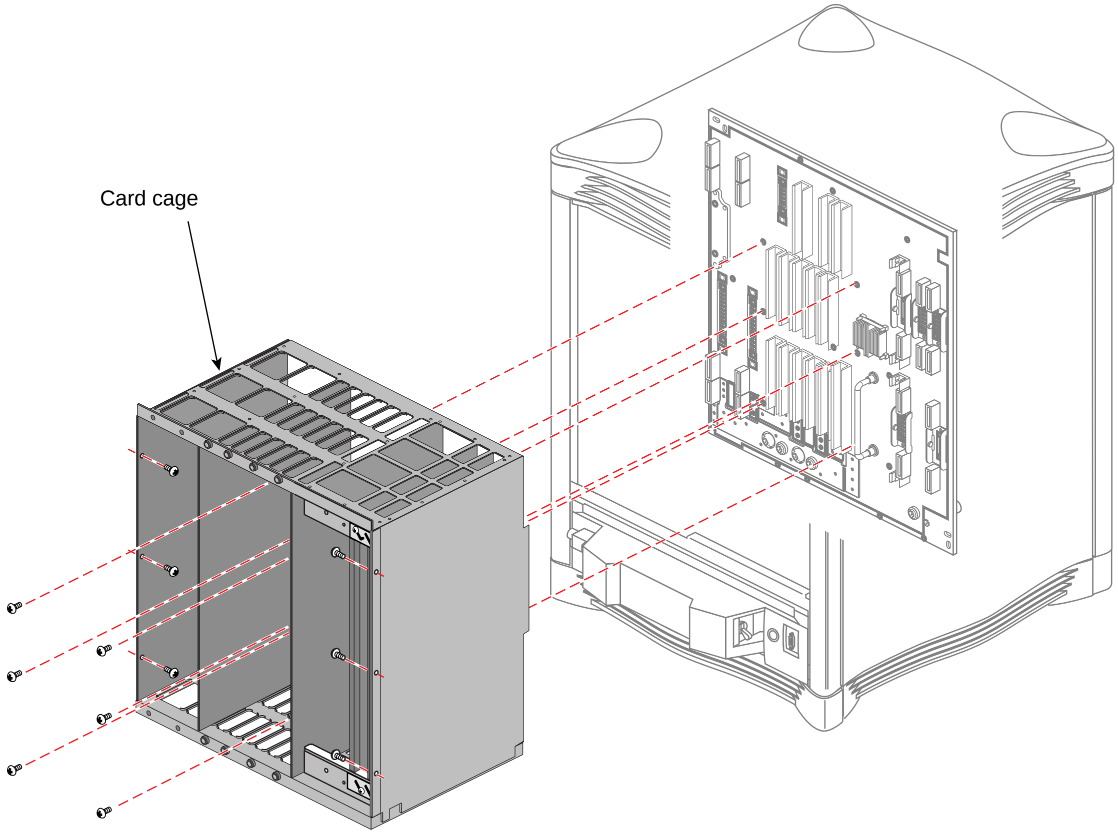

Figure 10-48 Removing the Onyx2 Deskside Card Cage

10.3.13 Removing the I/O Wall

To remove the I/O wall, see Figure 10-49 and follow these steps.

1. Remove the upper and lower screws that secure the I/O wall.
2. Slide the I/O wall out from its railings.

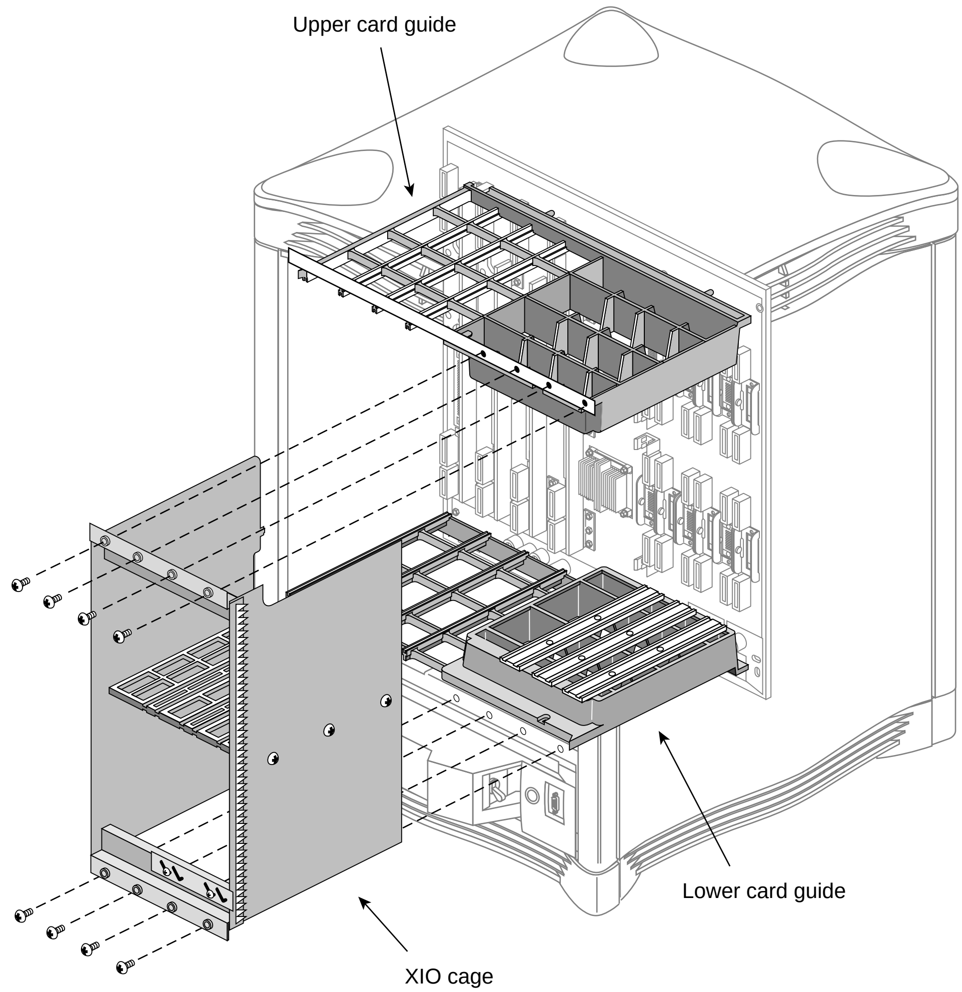

Figure 10-49 Removing the I/O Wall

10.3.14 Removing and Inserting the Midplane

To replace the midplane, detach all boards and other components that plug into the midplane, remove the defective midplane, and then insert the replacement midplane from the rear of the chassis, at an angle.

Warning: **Wait at least 2 minutes after you power off a rackmount system before you work on any part of the power supply or midplane/backplane. Because of the large capacitance in the system, a significant amount of the operating voltage remains on the midplane/backplane for up to 2 minutes after the system is powered off.**

10.3.14.1 Accessing the Midplane

To access the midplane, you must detach the components that plug into the midplane. Here are the general steps for an Origin2000 system.

Note: These steps refer to Origin2000 systems. If you replace the midplane in an Onxy2 system, use these procedures for general reference only and deviate from the procedures as necessary.

1. From the rear of the chassis, remove:
   - the BaseIO board
   - all XIO boards (see Section 10.3.5)
   - all Node boards (see Section 10.3.4)
   - all graphics boards (see Section 10.4.1)
2. From the front of the chassis, remove:
   - all router boards (see Section 10.3.3)
   - all disk and tape drives (see Section 10.3.1)

     Note: Number the disks; you must return them to their original positions.
   - the module system controller (MSC) and CD-ROM (see Section 10.3.2)
   - the fan tray (see Section 10.3.10)

3. After you remove all boards from the card cage area, remove:
   - the power supply (see Section 10.3.9)
   - the I/O sheet metal wall that separates the node boards and the Crosstalk boards (see Section 10.3.12 and Section 10.3.13)
   - the lower card guide (see Section 10.3.11)
   - the upper card guide (see Section 10.3.11)

10.3.14.2 Removing the Midplane

After you detach all the boards and required parts, you can remove the midplane. To remove the midplane, see Figure 10-50 and Figure 10-51, and then follow these steps:

Note: These steps refer to Origin2000 systems. If you replace the midplane in an Onxy2 system, deviate from these procedures as necessary.

1. Remove the screws that secure the midplane to the chassis. In Origin2000 systems, remove the screw that is located in the left center last.
2. Remove the midplane at an angle with one of the corners tilted to a 45 degree angle (approximately) as shown in Figure 10-50.

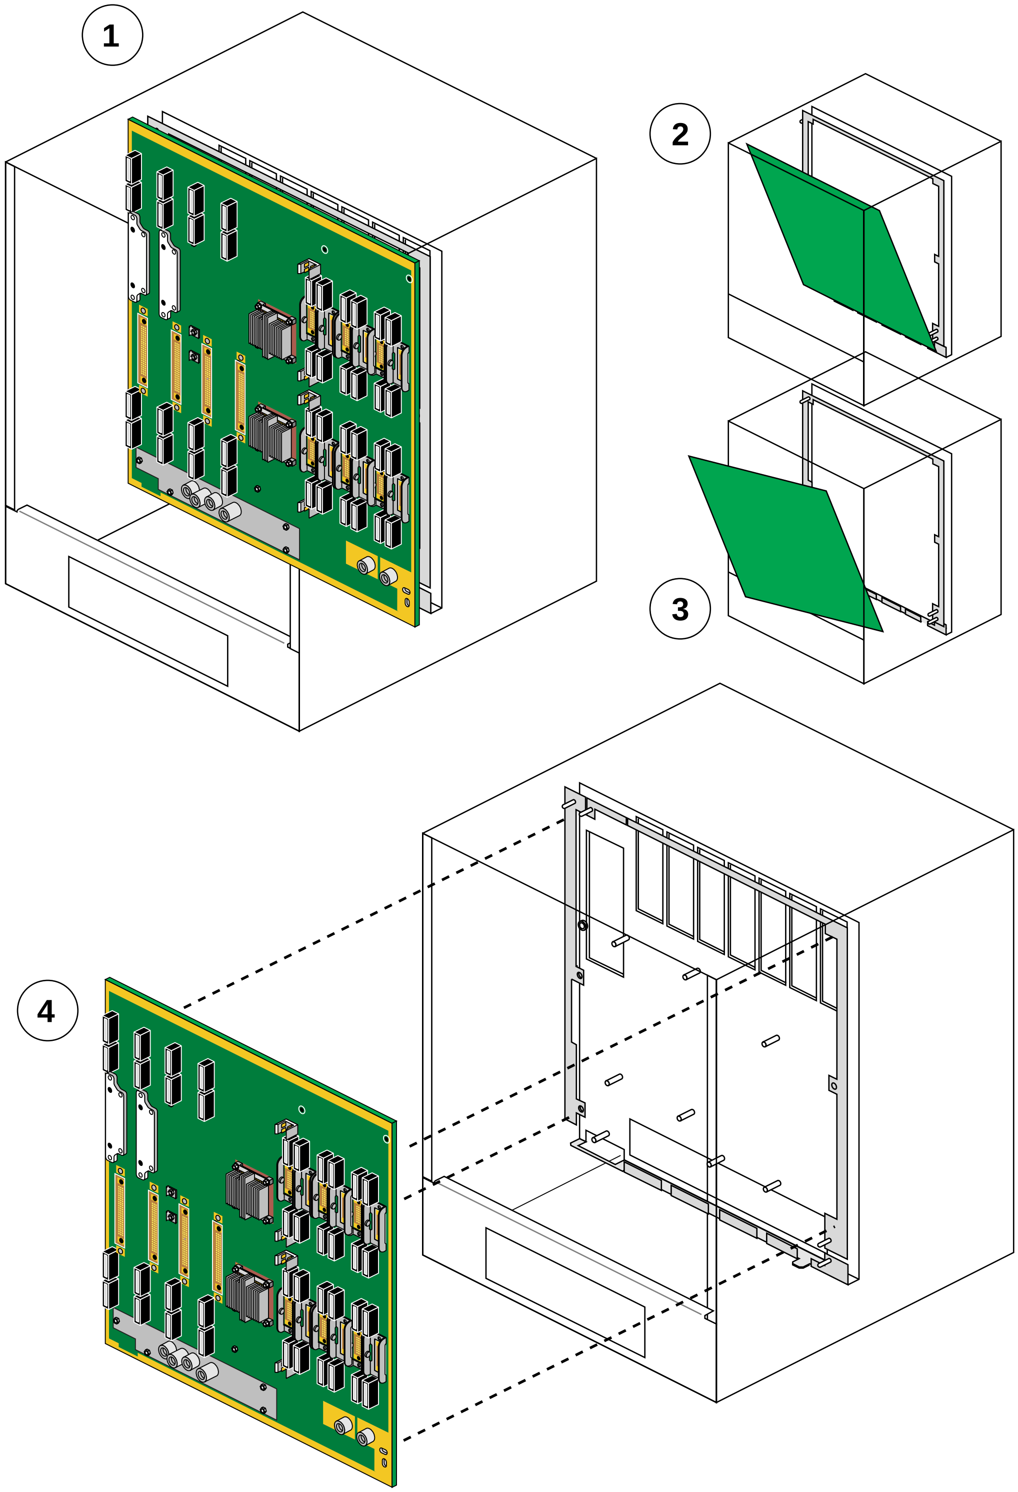

Figure 10-50 Removing the Origin2000 Midplane

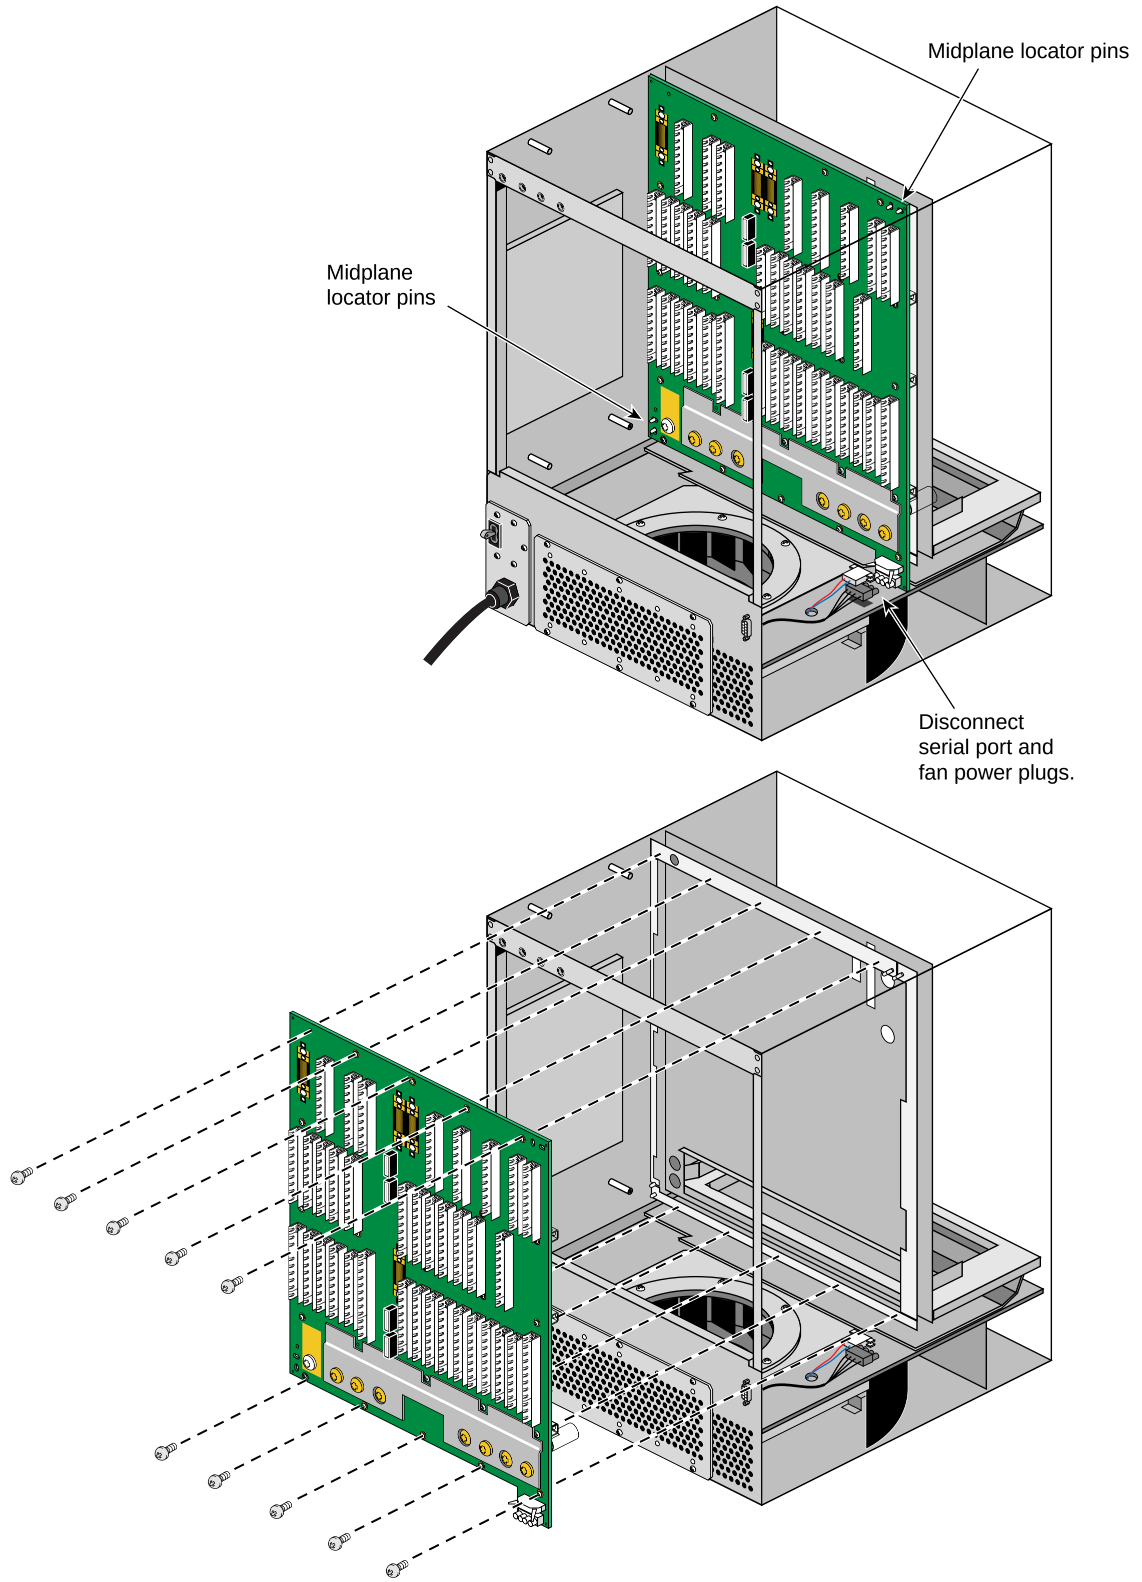

Figure 10-51 Removing the Graphics Module Midplane (Onyx2 Rack)

3. Set the midplane on an antistatic surface.

Caution: Be careful not to press against the midplane when it is lying on a surface. There are components on both sides of the midplane.

4. Locate and remove the coin-shaped NIC from the old midplane (see Figure 10-52 and Figure 10-53). Two spring clips hold the NIC in position.

10.3.14.3 Installing the Midplane

Follow this procedure to install the midplane.

1. Install the *old* midplane's removable NIC on the new midplane (see Figure 10-52 and Figure 10-53).

Note: The removable NIC contains the chassis ID and licensing information for the module. If you do not install the old midplane's NIC on the new midplane, software will not recognize the module and, therefore, will not function.

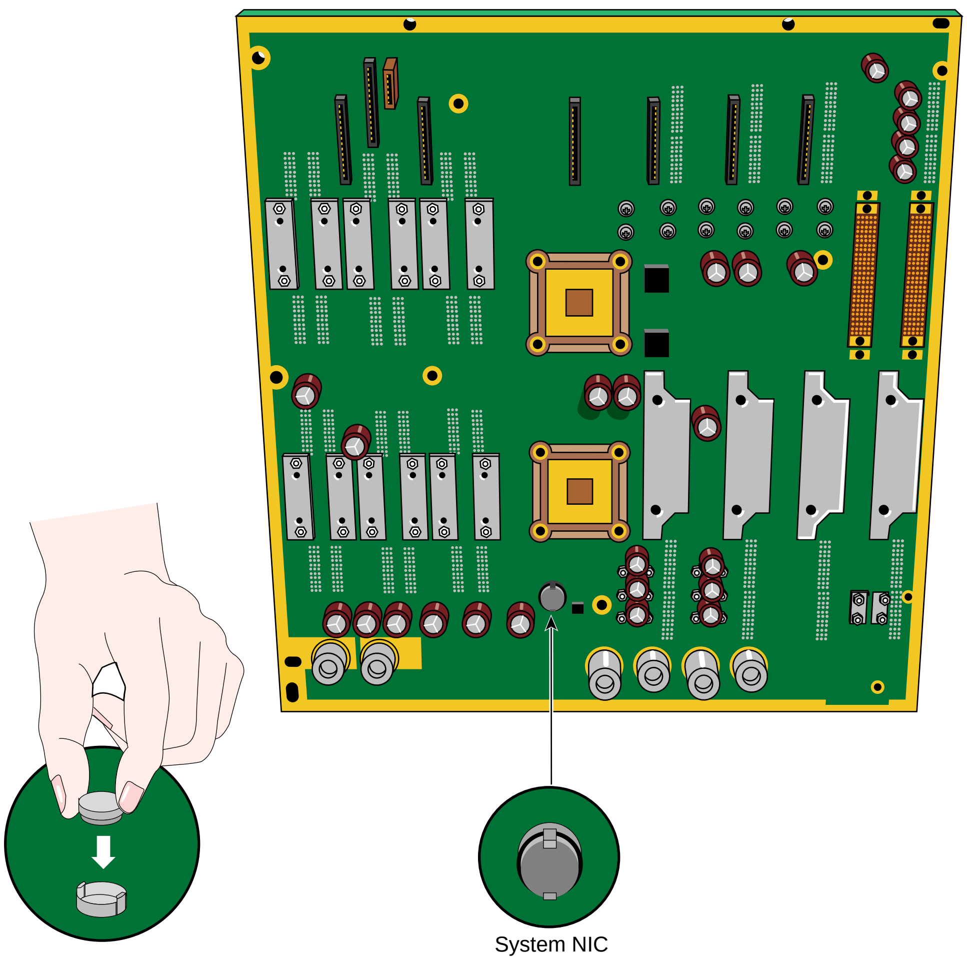

Figure 10-52 Origin2000 Midplane NIC

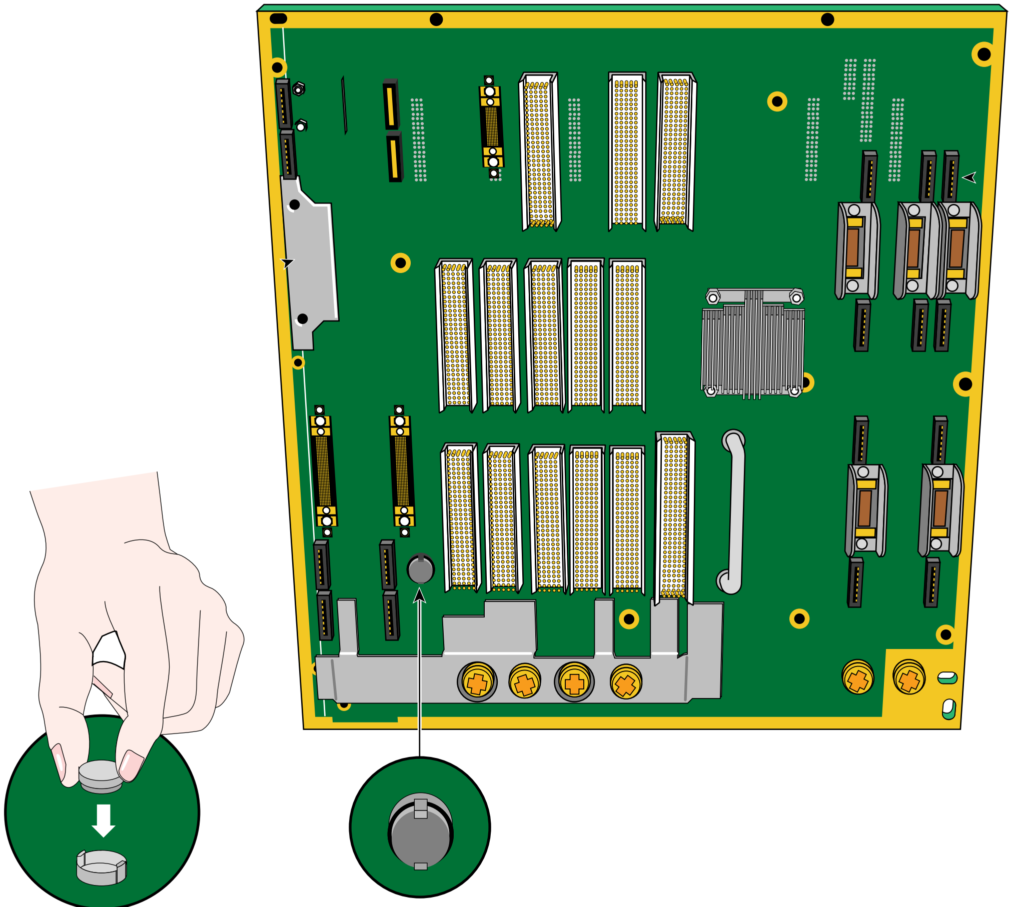

Figure 10-53 Onyx2 Midplane NIC

2. Angle the replacement midplane so that you can position it in the chassis (see Figure 10-54).
3. Place the midplane onto the corner locator pins to hold the midplane in place.
4. Secure the midplane to the chassis with screws.

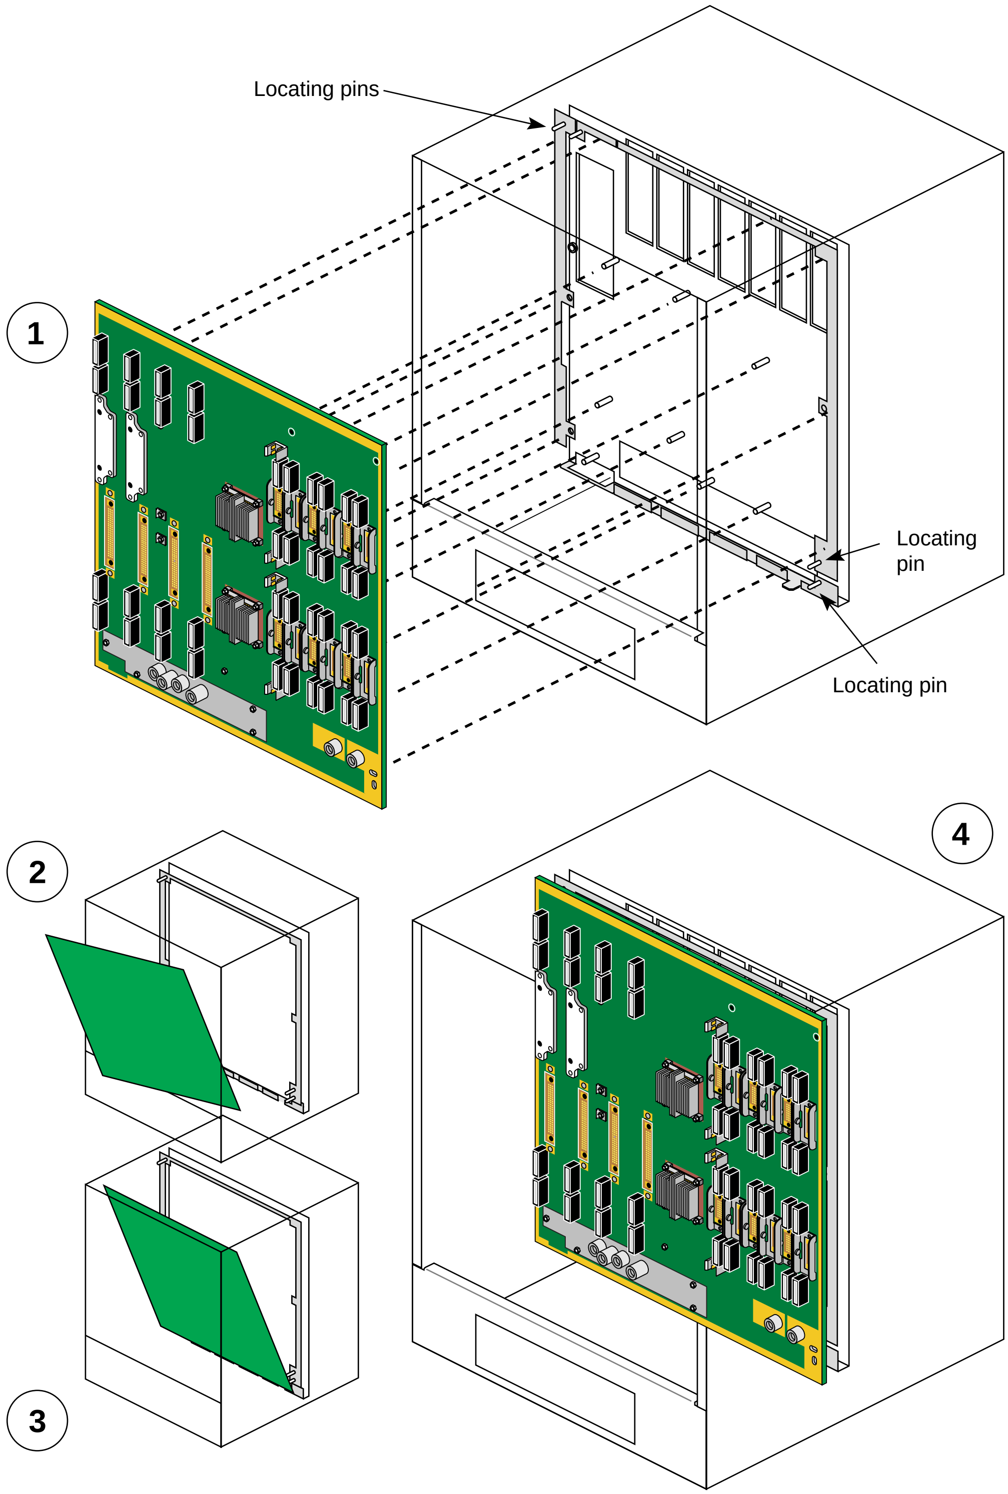

Figure 10-54 Installing the Origin2000 Midplane

10.4 Specific Procedures for Graphics Deskside and Rack Systems

This section provides removal and replacement procedures that apply only to graphics-related field-replaceable units (FRUs), such as:

- graphics board removal for Onyx2 system
- power supply removal for Onyx2 rackmount graphics module
- fan tray removal for Onyx2 rackmount module

To replace an FRU, first use Figure 10-3 through Figure 10-7 to identify the appropriate unit and its position in the chassis. Then proceed to the appropriate section and perform the steps.

10.4.1 Removing and Installing the Graphics Boards

This section describes how to remove graphics boards from the Onyx2 deskside and rackmount systems.

1. Power the system off as described in Section 10.2.1 or Section 10.2.2.
2. If you are removing the DG5 board, make a diagram or label the cables that are attached to the DG5 I/O panel.
3. Disconnect all cables from the board.
4. Loosen the captive screws at the top and bottom of the board and compression connectors (if applicable) as shown in Figure 10-55 and Figure 10-56.
5. Slide out the board.
6. To install a replacement board, align the new board with the card guide at the bottom of the original slot. Gently but firmly slide the board all the way into the slot.
7. Tighten the captive screw and compression connectors.
8. Reconnect (or connect) cables to the board.

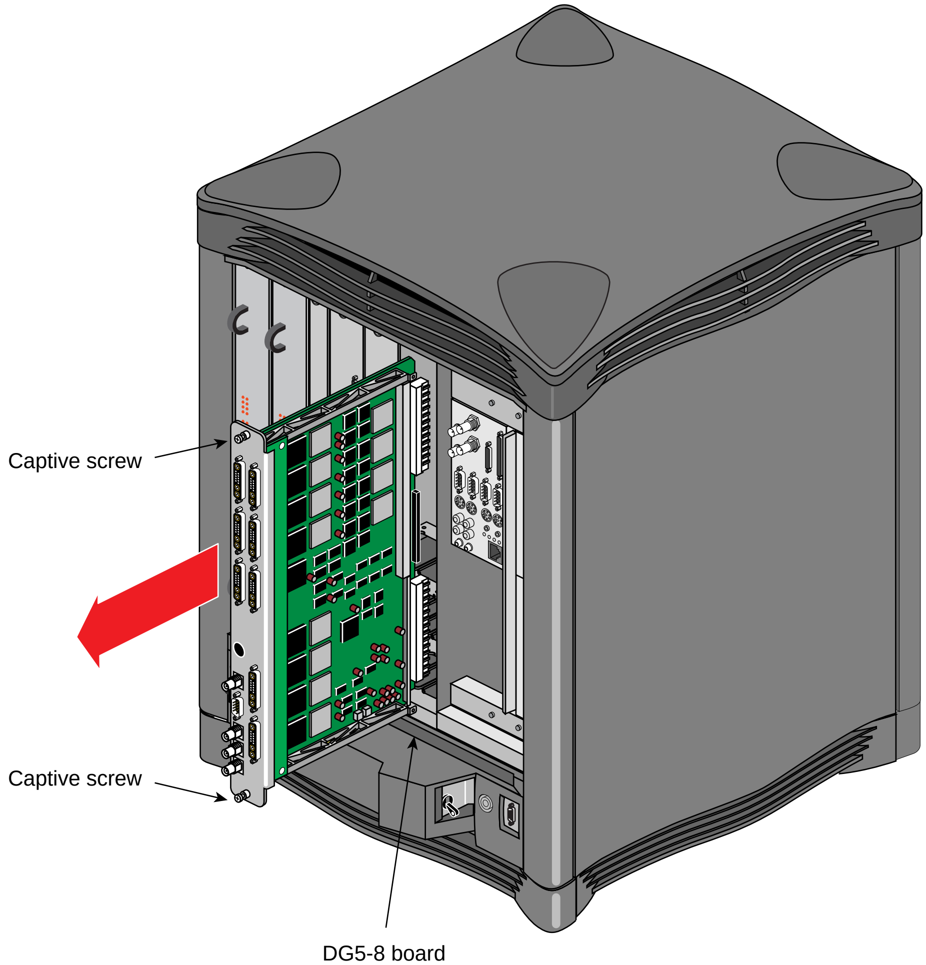

Figure 10-55 Graphics Board Removal and Replacement (DG5)

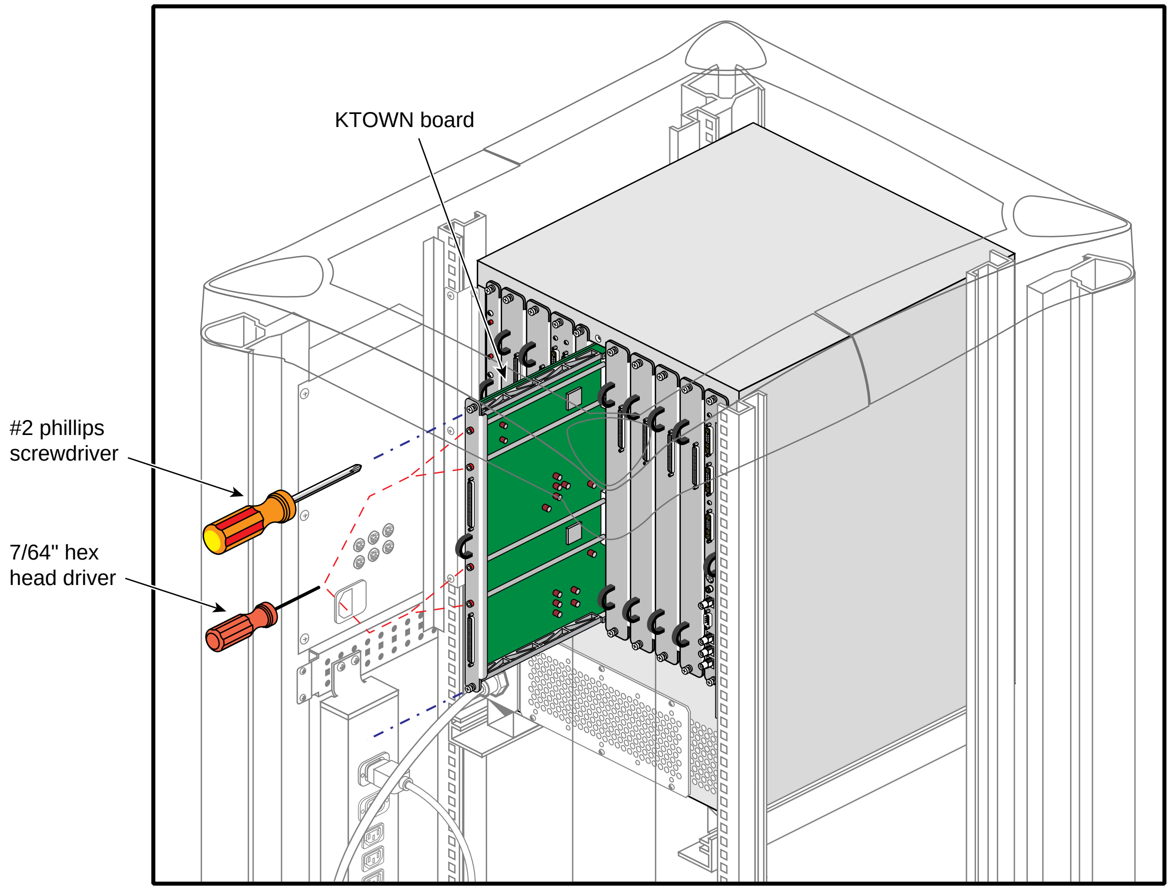

Figure 10-56 Graphics Board Removal (Ktown Board)

10.4.1.1 Onyx2 Rackmount Graphics Board Removal and Replacement

The graphics module can hold two graphics pipes (pipe 0 and pipe1). The IRIX kernel normally recognizes pipe 0 as the 4-RM board set (the graphics boards to the right of the Ktown board as you face the back the chassis). Pipe 1 is the 2 RM board set situated to the left of the Ktown board. The Ktown board has two external ports (port 0 and port 1). Port 0 is the top connector and corresponds to pipe 0. Port 1 is the bottom connector and is assigned to pipe 1.

When installing boards, it does not matter which pipe is populated first. A crosstown cable (see "CrayLink Interconnect, Xpress Link, and Crosstown Cables" in Chapter 3) connects between the Ktown board and a Crosstown board in an XIO slot.

10.4.1.2 Pipe Numbering in the Onyx2 Rack System

There are a maximum of four pipes in an Onyx2 multirack system (see Chapter 4, "System Configurations" for information on how the pipes are configured). Unlike the previous Onyx systems with the third cardcage option, the Onyx2 pipes are not numbered and configured in sequential fashion by the operating system (for example pipe 0, pipe 1, and pipe 2 and so forth).

In the Onyx2 rack, pipe numbers are limited to just 0 and 1. This numbering is adequate for a one- or two-pipe configuration. However, for configurations with more than two pipes, it poses a potential numbering problem.

Note: As stated in Section 10.4.1.1, pipe 0 is normally the 4-RM board set and pipe 1 is normally the 2-RM board set.

In the Onyx2 scheme, this numbering problem is resolved by *grouping* pipes with corresponding modules and Node boards. See the following example.

```
Module 1, Node 1 (master node) pipe 0, pipe 1 (from one two-pipe
graphics module)
Module 1, Node 2, pipe 0 and pipe 1 (from a second two-pipe graphics
module)
```

The pipes are not formally recognized as individual entities by the operating system. This software grouping of various hardware was done to help speed up the system initialization process. Reducing the number of components that the operating system needs to individually recognize enables the system to come up faster.

These groupings are important to recognize when you interpret readouts from hwgraph or hinv. For example, to determine the number of active Node boards using hardware graph, enter:

```
# ls /hw/module/1/slot
```

```
n1 n2
```

This output indicates two Node boards in module 1.

A corresponding hinv display example shows:

```
Board IO-GRAPHICS, module 1, slot n1
         adapter KONA
                 gfxpipe number 0, SGI-InfiniteReality Graphics
                 gfxpipe number 1, SGI-InfiniteReality Graphics

Board IO-GRAPHICS, module 2, slot n2
         adapter KONA
                 gfxpipe number 0, SGI-InfiniteReality Graphics
                 gfxpipe number 1, SGI-InfiniteReality Graphics
```

This hinv example shows a four-pipe configuration—one two-pipe configuration is on module 1 and the second pair is from module 2.

10.4.2 Removing the Graphics Module Power Supply

To remove the power supply from a graphics module, see Figure 10-57 and follow these instructions:

1. Unhinge the facade from the graphics module by lifting it up and pulling it straight out.
2. Remove the two screws that hold the top power supply door in place.
3. Pull down the handle, hold the power supply from the bottom, and carefully pull it out from the chassis.

Caution: Be careful not to drop the power supply on your feet!

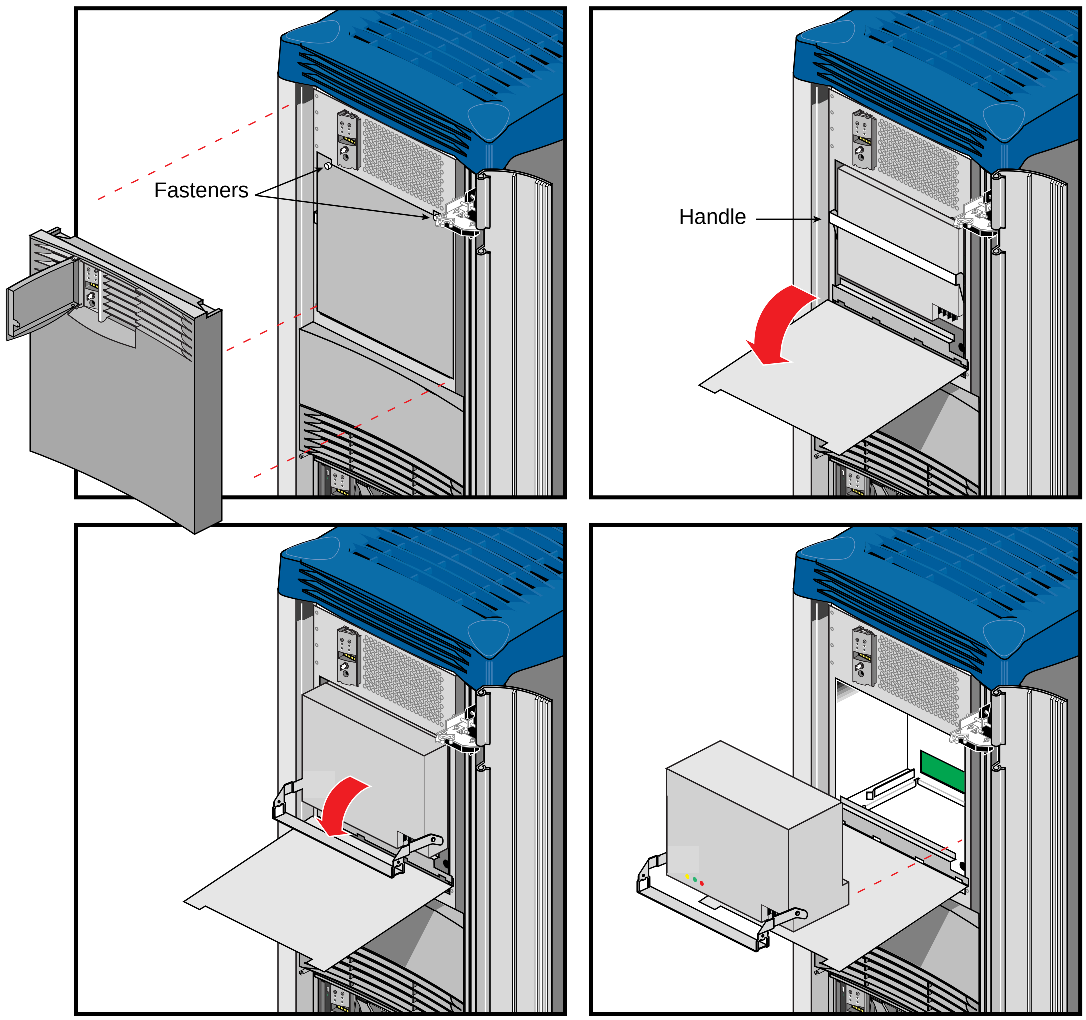

Figure 10-57 Removing the Graphics Module Power Supply

10.4.3 Removing and Replacing the Rack Graphics Module Fan Tray

The graphics module fan tray is the same type that is used in the previous-generation Onyx systems. See Figure 10-58 and follow these instructions to remove the fan tray.

1. Remove the nine screws that secure the fan tray to the back of the chassis (see Figure 10-58).
2. Pull the fan tray / blower straight out.

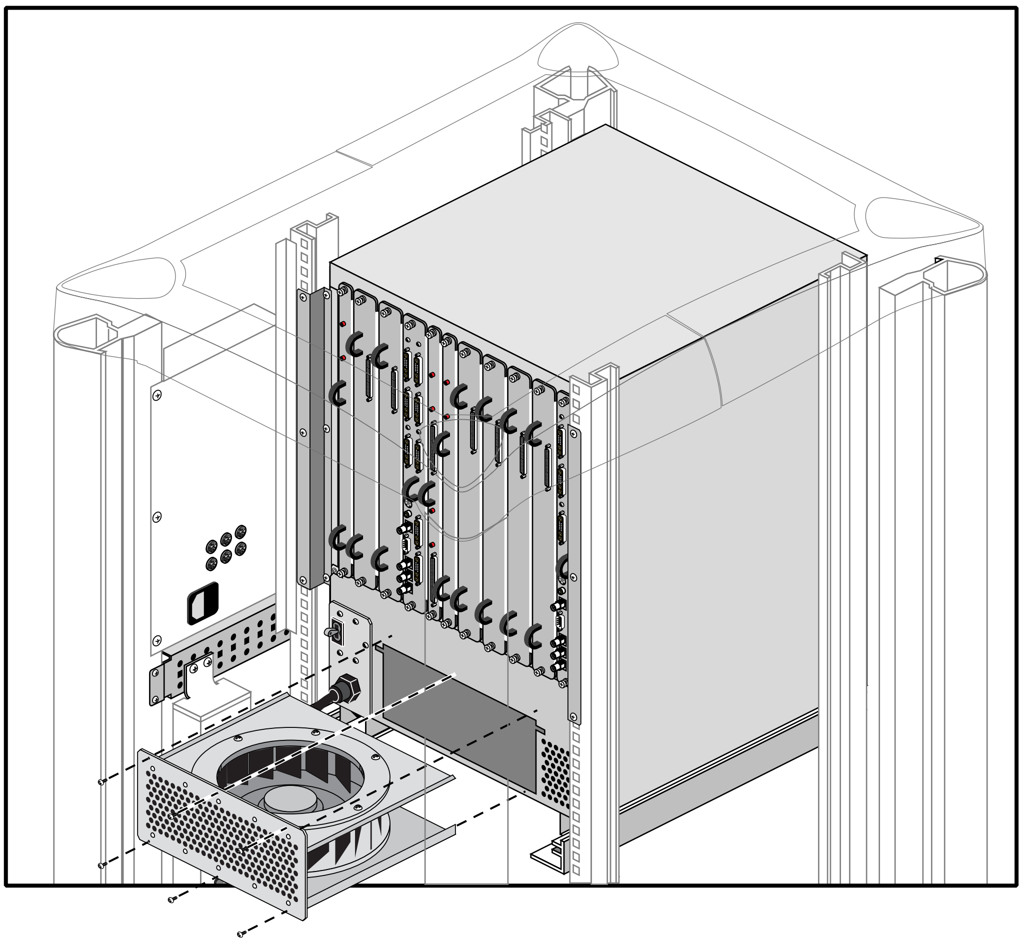

Figure 10-58 Removing the Graphics Module Fan Assembly

Appendix A

System Specifications

Table A-1 and Table A-2 provide technical specifications for the Origin2000 and Onyx2 rack system and Table A-3 and Table A-2 provide specifications for the Origin2000 and Onyx2 deskside system.

Table A-1 Rackmount Physical and Environmental Specifications

Parameter	Condition	Measurement
Dimensions		
Installed:	length	39” (99 cm)
	width	29” (74 cm)
	height	73” (185 cm)
Shipping:	length	81” (206 cm)
	width	47” (120 cm)
	height	49” (125 cm)
Weight:	minimum (empty rack)	300 lbs (136 kg)
	maximum (full rack)	750 lbs (340 kg)
	shipping (maximum)	900 lbs (408 kg)
Floor Loading:	minimum	38 lb / ft^2 (185 kg / m^2)
	maximum	95 lb / ft^2 (466 kg / m^2)
Air Temperature:	operating (< 5000 ft)	41° to 95° F (5° to 35° C)
	operating (> 5000 ft)	41° to 86° F (5° to 30° C)
	non-operating	–4° to 140° F (–20° to 60° C)
Thermal Gradient:	maximum	18° F (10° C) per hour
Altitude:	operating	10,000 ft (3,048 m) MSL, maximum
	non-operating	40,000 ft (12,192 m) MSL, maximum

Table A-2 Rackmount Electrical and Cooling Specifications

Parameter	Condition	Measurement
Voltage:		187-264 Volts, 1-phase
Watts (from-the-wall):	maximum	6500 watts
Power Factor		0.98
Inrush Current	maximum	400
Frequency		47-63 Hertz
Heat Output:		22,000 Btu / hr (1.84 ton AC load)

Table A-3 Deskside Physical and Environmental Specifications

Parameter	Condition	Measurement
Dimensions		
With skins:	length	24” (61 cm)
	width	20” (50.8 cm)
	height	26.5” (67.3 cm)
Without skins:	length	19” (48.3 cm)
	width	17.75” (45.1 cm)
	height	24” (61 cm)
Shipping:	length	31.5” (80 cm)
	width	24” (61 cm)
	height	41” (104.1 cm)
Weight:	minimum	120 lbs (54.4 kg)
	maximum	170 lbs (77 kg)
	shipping (max.)	190 lbs (86.2 kg)
Floor Loading:	minimum	36 lb / ft^2 (175 kg / m^2)
	maximum	51 lb / ft^2 (250 kg / m^2)
Air Temperature:	operating (< 5000 ft)	41° to 95° F (5° to 35° C)
	operating (> 5000 ft)	41° to 86° F (5° to 30° C)
	non-operating	–4° to 140° F (–20° to 60° C)
Thermal Gradient:	maximum	18° F (10° C) per hour

Table A-3 Deskside Physical and Environmental Specifications

Parameter	Condition	Measurement
Altitude:	operating non-operating	10,000 ft (3,048 m) MSL, maximum 40,000 ft (12,192 m) MSL, maximum
Humidity:	operating non-operating	10-90% (non-condensing) 10-95% (non-condensing)
Humidity Gradient:	maximum	10% relative humidity per hour
Acoustics:	typical	50 dBa
Vibration:	max. sustained, oper. max. peak, operating sensitive freq., oper.	5-10 Hz @ .01” total excursion, 10-500 Hz @ 0.1g 5-10 Hz @ .02” total excursion, 10-500 Hz @ 0.1g 8-33 Hz (varies with configuration)

Table A-4 Deskside Electrical and Cooling Specifications

Parameter	Condition	Measurement
Voltage:		87-264 volts, 1 phase
Wattage (from-the-wall):	Maximum at 220 volts: Maximum at 110 volts:	2600 watts 1900 watts
Power Factor:	minimum	0.98
Inrush Current:	maximum	140 amps
Frequency:		47-63 Hertz
Cooling Requirements (maximum):	At 220 volts: At 110 volts:	8,840 Btu / hr (0.74-ton AC load) 6,460 Btu / hr (0.54-ton AC load)
Wattage (from-the-wall):	Maximum at 220 volts: Maximum at 110 volts:	2600 watts 1900 watts
Power Factor:	minimum	0.98
Inrush Current:	maximum	140 amps
Frequency:		47-63 Hertz

Appendix B

Connector Pinouts

This appendix provides pinout information for the various connectors in the Origin2000 and Onyx2 system.

B.1 Common Server and Graphic BaseIO Connectors

The Origin2000 and Onyx2 systems contain the following common connectors on the BaseIO board:

- 10 / 100 Base-T Ethernet
- RS-232 or RS-422 serial ports
- single-ended, 68-pin SCSI connector

B.1.1 The Ethernet Interface Connection

The system comes with a single 10 / 100 Base-T 8-pin Ethernet connector. Optional boards supporting additional Ethernet connectors are available.

Table B-1 shows the cable pinout assignments for the Ethernet 10 / 100 Base-T Ethernet port.

Table B-1 Ethernet 100-BASE T Ethernet Port Pin Assignments

Pin	Assignment
1	TRANSMIT+
2	TRANSMIT–
3	RECEIVE+
4	(Reserved)
5	(Reserved)
6	RECEIVE–
7	(Reserved)
8	(Reserved)

Figure B-1 shows the location of the standard Ethernet connector on the rackmount system.

There are two LEDs on the RJ-45 Ethernet. The top (green) LED lights only when the system is transmitting. The bottom (yellow) LED lights whenever it sees any packet on the wire, including packets not destined for your system.

The four LEDS above the RJ-45 Ethernet connector have the following functions:

- The yellow LED on the far left (LED 1) lights to indicate SCSI activity on the BaseIO single-ended SCSCI connector.
- The green LED (LED 2) lights to indicate 100 MB-per-second packet activity.
- The yellow LED on the right (LED 3) indicates when the Ethernet is operating at full duplex rates of transfer or receive.
- The green LED on the far right (LED 4) shows the Ethernet link test. It lights when link state is valid.

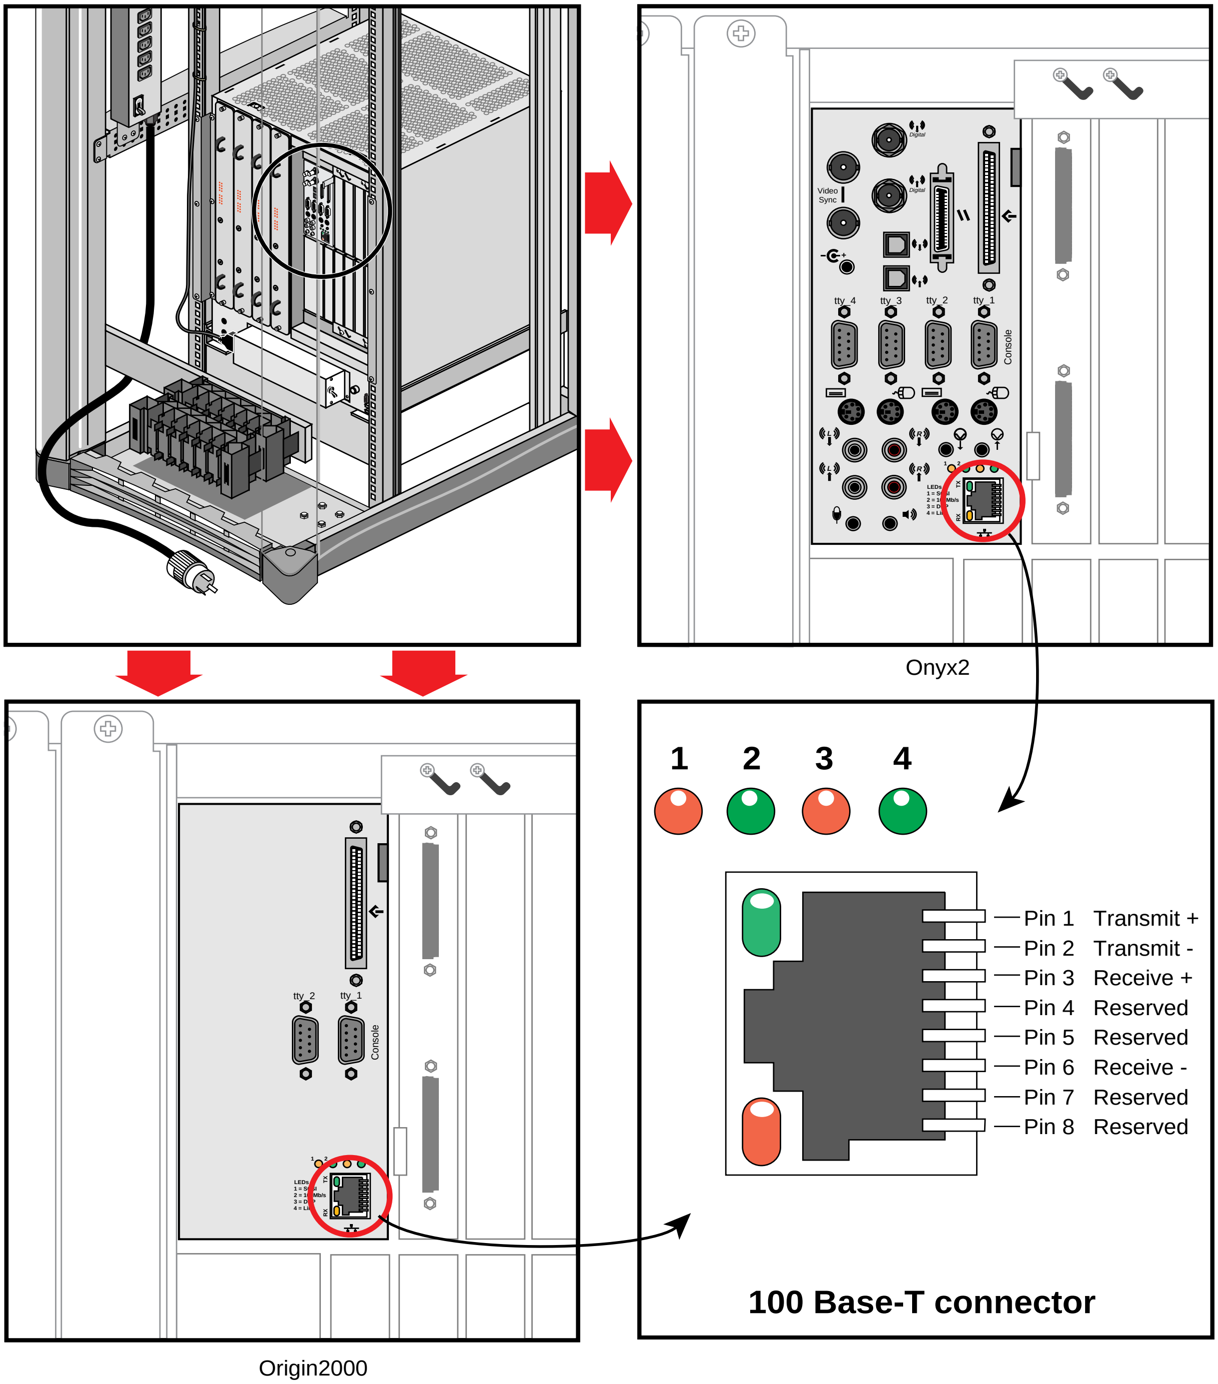

Figure B-1 Standard Ethernet

B.1.2 Standard Serial Ports

Each Origin2000 or Onyx 2 system comes with two standard 9-pin serial ports. These ports can support either RS-232 or RS-422 interface devices. Figure B-2 shows the location and pinouts for a serial port. Optional additional serial ports are also available.

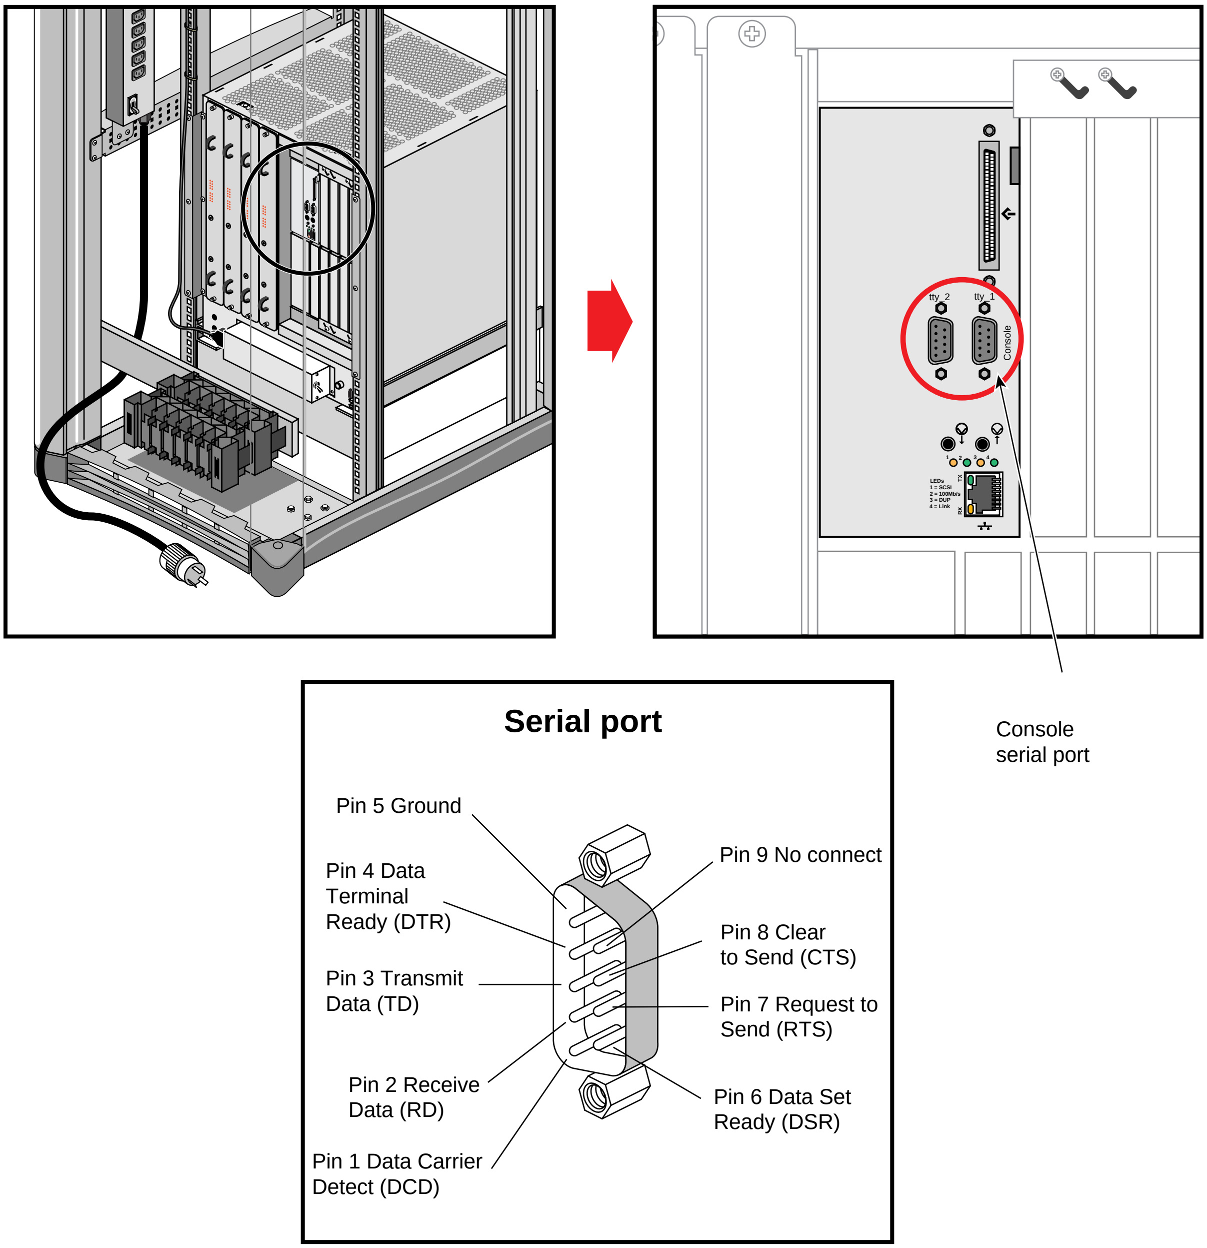

Figure B-2 Serial Port Location and Pinouts

Note: This connector departs from traditional Silicon Graphics usage:

- Pin 2 (not pin 3) is Receive Data (RD)
- Pin 3 (not pin 2) is Transmit Data (TD)
- Pin 5 (not pin 7) is signal ground
- gender is female (not male)

The RS-232 standard recommends the use of cables no longer than 50 feet (15.2 meters). This standard should also be applied to the RS-422 connector. Longer runs increases the possibility of line noise which can affect data transmission and cause errors. For cable runs longer than 50 feet (15.2 meters), use an appropriate extender device.

Do not run cables through areas that are electrically noisy, such as paces s where large electric motors, welding apparatus, or X-ray machines operate. Bury outside wiring in conduit, as lighting strikes can damage the system.

The Origin2000 and Onyx2 systems use standard PC-style serial cables for connection to an ASCII monitor and for cabling the FFSC and ELSC. You cannot use earlier Silicon Graphics serial cables. The DIN-8 is identical to the Indy DIN-8. The DB9 is PC standard; standard PC serial cables work with Origin2000 and Onyx2 systems. Table B-2 summarizes the differences between the DIN-8 connectors for various systems.

Table B-2 DIN-8 Connectors for Various Systems

Item	Challenge	Indy	Indigo/Indigo^2	Origin2000/Onyx2 System Controller (Rear)
Pin 1	DTR	DTR / HSKo	DTR / HSKo	DTR
Pin 2	CTS	CTS / HSKi	CTS / HSKo	CTS
Pin 3	Stereo sync	TD / TxD-	TD / TxD	TXD
Pin 4	RD	GND / GND	GND / GND	GND
Pin 5	TD	RD / RxD-	RD / RxD	RXD
Pin 6	GND	RTS / TxD+	RTS / TxD+	RTS
Pin 7	GND	DCD / GPI[a]	DCD / GPI**	DCD
Pin 8	10 V supply	GND / RxD+	GND / RxD+	GND
Speed	38.4 Kb	38.4 Kb	38.4 Kb	9600b

a. General Purpose Input

B.1.3 The Standard SCSI Connector

A single, external 68-pin SCSI connector is provided on the BaseIO panel (see Figure B-3). This connector supports both Ultra SCSI and SCSI-2 devices. The connector sends single-ended SCSI signals only.

Optional additional SCSI ports can be implemented using XIO option boards.

The hyphen preceding a signal name indicates that the signal is low. Note that 8-bit devices that connect to the P cable leave the following signals open: -DB(8), -DB(9), -DB(10), -DB(11), -DB(12), -DB(13), -DB(14), -DB(15), -DB(P1). All other signals are connected as shown in Table B-3.

Table B-3 68-pin SCSI Pin Assignments

Signal Name	Pin Number	Pin Number	Signal Name
Ground	1	35	-DB(12)
Ground	2	36	-DB(13)
Ground	3	37	-DB(14)
Ground	4	38	-DB(15)
Ground	5	39	-DB(P1)
Ground	6	40	-DB(0)
Ground	7	41	-DB(1)
Ground	8	42	-DB(2)
Ground	9	43	-DB(3)
Ground	10	44	-DB(4)
Ground	11	45	-DB(5)
Ground	12	46	-DB(6)
Ground	13	47	-DB(7)
Ground	14	48	-DB(P)
Ground	15	49	Ground
Ground	16	50	Ground
TERMPWR	17	51	TERMPWR
TERMPWR	18	52	TERMPWR
Reserved	19	53	Reserved
Ground	20	54	Ground
Ground	21	55	-ATN
Ground	22	56	Ground
Ground	23	57	-BSY

Table B-3 (continued) 68-pin SCSI Pin Assignments

Signal Name	Pin Number	Pin Number	Signal Name
Ground	24	58	-ACK
Ground	25	59	-RST
Ground	26	60	-MSG
Ground	27	61	-SEL
Ground	28	62	-C / D
Ground	29	63	-REQ
Ground	30	64	-I / O
Ground	31	65	-DB(8)
Ground	32	66	-DB(9)
Ground	33	67	-DB(10)
Ground	34	68	-DB(11)

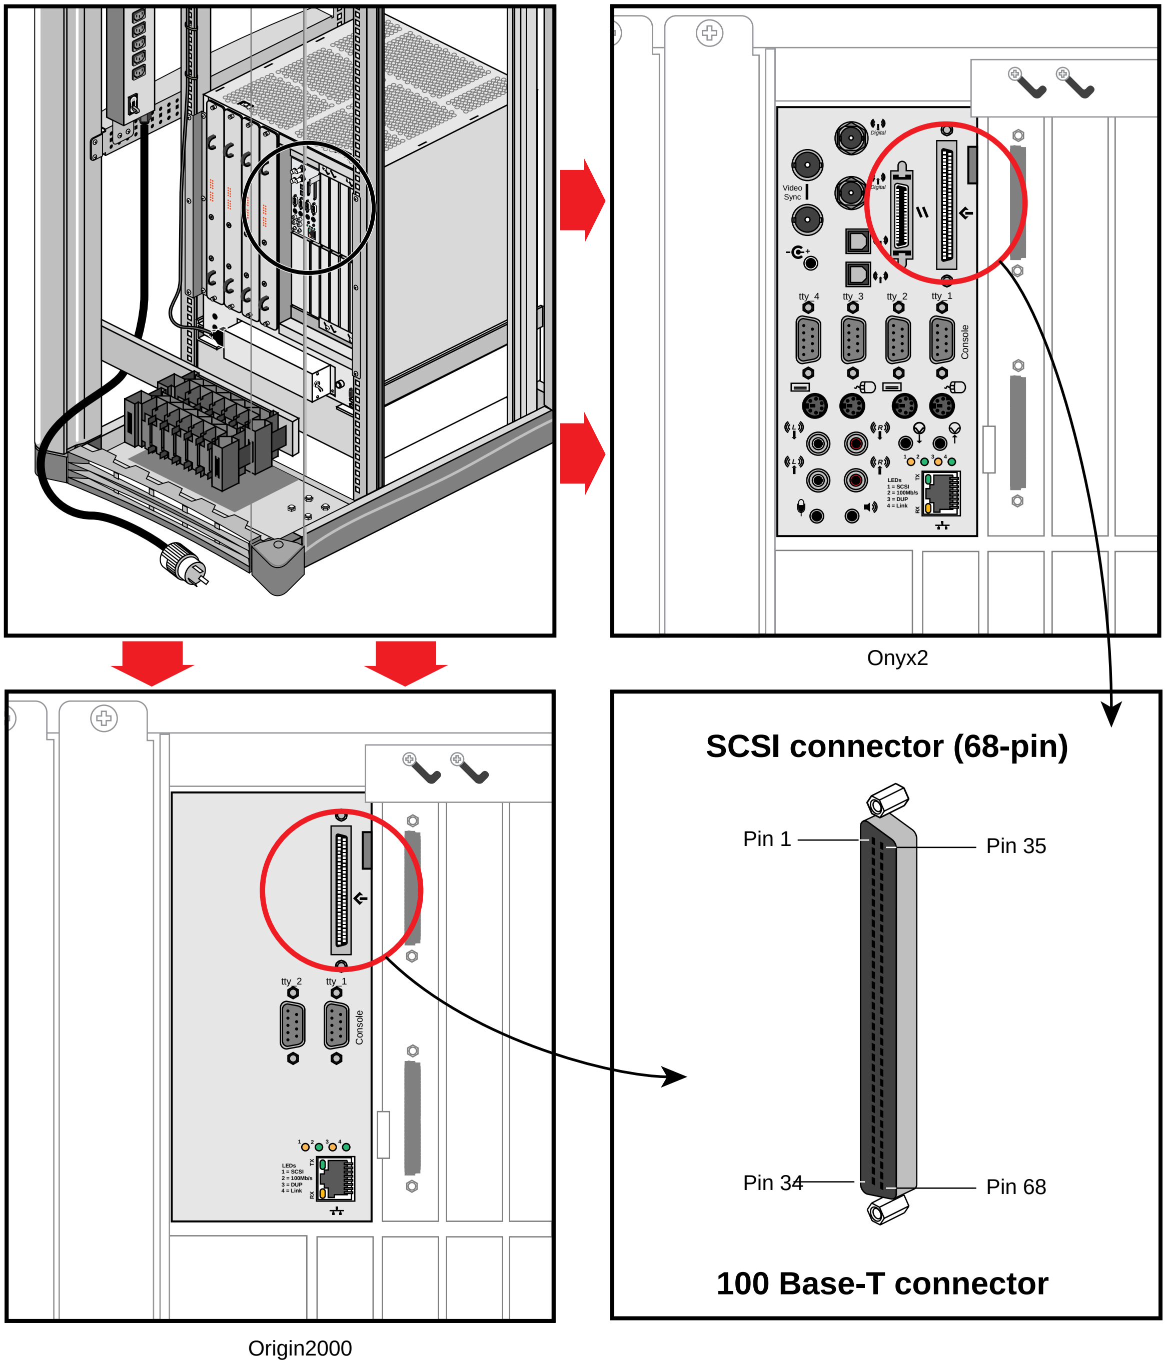

Figure B-3 68-Pin SCSI Connector

B.2 The Graphics BaseIO Interface Panel

The graphics BaseIO assembly (IO6G) is a graphics-oriented set of interface connectors that comes standard with each Onyx2 system. The connectors on the BaseIO include support for Ethernet, two keyboards, analog and digital audio, serial and parallel connectors, and others. The following sections provide location and pinout information on these connectors.

B.2.1 The Parallel Port Connector

The BaseIO board supports one IEEE 1284-C 36-pin parallel port connector. The location of this connector is shown in Figure B-4. Pinouts for the parallel port connector are listed in Table B-4.

Suitable cables for use with this port should be marked “IEEE 1284-compliant.” For most parallel printers, you can use a cable with an IEEE 1284-C connector at the Onyx2 end and an IEEE 1284-B connector (also known as a Centronics-style) at the printer end.

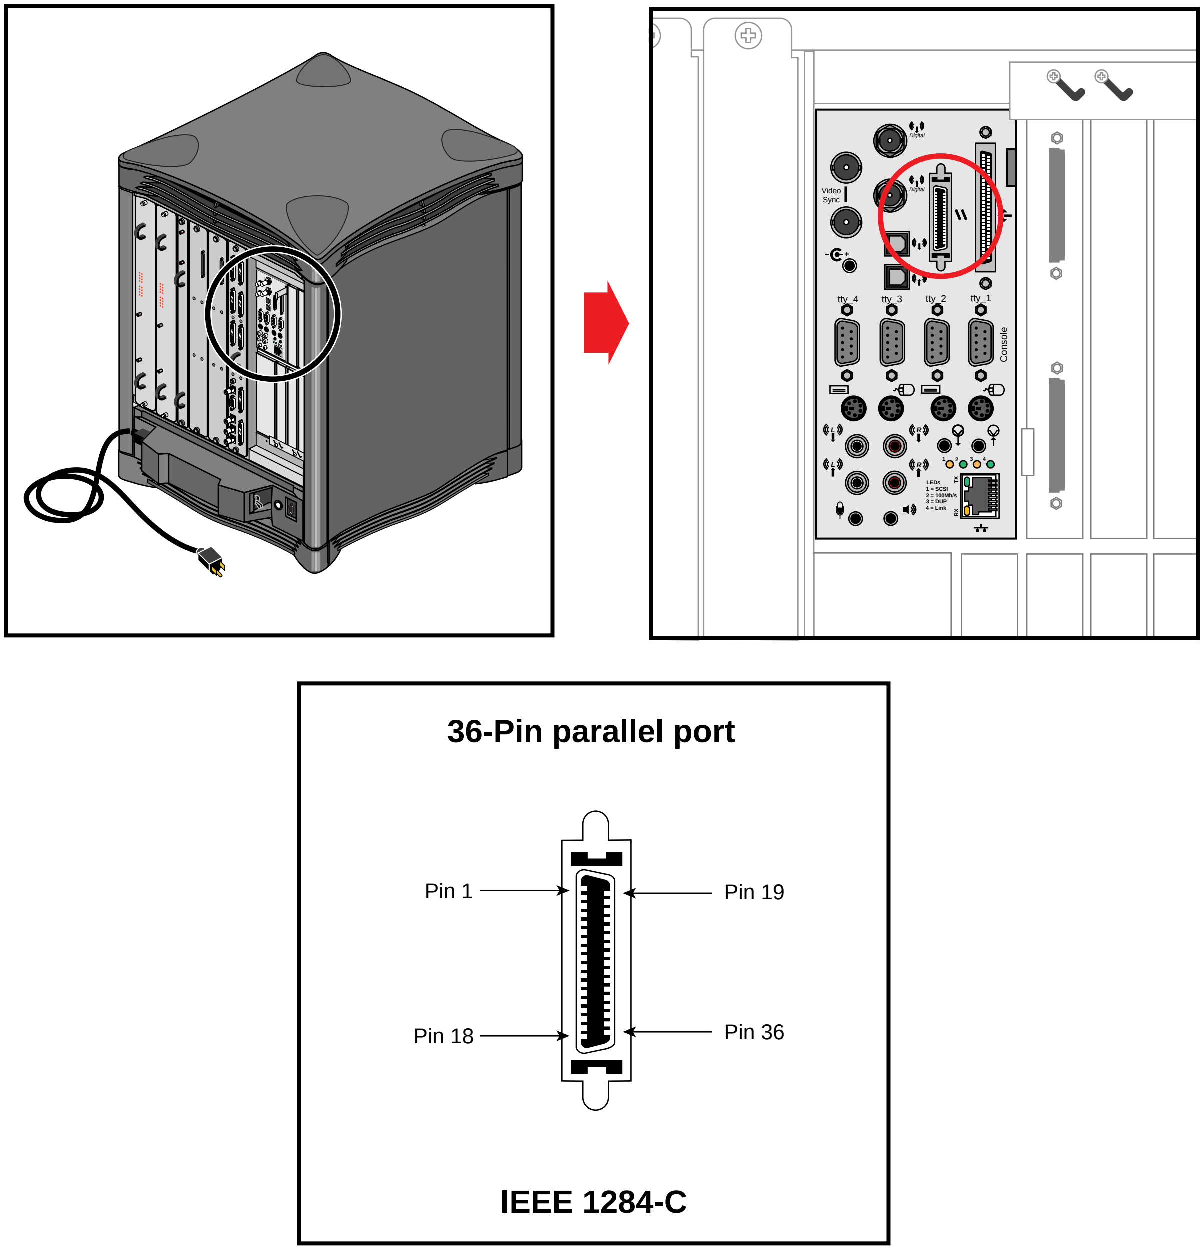

Figure B-4 Parallel Printer Port Location

Table B-4 Pinouts for the 36-Pin Parallel Port Connector

Pin	Signal	Source
1	Busy	Printer
2	Select	Printer
3	nAck	Printer
4	nFault	Printer
5	PError	Printer
6	Data 1 (LSB)	Bidirectional
7	Data 2	Bidirectional
8	Data 3	Bidirectional
9	Data 4	Bidirectional
10	Data 5	Bidirectional
11	Data 6	Bidirectional
12	Data 7	Bidirectional
13	Data 8 (MSB)	Bidirectional
14	nInit	Host
15	nStrobe	Host
16	nSelectIn	Host
17	nAutoFd	Host
18	Host Logic High	N / A
19	Signal ground (Busy)	N / A
20	Signal ground (Select)	N / A
21	Signal ground (nAck)	N / A
22	Signal ground (nFault)	N / A
23	Signal ground (pError)	N / A
24	Signal ground (Data 1)	N / A
25	Signal ground (Data 2)	N / A
26	Signal ground (Data 3)	N / A
27	Signal ground (Data 4)	N / A
28	Signal ground (Data 5)	N / A
29	Signal ground (Data 6)	N / A
30	Signal ground (Data 7)	N / A

Table B-4 (continued) Pinouts for the 36-Pin Parallel Port Connector

Pin	Signal	Source
31	Signal ground (Data 8)	N / A
32	Signal ground (nInit)	N / A
33	Signal ground (nStrobe)	N / A
34	Signal ground (nSelectIn)	N / A
35	Signal ground (nAutoFd)	N / A
36	Peripheral logic high	Printer

B.2.2 Mouse and Keyboard Ports

Each Onyx2 deskside workstation comes with two keyboard and mouse connectors. Figure B-5 shows the location of the connectors and their pinouts.

There are two sets of keyboard and mouse connectors on the rear of the graphics BaseIO panel provided with each Onyx2 system. If your system uses one keyboard and mouse, attach them to the primary keyboard and mouse connector ports. These primary ports are located on the right side of the BaseIO panel.

You can plug the keyboard and mouse cables directly into the BaseIO panel. However, in cases where your monitor, keyboard, and mouse are located away from the system, use the included extension cable. Each system comes with a 24-foot (7.3 m) keyboard and mouse extension cable.

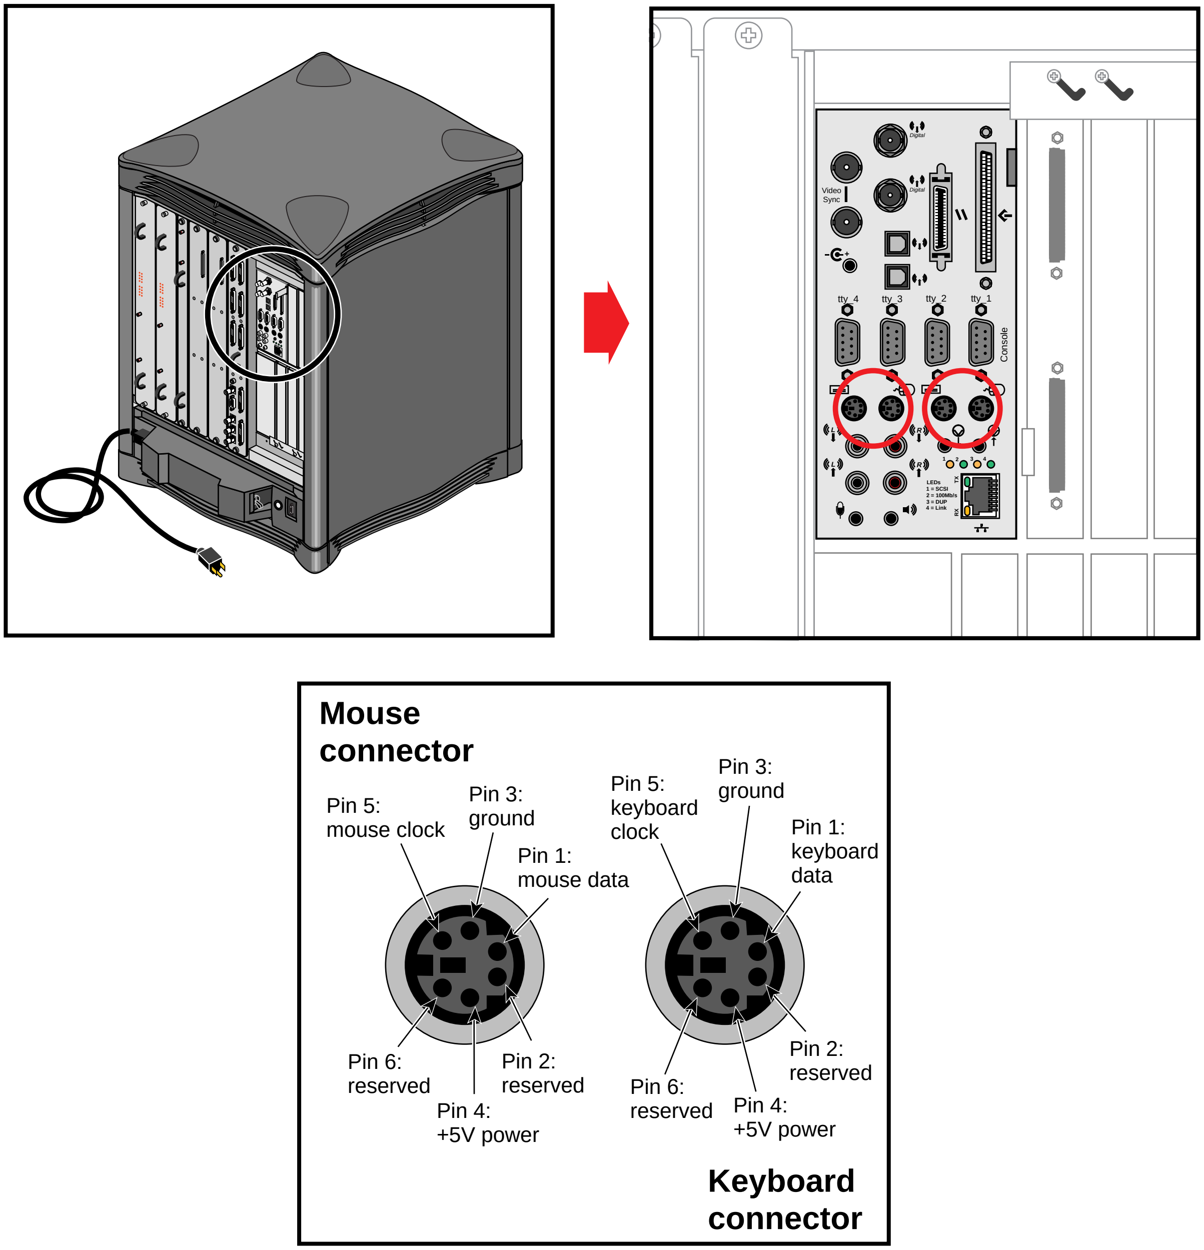

Figure B-5 Keyboard and Mouse Locations and Pinouts

Table B-5 shows the cable pinout assignments for the keyboard port.

Table B-5 Keyboard Port (6-Pin MINIDIN) Pin Assignments

Pin	Assignment
1	KEYBOARD DATA
2	(Reserved)
3	GROUND
4	KEYBOARD POWER (+5 V)
5	KEYBOARD CLOCK
6	(Reserved)

Table B-6 shows the cable pinout assignments for the mouse port.

Table B-6 Mouse Port (6-Pin MINIDIN) Pin Assignments

Pin	Assignment
1	MOUSE DATA
2	(Reserved)
3	GROUND
4	MOUSE POWER (+5 V)
5	MOUSE CLOCK
6	(Reserved)

B.2.3 Analog Stereo In and Out (RCA-Type) Ports

Table B-7 shows the cable pinout assignments for the line level audio (RCA-type) ports.

Table B-7 Analog Composite Video Port Pin Assignments

Pin	Assignment
(sleeve)	GROUND
(tip)	Line level audio

You may connect audio equipment to the line level inputs and outputs using standard shielded RCA-type connectors (see Figure B-6). For best results, always route these analog signal cables away from power cords.

The right channel is color coded red, and the left channel is white.

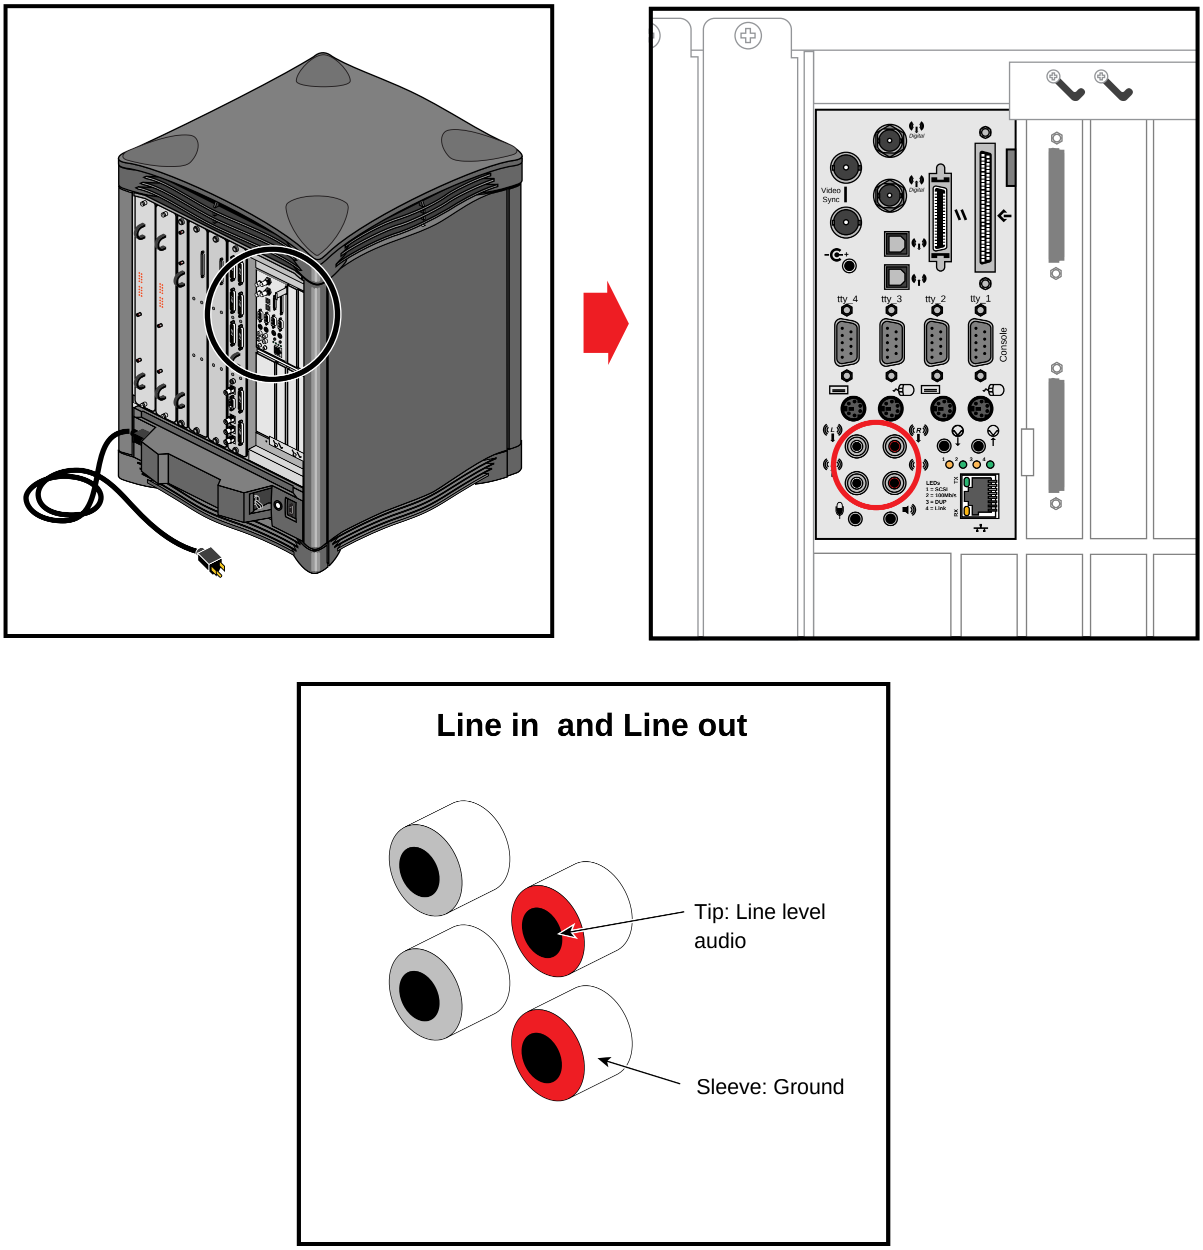

Figure B-6 Analog Stereo Port Locations

B.2.4 Optical Digital Audio Interface Connectors

Just above serial port two are the single-jack ADAT optical connectors. These ports can be used with multi-track digital audio recording input and output devices.

These connections support optical input and output of eight channels at up to 24 bits and up to 48 Hz sample rates.

Use standard plastic fiber interconnecting cables. You will need two cables; one for input and one for output. The Onyx2 system ships with connector cover plugs over the input and output ports. These must be removed before using the optical connectors. Retain these dust covers for use when shipping or if you discontinue ADAT use.

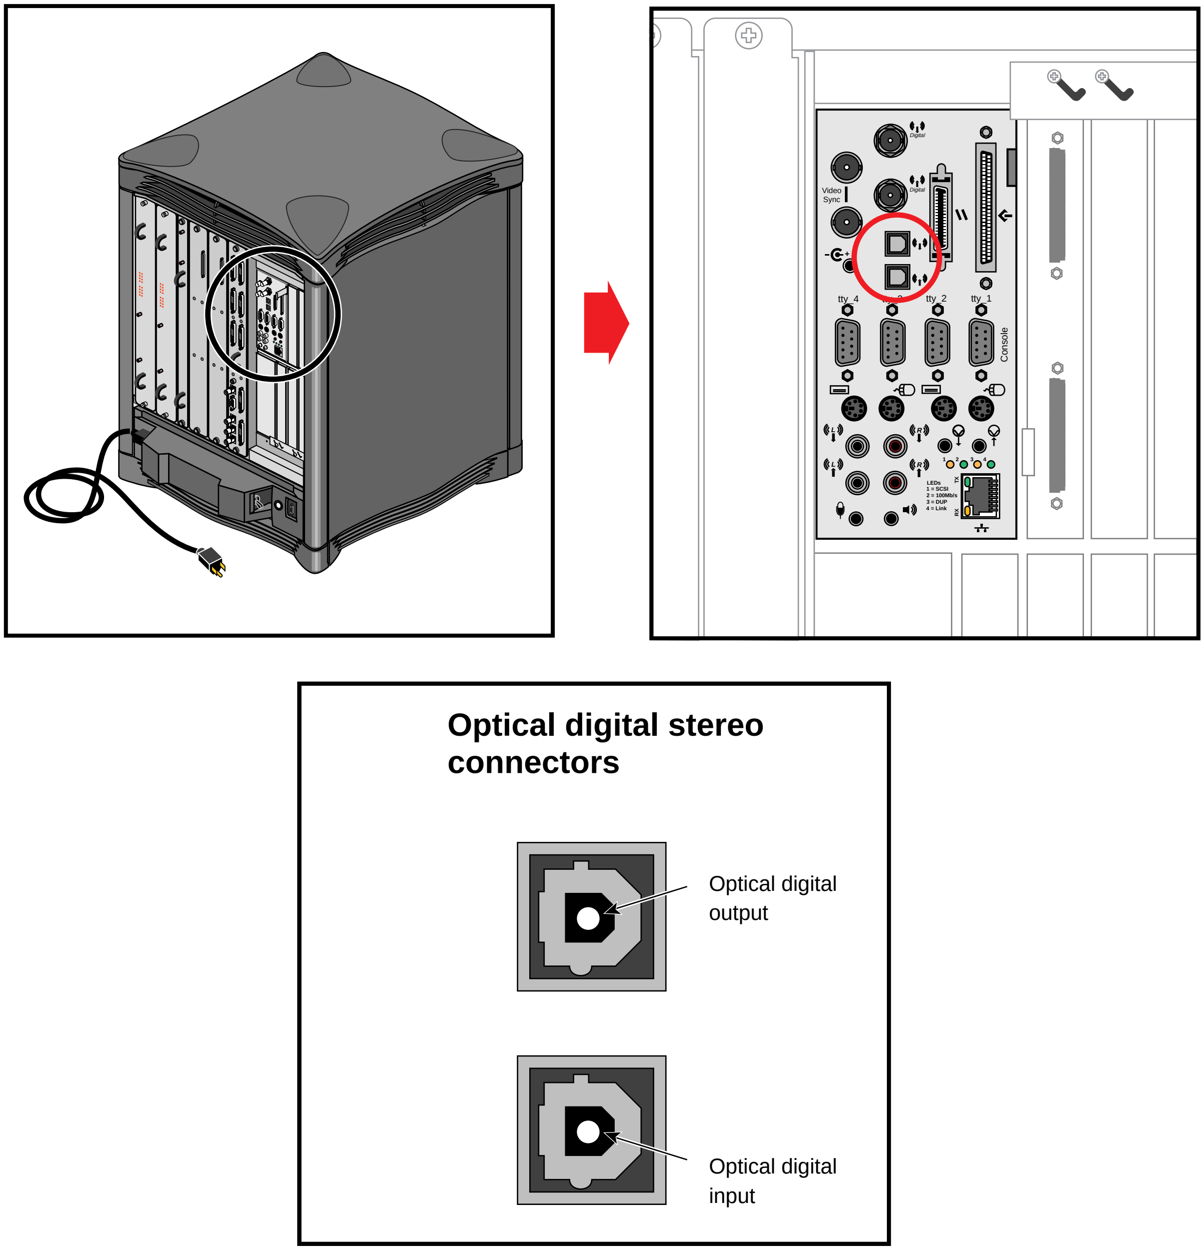

Figure B-7 Optical Digital Audio Interface

B.2.5 Loopthrough and Digital Audio Connectors

Figure B-8 shows the loopthrough and digital audio connectors.

The AES and EBU digital audio connectors support 75 ohm signals at a nominal 1.0 volts (peak-to-peak) signal level. You should use 75 ohm coaxial cable with standard BNC connectors for interconnections (such as with digital video recorders).

Some equipment supporting AES and EBU digital audio signals uses 3-pin XLR connectors that support balanced 110 ohm signals. To successfully interconnect with equipment of this type, install a digital audio “BALUN” adapter at the equipment’s XLR connector points. The BALUN adapter connects the 3-pin XLR to a 75-ohm BNC connection. The 75 ohm coaxial cable then connects between the Onyx2 system and the BALUN adapter.

Note: 110 to 75 ohm digital audio BALUN adapters come in male and female versions. You need one of each type when using both the input and output AES and EBU connectors.

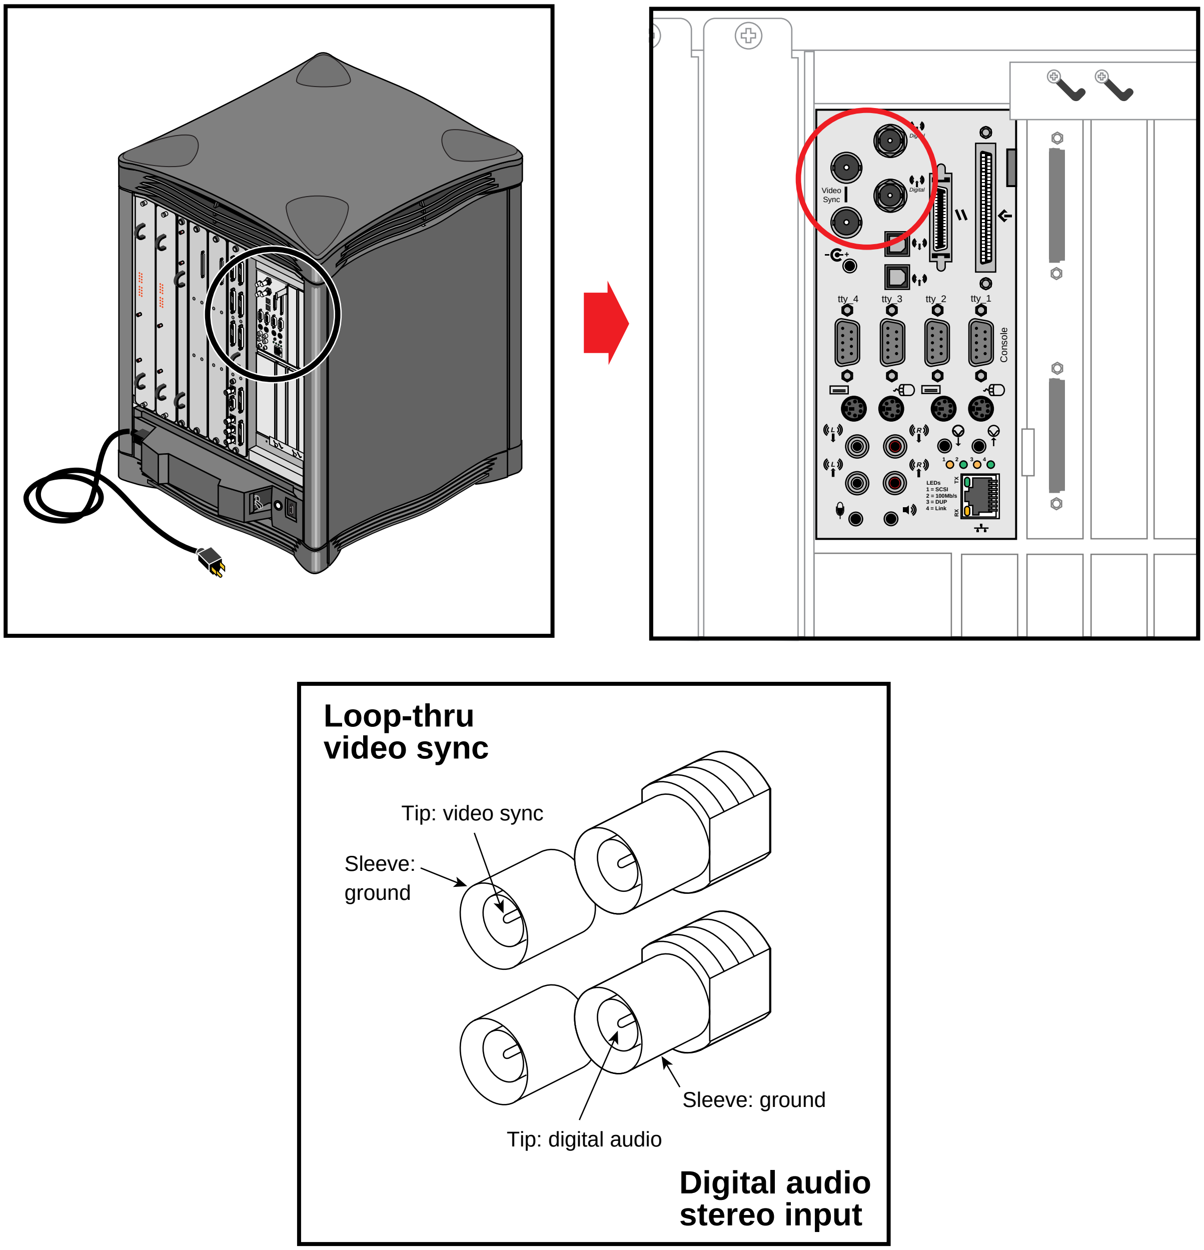

Figure B-8 Loopthrough and Digital Audio Connectors

B.2.6 Speaker and Microphone Connections

The Onyx2 BaseIO panel uses a 30-foot (9.1 m) three-connector bundled cable to make connection to a microphone and a pair of speakers (included with your Onyx2 system).

Figure B-9 shows the connection points on the rear of the speakers.

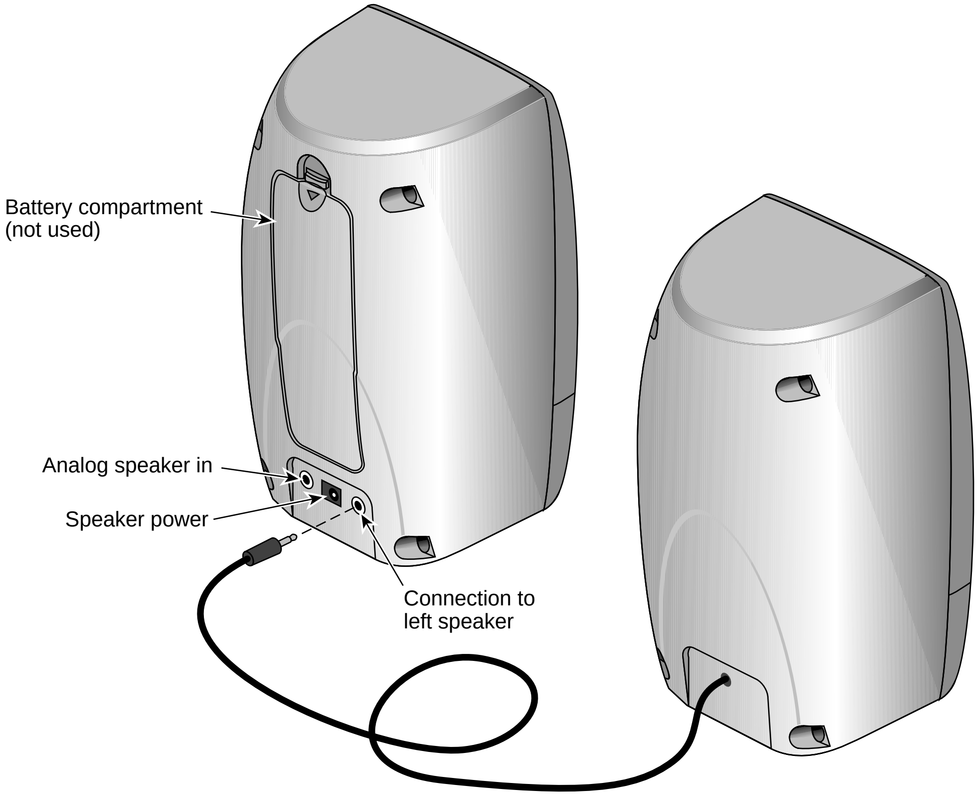

Figure B-9 Cable Connection Locations on the Speakers

The analog speaker connector plugs into the BaseIO board and the other end goes to the analog speaker plug.The analog speaker power connector goes from the BaseIO to the middle plug on the back of the right speaker (see Figure B-10).

The microphone connector goes from the BaseIO panel and you plug in your (included) microphone at the other end (look for the microphone logo on the connector). Note that there is also a 10-foot (2 m) “speaker only” cable included with your Onyx2 system. This 10-foot cable does not supply a microphone plug and limits where you may place the speaker set.

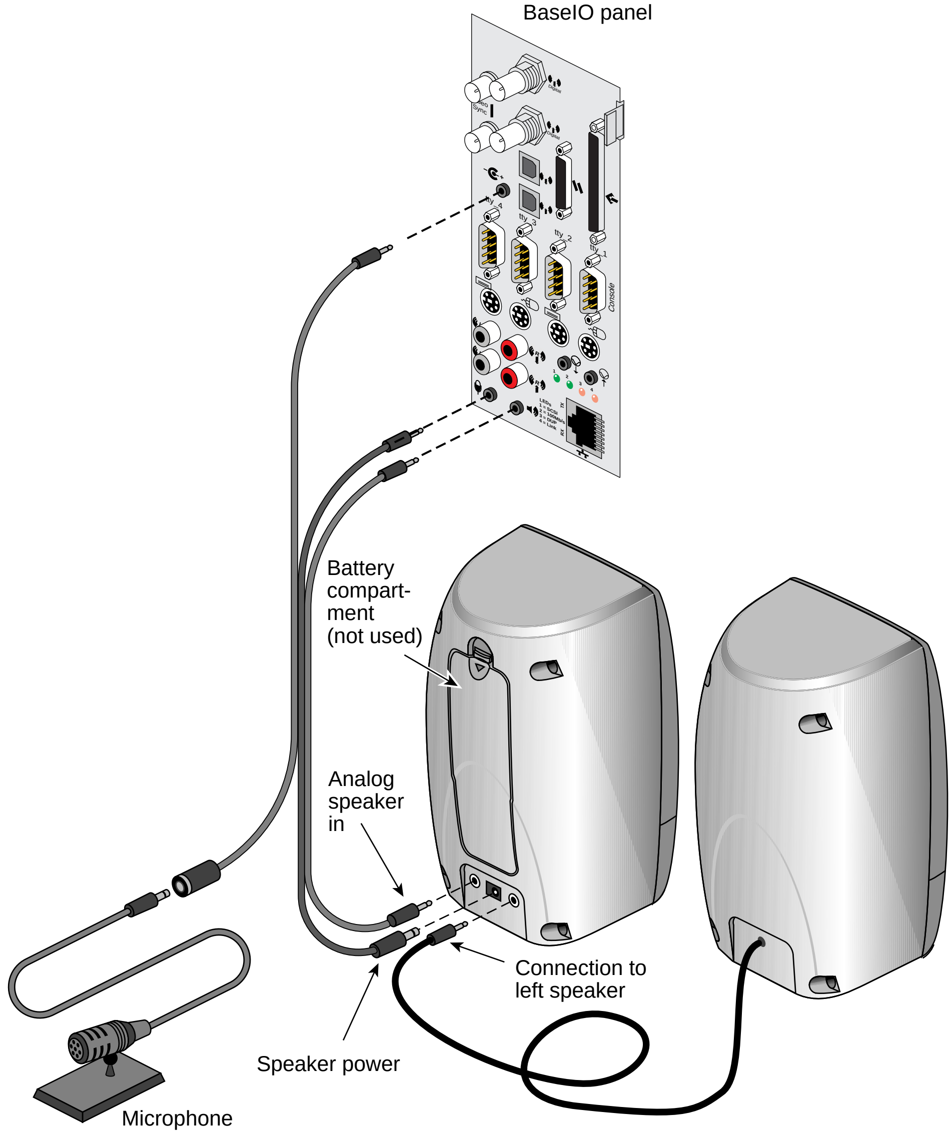

Figure B-10 Speaker and Microphone Connections to the BaseIO

Appendix C

MMSC Command Language

The multimodule system controller (MMSC) controls a single Origin2000 or Onyx2 rack system.* When multiple racks are connected to form a single system, each rack has its own MMSC; these MMSCs are then connected via a private Ethernet.

Note: The MMSC is also known as the full-feature system controller (FFSC).

Each MMSC can accept commands from a number of different sources, including consoles connected via direct serial connection or modem, individual modules (bays) in a rack via their module system controllers (MSCs), and other MMSCs via the private Ethernet. In addition, one rack in an Origin rack system has a display with several buttons that can be used to generate commands.

The commands that are issued to an MMSC can directly affect the operation of the MMSC or can be passed along to individual bays, either in the same rack as the MMSC or in a different rack. This appendix describes the syntax and contents of these MMSC commands.

Note: When referring to a rack, the term module sometimes means any module in a system, and it sometimes means just one of the two modules in a specific rack in a system. To avoid ambiguity, this appendix uses the word module when the logical definition (that is, any module in a system) is intended, and uses the term bay when the physical definition (a module in a specific rack) is intended.

This appendix explains

- how to specify destinations
- command syntax
- command set

* This information is from a Web page written by Rob Bradshaw.

C.1 Specifying Destinations

Many commands are intended for delivery to specific racks (MMSCs) or bays (MSCs), so it is important to have a way to specify particular destinations.

At the lowest level, racks are addressed with small integers starting with 1, and MSCs within a rack (bays) are addressed with letters that represent their position: U (upper) and L (lower) for the Origin rack.

One or more rack or bay addresses can be combined in a list. A list of rack addresses consists of one or more addresses or ranges separated by commas, with no intervening white space. A range, in turn, is a series of contiguous values specified as two addresses separated by a hyphen, again with no intervening white space. The following examples are valid lists of rack addresses:

```
1
1,3
1-f
1-4,7,39-3b
```

A list of bay addresses also consists of one or more individual addresses separated by commas, with no intervening white space. For bay addresses, comma separators are optional. However, ranges are not meaningful for lists of bay addresses, and so are not allowed.

C.1.1 Physical Destinations

A physical destination refers to one or more specific bays without regard to any logical module or partition designations. These destinations are typically used for maintenance commands, such as powering off a single bay to replace a board. A complete physical destination specification consists of one or more rack / bay pairs; for example:

rack `rlist` **bay** `blist`

The keywords **rack** and **bay** can be abbreviated as **r** and **b**, respectively.

A rack / bay pair selects each of the bays in *blist* on each of the racks in *rlist* (somewhat like a product). Specifying more than one pair extends the selection accordingly (somewhat like a sum). Thus, the long destination

```
rack 1-3 bay u,L    rack 4 bay L
```

is equivalent to

```
rack 1 bay u
rack 1 bay l
rack 2 bay u
rack 2 bay l
rack 3 bay u
rack 3 bay l
rack 4 bay l
```

Instead of specifying a list of addresses for *rlist* or *blist*, you can also use one of several keywords, as summarized in Table C-1.

Table C-1 Address List Keywords

Keyword	Abbreviation	Use
all	*	Selects all known *online* addresses. Skips addresses that are believed to be offline. If the MMSC's recorded online or offline addresses seem to be inaccurate, use the *scan* command (see below) to update them. A command can still be directed to a supposedly offline address by explicitly specifying it rather than by using the *all* keyword.
local	.	This keyword is valid only for an *rlist*. Selects the local rack, that is, the rack that contains the MMSC that initially accepts the command.

If the rack portion of a physical destination is omitted but the bay portion is not, then *rack local* is implied. If the bay portion is omitted but the rack portion is not, then *bay all* is implied. Therefore, the following examples are all equivalent in a four-rack system with two bays per rack:

```
rack all bay ul
rack 1,2-4
r * b *
```

C.1.2 Logical Destinations

A logical destination refers to individual modules by their module numbers, rather than their physical positions in a system. Logical destinations are used more commonly than physical destinations because IRIX and the various PROMs use these types of addresses. A logical destination consists of the keyword *module* followed by a module list:

```
module mlist
```

The keyword *module* can be abbreviated as *m*.

As with physical destinations, mlist can be replaced with the keyword *all* to specify all known modules. There is a special case in which no modules are defined (this might occur, for example, if none of the modules have been powered on). In this case, *module all* is equivalent to *rack all bay all*.

Many MMSC commands pertain only to a particular rack and not to a specific bay within that rack. If a logical address is used with such a command, the bay address that is implied by the module number is ignored.

Both physical and logical destinations may be specified for the same command, as long as one is not embedded in another. Thus, the following addresses are valid:

```
rack * bay u module 3,5-7
r 1 b ul    m 8   rack 3 bay l
```

However, the following addresses are not valid or they may not give the expected results:

```
rack module 2 3 bay u    (invalid)
r 3 m 5 b u   (valid, but equivalent to:  r 3 b *   m 5   r . b u)
```

When the MMSC executes a command, logical destinations are converted internally to physical destinations using the module number mappings that the MMSC has at the time that the command is executed. If the MMSC's mappings are no longer valid (for example, if a module has been powered off or is otherwise unavailable), then commands may time out or be directed to the wrong module. The scan command can be used to update the MMSC's mappings manually if necessary.

C.1.3 Default Destinations

Although you can indicate a specific destination for any command, doing so can be cumbersome in some cases, such as when several commands go to the same complicated list of destinations. Therefore, the MMSC allows you to specify a default destination by entering a destination with no additional command or by using the *dest* command. Any subsequent commands that expect a destination use this default value if no other destination is specified.

The initial default destination is *module all*.

C.2 Command Syntax

Other than the destination specification, the general syntax for MMSC commands is fairly simple. In general, an MMSC command looks like this:

[*escape*] [*dest*] *command* [*args*] <Enter>

where

escape — is an MMSC escape character, Ctrl+T by default. In console and MSC pass-through modes, this character indicates the beginning of an MMSC command. See Section C.3 for further details.

dest — is a destination specification.

command — is the name of an MMSC command, as described in "Command Set." If *command* is not a valid MMSC command, it is assumed to be an MSC command. Along with *args*, it is passed to the bay(s) addressed by *dest* without further translation.

args — are zero or more arguments to *command*.

Enter — terminates a command; commands are always terminated with a carriage-return character. Note that neither the command nor its arguments can contain an embedded carriage-return character.

If an MMSC escape character is entered within an MMSC command, all previous characters since the last MMSC escape character are discarded. This feature can be useful if the current state of a console or other MMSC connection is currently unknown. The kill character (Ctrl+U by default) has the same effect.

C.3 Command Set

This section describes the valid MMSC commands. If a command other than one of these is specified, it is assumed to be an MSC command and is passed along to the addressed MSC. Many commands accept an optional destination specification, while others do not. Unless otherwise specified, each of these commands can be executed by a user at the basic authority level.

authority [*level* [*pw*]]

If *level* is not specified, then the authority level of the console from which the command was entered is displayed.

If *level* is specified, the authority level of the console from which the command was entered changes to *level*. The valid authority levels are *basic*, *supervisor*, and *service*. For more information on these, see *Security on the Origin2000 Multi-Module System Controller* (http://uniscan.engr.sgi.com/~rdb/FFSC/security.html).

If the *supervisor* or *service* level is selected and this *level* has a password, then the appropriate password (*pw*) must be specified after *level*. If the password is not correct, an error message displays and the current authority level does not change.

This command is only meaningful when it is entered from terminals.

bs char

bs **?**

Sets the backspace character to *char*, which can be a single literal character, an integer value that represents a single ASCII character, or a control sequence that consists of a caret followed by a single character (for example, ^H). Note that this only affects the backspace character in MMSC or MSC commands. It does not affect the backspace character used in normal pass-through operation. The default backspace character is Ctrl+H.

If *bs* **?** is specified, the current backspace is printed. Thus, in order to use the **?** character as the actual backspace character, you must specify it using its numeric ASCII code.

This command is only meaningful when it is entered from terminals.

cecho [**on** | **off**]

Controls the echoing of characters when an MMSC command is typed. When echo is *off*, any characters that are received after an MMSC escape character are buffered internally by the MMSC but are not echoed back to the input source. When echo is *on* and an MMSC escape character is received, an MMSC prompt and the character that follows the escape character are echoed to the input source. All additional input characters up to the carriage return are echoed as well. Once the carriage return is received, input processing returns to its original state.

If *cecho* is entered without arguments, the current *cecho* mode is toggled. The default setting for *cecho* is *on*. This command is meaningful only when it is entered from terminals.

com port

com port **cmd on/off**

com port **function** *func*

com port **oob on/off**

com port **rxbuf/txbuf** *value*

com port **speed** *baudrate*

Sets or displays communications settings of the various serial ports on the MMSC.

port is a number from 1 to 6 that corresponds to a particular serial port on the addressed MMSC, as shown in Table C-2.

Table C-2 Port Definitions

Port	Label	Default Function
1	console	terminal
2	upper bay	upper
3	lower bay	lower
4	BaseIO tty1	system
5	alternate console	altcons
6	test	debug

If no other arguments are specified, the current communication settings for the selected port are displayed.

The *cmd* subcommand indicates whether or not MMSC and MSC commands are accepted from a port. If **off** is specified, then MMSC and MSC commands are not accepted from the port. For ports associated with a terminal (ports with functions terminal or altcons), this causes the MMSC escape character to be ignored when the port is in the CONSOLE input mode. This can be useful if a particular console is in an insecure location, for example.

For ports associated with the system (such as functions system or daemon), this command disables the Out Of Band function FFSC_COMMAND. If *on* is specified, then MMSC and MSC commands are accepted from the port. The only way to place a port that has specified *cmd off* into some other input mode is to “steal” the console from a different port (see the *steal* command description). Thus, be very careful that at least one console specifies *cmd on*. The default setting for ports 1 and 5 is *cmd on*, and for ports 2, 3, 4 and 6 is *cmd off*.

The *function* subcommand is used to set the function associated with the selected port to *func*. Table C-3 lists valid functions.

Table C-3 Valid Functions

func	Function of the Associated Port
terminal	Communication with the user terminal device.
upper	Communication with the MSC in the upper bay of the rack.
lower	Communication with the MSC in the lower bay of the rack.
system	Communication with the operating system. This is the port through which both the user and the MMSC communicate with IRIX. It is typically connected to the TTY1 port on the master BaseIO card.
altcons	Remote service port. This is typically connected to a modem that is used for communication with an SGI service center.
daemon	Communication with a system controller daemon on IRIX, such as ffscd. Such a daemon would ordinarily be used for generating bar graph data on the MMSC display. This port would typically be connected to a second serial port on BaseIO card.
debug	MMSC debugging log. This is mostly useful if some sort of MMSC error has occurred. In that event, the debugging log may contain additional information.

A function can be assigned to only one port. If func specifies a function that is already assigned to another port, then the other port will have its port changed to an “unassigned” state. Ensure that, at minimum, the TERMINAL function is assigned to a port, or else it may become impossible to communicate with the MMSC.

The *oob* subcommand indicates whether out-of-band (OOB) data received from a port should be intercepted and processed. OOB messages are used by programs such as ffscd to display bar graphs and other information on the MMSC display. The actual OOB message protocol is not described in this document. If on is specified, then OOB message processing will be done for data received by the port. For ports associated with the system (ports with functions SYSTEM or DAEMON), any data preceded by the out-of-band prefix character will be interpreted as an OOB message. The MMSC will perform the requested action and respond as specified in the OOB message protocol. For ports associated with a terminal (namely ports with functions TERMINAL or ALTCONS), OOB message processing consists of simply doubling the out-of-band prefix character before sending it to the system. This prevents it from being interpreted by IRIX as the beginning of an out-of-band message. If off is specified, then out-of-band messages that are received from the port are ignored and passed through without alteration. The default setting for all ports is oob off.

The *rxbuf* and *txbuf* subcommands are used to change the size of the serial port’s receive and transmit buffers, respectively. The default for both is 4096. If the system and terminal ports have different speeds, it may be necessary to increase the size of the transmit buffer on the slower port or else data may be lost. In extreme cases, serial buffer overflows may take a serial port out of service, so avoid underestimating the buffer size.

The *speed* subcommand is used to set the baud rate of the selected port to baudrate. Valid baud rates include 300, 1200, 2400, 4800, 9600, 19200, 38400, 57600, and 115200.

Function changes do not take effect until the MMSC is reset (see reset_mmsc). Speed changes take effect immediately. The user must have service authority to specify the *cmd*, *function*, *oob,* or *speed* subcommands.

This command accepts an optional destination specification.

console [*args*]

cons [*args*]

If *args* are present, this command sends them to the system console port and otherwise ignores them. *args* is assumed to start at the first nonblank character following the command.

If *args* is not specified, console pass-through mode is entered. Any subsequent input is passed through to the system console port and any output from that system console port is echoed here. There are two exceptions:

1. Any input preceded by the MMSC escape character, up to and including the first subsequent carriage-return character, is processed by the MMSC rather than being passed through. Any output that is generated by this MMSC command may be discarded or echoed, depending on the current *rmsg* setting.
2. Two MMSC escape characters in a row cause a single MMSC escape character to be sent to the system console port, then otherwise be ignored by the MMSC.

dest

Displays and / or modifies the current default destination. If no address was specified prior to the dest keyword, then the console's current default destination is displayed. Otherwise, the specified address becomes the new default destination for the console on which the command was entered. Any subsequent command that specifies no address will automatically use the default destination instead.

direct

Enter direct input mode on both the current console and the "other" console on the MMSC. If the device on which the command is entered happens to be the normal terminal device, then the "other" console is the alternate console device. Likewise, if the command is entered from the alternate console device, the "other" console is the normal terminal device.

All input from the current console is sent directly to the other console and vice versa. The only exception is that typing the exit character on the current console causes both consoles to leave direct mode and return to their usual input modes (console mode for the terminal and RAT mode for the alternate console).

Direct mode is most useful for programming a modem that is attached to the "other" console.

A user must have service authority to execute this command.

downloader **disable**

downloader **enable**

This command sets or clears the “serial downloader” flag in the addressed MMSC’s initialization PROM. The command *downloader enable* sets the flag, and the command *downloader disable* clears it.

When the “serial downloader” flag is set, the initialization PROM on the MMSC’s single-board computer loads the “serial downloader” rather than the normal MMSC firmware. The serial downloader can be used to load a new firmware image into the MMSC’s flash RAM.

Note that once the MMSC is initialized with the “serial downloader” flag set, there is no way to return to normal MMSC operation until a new MMSC firmware image is loaded. If you inadvertently set the “serial downloader” flag on the wrong rack, use the *downloader disable* command to reset the flag as soon as possible. If an MMSC is reset, either with the *reset_mmsc* command or because it has lost power, you must follow the procedure in Using the Serial Downloader to restore normal operation.

The *downloader enable* command is typically used to place an MMSC that does not have an attached display into serial downloader mode. A rack that does have a display can be placed into serial downloader mode by holding down the bottom two buttons of the display (ENTER and DOWN) while the MMSC is powered on.

This command accepts an optional destination specification.

A user must have service authority to execute this command.

end char

end **?**

Sets the end-of-command character to *char*, which can be a single literal character, an integer value that represents a single ASCII character, or a control sequence that consists of a caret followed by a single character (for example, ^T). In the console or MMSC input modes, the end-of-command character signals the end of an MMSC command. The default end-of-command character is Ctrl+M (carriage return).

If *end* **?** is specified, then the current end-of-command character is printed. Thus, in order to use the **?** character as the actual end-of-command character, you must specify it using its numeric ASCII code.

This command is meaningful only when it is entered from terminals.

esc char

esc **?** — Sets the MMSC escape character to *char*, which can be a single literal character, an integer value that represents a single ASCII character, or a control sequence that consists of a caret followed by a single character (for example, ^T).

The MMSC escape character signals the beginning of an MMSC command while in console mode. The default MMSC escape character is Ctrl+T. Note that changing the MMSC escape character does not affect the MSC escape character, which is always Ctrl+T.

If *esc* **?** is specified, the current MMSC escape character is printed. Thus, in order to use the **?** character as the actual MMSC escape character, you must specify it using its numeric ASCII code.

This command is meaningful only when it is entered from terminals.

exit

exit char

exit **?** — If neither *char* nor **?** is specified, this command leaves the current input mode and returns to the console mode. If you are already in console mode, this command has no effect.

In MSC and MMSC mode, typing the exit character has the same effect. The only way to leave MSC mode is by typing the exit character, because MMSC commands are not accepted in MSC mode.

If *char* is specified, this command sets the exit character to *char*, which can be a single literal character, an integer value that represents a single ASCII character, or a control sequence that consists of a caret followed by a single character (for example, ^E).

The default exit character is Ctrl+E. An alternative might be Control-D, the standard Unix EOF character that is often used to exit shells.

If *exit* **?** is specified, the current exit character is printed. Thus, in order to use the **?** character as the actual exit character, you must specify it using its numeric ASCII code.

flash [**from system**]

flash **from console**

Flashes a new firmware image into non-volatile storage on the MMSC for the addressed rack. An error occurs if more than one rack is addressed by this command. The bay portion of the destination is ignored.

In the first form, *flash from system* (the words “from system” are optional), the image is provided to the MMSC from the console of a running IRIX system using the flashmmsc(1m) command. The flashmmsc IRIX command and the flash MMSC command must both be issued from the same terminal device. Typically, the flashmmsc command is issued first with the -m option, then an MMSC-escape character is typed and the flash command entered from the same terminal.

In the second form, *flash from console*, the image is read from the terminal device itself. This is useful when the user terminal is in fact a terminal emulator (perhaps on a PC) capable of XMODEM file transfer. For this form of the flash command, issue the flash command to the MMSC first, then arrange for an XMODEM or XMODEM-1K file transfer of the firmware image from the user terminal. The use of XMODEM-1K may improve the transfer time by as much as 50%.

Be forewarned that this command takes a long time to run. In addition to the time to download the firmware image over a serial line (MMSC firmware images are approximately 1MB in size) 100 seconds is required for clearing out the flash RAM prior to installing the new image. Every effort should be made not to interrupt the MMSC while this command executes. Early versions of the MMSC do not have sufficient flash storage to hold two separate images; therefore an attempt to flash a new image will necessarily wipe out the old firmware image, even if the attempt is unsuccessful. Do not reboot an MMSC that has tried and failed to flash a new firmware image. If this does occur, there is an emergency procedure to flash new firmware onto it if necessary. For more information on this procedure, see Flashing MMSC Firmware.

This command accepts an optional destination specification.

A user must have service authority to execute this command.

help [*cmd* | **ALL**]

Displays information about the MMSC commands. If no argument is specified, a list of available commands is printed. If a command is specified, more specific information about that particular command is printed. If *ALL* is specified, specific information about all available commands is printed.

This command is meaningful only when it is entered from terminals.

kill char

kill **?** Sets the kill character to *char,* which can be a single literal character, an integer value that represents a single ASCII character, or a control sequence that consists of a caret followed by a single character (for example, ^U). The kill character is used while typing an MMSC / MSC command in console mode to cancel any characters that were typed since the last MMSC escape character. In MMSC mode, the kill character cancels any characters that were typed since the last prompt was printed.

The default kill character is Ctrl+U.

Note that this command does not affect the normal UNIX character that is associated with the console when it is in pass-through mode.

If *kill* **?** is specified, the current kill character is printed. Thus, in order to use the **?** character as the actual kill character, you must specify it using its numeric ASCII code.

This command is meaningful only when it is entered from terminals.

log [*log*] **clear**

log [*log*] [**dump**] [*num*]

log [log] **disable** | **enable**

log [*log*] **info** | **?**

log [log] **lines** *num*

log [*log*] **length** *num*

Manipulate the various message logs maintained by the MMSC.

The *log* argument specifies which log is to be manipulated. It may be one of the following values:

Table C-4 Log Argument Values

log	Function
msc	A log of all messages and other output from the addressed MSC.
system	A log of all output generated by the Base I / O board on the addressed rack, typically consisting of output from the operating system.
terminal	A log of all output to the main terminal console that is attached to the addressed rack. This is different from the SYSTEM log in that it may contain output from the MSC and MMSC input modes, as well as any messages received from an MSC (typically generated by various PROM's during system initialization).
altcons	A log of all output to the alternate console that is attached to the addressed rack.
debug	A log of all debugging messages produced by the addressed rack's MMSC.
display	A log of each command issued by the MMSC display and the corresponding response.

If *log* is not specified, the default is msc.

The *clear* subcommand causes the contents of the specified log to be discarded. Logging continues as usual unless it is disabled with the *disable* subcommand.

The *dump* subcommand causes the contents of the specified log to be dumped on the current terminal. The data is displayed as paged output (see pager for details). If *num* is specified, then only the last num lines of the log are dumped. This is the default subcommand if none is specified.

The *disable* subcommand stops the logging of data to the specified log. Any data that is currently in the log remains unchanged. The *enable* subcommand resumes logging operations.

The *info* subcommand shows information about the specified log, such as its size and its enable / disable state. The *?* subcommand does the same thing.

The *lines* subcommand sets the maximum size of the specified log to num lines, where a "line" is defined as a sequence of characters that ends with either CR / LF or LF / CR. The actual number of lines that a log can hold depends on the average line length (see below). The change will not take effect until the MMSC is reset (see reset_mmsc). If default is specified for num then the actual value is taken from the LOG_DFLT_NUMLINES environment variable.

The *length* subcommand sets the average line length of lines in the buffer to num. The total amount of storage set aside for log data is the product of the log's lines and length values. If the length value is too small, then the log might wrap before it contains the maximum number of lines. If the length value is too large, it wastes storage. A change to the length value does not take effect until the MMSC is reset (see reset_mmsc). If default is specified for num then the actual value is taken from the LOG_DFLT_LINELEN environment variable.

The *lines*, *length*, and *enable / disable* settings for each log are separate from the other logs. They are saved in NVRAM and are restored after the MMSC is reset.

This command accepts an optional destination specification.

A user must have service authority to execute the *lines*, *length*, *enable,* or *disable* subcommands.

mmsc [*args*] If *args* are present, they are sent along as a command to the MMSC on the addressed rack(s). *args* may be a valid MMSC command only; MSC commands are not recognized in this case and will result in an ERROR CMD response. This usage is useful for avoiding ambiguity with MSC commands. It can also be used to send a normally “local” command (e.g. esc) to a remote MMSC. args is assumed to start at the first non-blank character following the command.

If *args* is not specified, enter MMSC mode. In MMSC mode, the MMSC prompt:

MMSC>

is displayed, and output from the IO6 is held and / or discarded. Any input from the terminal keyboard is handled directly by the local MMSC, using the address specified with the *mmsc* command as the default destination. To leave MMSC mode, use the command exit or type the current exit character.

Ordinarily, the MMSC escape character has no special meaning in MMSC mode and is processed like any other character. However, if the first character of a line is the MMSC escape character, it is discarded. This allows automated tools to type the same sequence of characters in order to issue an MMSC or MSC command from either MMSC or console mode.

This command accepts an optional destination specification.

mmsg [**on** | **terse** | **off**]

mmsg rack [*rackid*]

mmsg altrack [*rackid*]

The first form of this command is only meaningful when it is entered from a terminal, and it ignores any destination. The second and third forms of this command do honor a destination. mmsg controls the display of unsolicited messages from the MSC (e.g. those sent to the MMSC in response to an MSC “acp” command). It only affects the console on which it was entered. When mmsg is *off*, unsolicited MSC messages are discarded silently. When mmsg is *on*, all unsolicited MSC messages are echoed to the input source on their own lines (i.e. preceded and followed by carriage-return / linefeed as needed), prefixed with a string that identifies the originating MSC. The prefix is of the form:

```
R<rack-addr><bay-addr>-
```

For example:

```
R1U-This is a message from rack 1 bay U
```

When mmsg is *terse*, unsolicited MSC messages are still echoed on their own line, but the identifying prefix is omitted.

If mmsg is entered without arguments, the current mmsg mode is cycled in some undefined, but consistent, order.

If mmsg is entered in the second form, then any unsolicited MSC messages generated by the addressed racks are sent to the TERMINAL device that is attached to the rack specified by rackid. A rackid of none causes MSC messages to be discarded, and so is functionally equivalent to mmsg off.

Because an MMSC is not able to detect the presence of a BaseIO card on a particular rack, it is necessary to run this command at least once when the system is first set up in order to receive unsolicited MSC messages. The command would typically be of the form "rack all mmsg rack rackid" mmsg rack without a rackid returns the ID of the rack that is currently designated to receive unsolicited MSC messages.

The third form of mmsg is similar to the second form except that it specifies a rack whose ALTERNATE console port will receive unsolicited messages. The rackids specified for rack and altrack do not have to be the same. Either or both may be none.

The default mmsg mode is off. The default rack to receive unsolicited MSC messages on both the TERMINAL and ALTERNATE ports is 1.

This command accepts an optional destination specification.

msc [*args*]

If *args* are present, this command sends them to the MSC on the addressed bay(s), but otherwise ignores them. They are prefixed with the MSC escape character Ctrl+T. This command can be useful for forcing a command to be passed on to the MSC, such as when the command has the same name as an MMSC command. *args* is assumed to start at the first non-blank character following the command.

If *args* is not specified, MSC mode is entered. MSC mode is (supposed to be) functionally equivalent to having a direct connection to the MSC port: all keyboard input is echoed directly to the addressed MSC, and all output from the MSC is echoed to the user terminal without modification. There are two exceptions:

- The current exit character causes an exit from MSC mode and returns you to the previous input mode. (There is no way to send an exit character to the MSC in MSC mode.)
- Any output from the MSC that begins with a Ctrl+T character is assumed to be a message to the MMSC; all characters up to and including the next carriage return are handled by the MMSC on the same rack as the MSC.

When the *msc* command is used this way, it addresses exactly one MSC, or an error occurs.

Note that when an MSC is attached to a console with the *msc* command, that MSC is no longer available to perform commands issued from other consoles. For example, if the alternate console is in MSC mode, then any MSC commands issued from the normal terminal or the display are rejected. The error message in this case is typically ERROR INUSE.

In MSC mode, output from the system is held or discarded. Messages from other MSCs are displayed only after a carriage return / line feed sequence is echoed to the user terminal, or if the terminal is inactive for some period of time (approximately 2 seconds).

This command accepts an optional destination specification.

nap_time [*value* | **default**]
If *value* is specified, then the nap interval for the console on which the command is entered is set to *value* microseconds. The nap interval is the frequency with which the MMSC will attempt to perform various clean-up tasks when the console is otherwise idle. The most visible effect of these clean-up tasks is that any incomplete messages that were received from an MSC are printed. (An "incomplete message" is one that is not terminated with a CR / LF combination.)

If *default* is specified instead of a value, the nap interval is reset to the system default, which is normally 2 seconds.

If neither *value* nor *default* is specified, then the current nap interval is displayed.

A user must have service authority to execute this command.

This command is only meaningful from a terminal.

options [*value*] Sets the option flags to *value*. The currently defined debugging flags, which can be logically ORed in almost any combination, include:

- 0x00000001: Do not display the MMSC prompt in console mode until the character after the MMSC escape character is typed. This option can be useful if you expect to send the MMSC escape to the system (by typing it twice) from time to time.
- 0x00000002: If the end-of-command character (usually carriage return) is typed immediately after the MMSC escape character, enter MMSC mode on the local rack. This option is equivalent to typing the command *rack local mmsc*. This option can be useful for people who use *mmsc* mode frequently.
- 0x00000004: The MMSC always sends the command *ech* 0 to the MSC before it takes control of it, so when control is given to the user after entering MSC mode, echoing is not normally enabled. When this option is set, an *ech* 1 command is sent to the MSC prior to entering MSC mode. This command should cause the MSC to echo keyboard input as it is typed.

- 0x00000008: Ordinarily, blank messages from the MSC are ignored, even with mmsg on. When this option is set, blank messages from the MSC are displayed like any other message.
- 0x80000000: Sets the debugging flag and enables additional mmsc commands and features, which are known collectively as “Rob Mode.” For details, refer to http://uniscan.engr.sgi.com/~rdb/FFSC/rob_mode.html.

This command is meaningful only when it is entered from terminals.

pager {**back** | **fwd** | **quit**} *char*

pager [**info** | **?**]

pager **lines** *val*

pager {**on** | **off**} The *pager* command controls the built-in pager that the MMSC uses to display large blocks of output, such as logs or help messages. It is conceptually similar to the standard UNIX command more.

The *back* subcommand specifies the character that is typed to scroll backwards through the paged output. char may be a single literal character, an integer value that represents a single ASCII character, or a control sequence that consists of a caret followed by a single character (for example, ^U). Notice that the only way to specify a <space> character is with its integer value, 32. The default back character is b.

Similarly, the *fwd* subcommand specifies the character that is typed to scroll forwards through the output, and the *quit* subcommand specifies the character that is typed to discontinue the output. The default fwd character is <space>, and the default quit character is q.

The *info* subcommand displays information about the current pager settings. The ? subcommand does the same thing. This is the default if no subcommand is specified.

The *lines* subcommand is used to set the number of lines in a single page of output. After val lines are displayed, a trailer line is displayed and output is halted until the fwd, back, or quit character is typed. Specifying default for val restores the default value. The default setting is controlled by the PAGE_DFLT_LINES environment variable, which is typically 23 lines (leaving room for a single trailer line on a standard 24-line terminal).

The *off* subcommand turns off paged output. Any output that is ordinarily paged is instead sent to the terminal in one single, large block. The *on* subcommand reverses this.

The *back*, *fwd*, *quit* and *on*/*off* settings for each console are separate from the other consoles. They are saved in NVRAM and are restored after the MMSC is reset.

password {**set** | **setmmsc**} *passwd string*

password **unset** *passwd*

Changes the passwd password. Unless otherwise specified, you must be at the supervisor or service authority level to use this command. passwd may be one of the following values:

- msc: The password on the addressed MSC. This is the same password that is specified with the MSC *pas* command. This is passed along to the MSC before any restricted commands when the user is in the supervisor or service authority level.
- supervisor: The password used to enter the supervisor authority level.
- service: The password used to enter the service authority level. You must be at the service authority level to change this password.

The *set* subcommand changes the password passwd to string. If passwd is msc, then the password is changed on the MSC first by passing the command "pas s string" to the MSC. If the command does not succeed, then the password is not changed on the MMSC either.

The *setmmsc* subcommand is identical to the *set* subcommand, except when passwd is msc. In that case, the password is not set on the MSC first. This is useful if the password was changed on the MSC without the MMSC's knowledge (for example, while in the MSC input mode).

The *unset* subcommand is used to remove the password associated with passwd. Note that the MSC does not have a notion of "removing the password", so specifying a passwd of **msc** affects the MMSC only. In this case, no attempt is made to set or revoke the MSC supervisor mode prior to sending commands to it.

This command accepts an optional destination specification.

printenv [**all**] [**default**]

Displays the names and values of any environment variables that have non-default settings. If *all* is specified, all of the known environment variables are listed. Those with non-default values will have their values printed as well. If *default* is specified, the default value for each variable will also be displayed. The data will be displayed as paged output (see pager for details).

This command accepts an optional destination specification.

rackid [*value*] If *value* is specified, this command changes the rack ID of the addressed rack from its current setting to *value*. This setting is saved in NVRAM and should persist across power cycles and resets. The change becomes effective immediately. This feature can cause trouble if you are in MMSC or MSC mode when *rackid* is issued, because the rack ID that was previously addressed by MMSC / MSC mode is no longer valid.

If *value* is not specified, this command ignores the destination and returns the current rack ID of the local rack.

This command accepts an optional destination specification. A user must have service authority to execute this command.

rat Places the console into RAT (Remote Access Tool) mode. RAT mode is a combination of Console, MMSC and MSC input modes that is supposed to make the MMSC behave somewhat like an MSC. It is intended mainly for use by certain automated service tools. It has the following characteristics:

- User input is not echoed.
- Any input that is not preceded by the MMSC-escape character is discarded.
- User input is treated as if it was entered in MMSC mode: if it is a valid MMSC command, it is processed as such; otherwise, it is passed to the addressed MSC(s).
- Response and ELSC messages are printed in the same way as for Console mode.

reset_mmsc Restarts the addressed MMSC(s). This may be useful after flashing a new firmware image onto the MMSC so that the changes can be registered. This can also be used if the MMSC and / or system console has hung even though the system itself is still running. Of course, if it is used for this second purpose, then it is appropriate to file a bug report / service call.

This does not affect IRIX or any other part of the system. Only the addressed full-featured system controllers are rebooted.

For historical reasons, the command reboot_mmsc does the same thing as reset_mmsc.

This command accepts an optional destination specification.

A user must have service authority to execute this command.

reset_nvram Resets the contents of non volatile storage on the addressed MMSC(s) to default values. Use this command with caution.

Values that are stored in non-volatile storage include the rack ID, MMSC serial port speeds, and the mmsg rack settings.

This command accepts an optional destination specification.

A user must have service authority to execute this command.

rmsg [**on** | **error** | **off**]

Controls the echoing of responses to MSC or MMSC commands. It only affects the console on which the command was entered. When *rmsg* is *off*, the response string that is generated by an MSC or MMSC command is discarded silently. When *rmsg* is *on*, all command responses are echoed to the input source on their own lines (preceded and followed with carriage return / line feed as needed). When *rmsg* is *error*, only error responses (those not prefixed by OK or OFFLINE) are echoed.

If *rmsg* is entered without arguments, the current *rmsg* mode is cycled in some undefined, but consistent, order.

In either error or on mode, the terminal does not return to its normal pass-through mode until a response is actually received or some timeout is exceeded.

The default setting for *rmsg* is **error**. This command is meaningful only when it is entered from terminals.

scan Checks for the presence or absence of the addressed MSCs and updates the internal table of module numbers to physical address mappings. This may be necessary if a module becomes unavailable or its module number changes without the MMSC's knowledge.

Unlike most commands, if "bay all" is specified in the destination for this command, then all bays in the specified racks will be probed, even if they were originally thought to be offline. This is useful for determining whether a formerly offline MSC was brought online. "module all" behaves normally and so is not useful for finding new modules.

A user must have service authority to execute this command.

setenv var [[=] *value*]

Changes the setting of environment variable var to value. value must normally be an integer. The "=" between var and value is optional. If it is specified, there should be no whitespace between it and *var* or *value*. If *value* is not specified, then the current setting of the environment variable var is displayed.

Only environment variables known to the MMSC may be stored or displayed. Table C-5 lists valid environment variables:

Table C-5 Environment Variables

Variable	Function	Default value
DEBOUNCE_DELAY	The debounce delay (in microseconds) used by the MMSC display switches. It may be helpful to increase this value if single key presses are being registered multiple times. On the other hand, if some key presses seem to ignored when pressing keys rapidly, it may help to decrease this value.	200000
LOG_DFLT_NUMLINES	The default number of lines in a log. Once a log is filled with this many lines of data, older lines are discarded to make room for new lines. The MMSC must be reset (see reset_mmsc) for this setting to take effect.	200
LOG_DFLT_LINELEN	The default average line length for a log. The total amount of storage allocated for logged messages is NUMLINES * LINELEN. If many lines in a log are larger than LINELEN characters in length, then it may be necessary to discard older lines even before the log contains NUMLINES lines. The MMSC must be reset (see reset_mmsc) for this setting to take effect.	80
PAGE_DFLT_LINES	The default number of lines per page when paging output.	23
PWR_DELAY	The MSC "pwr u" and "pwr c" commands may be intercepted in order to sequence the power to two or more racks one rack at a time to avoid an excessive power surge. (See Intercepted MSC Commands, below.) This variable sets the number of microseconds to wait between "pwr" commands.	5000000

This command accepts an optional destination specification.

steal This command is only meaningful when it is entered from a terminal or alternate console device, and only when it is in MMSC mode. If another device is currently in console mode, it is placed into MMSC mode. Then the device on which the command was entered is placed into console mode. In effect, the system console is "stolen" by the device on which the command was entered.

A user must have supervisor authority to execute this command.

unsetenv var — Restores the default value for environment variable *var*. This command accepts an optional destination specification.

unsteal — This command is only meaningful when it is entered from a terminal or alternate console device, and only when it is in console mode. It places the "other" device into console mode, effectively undoing a *steal* command. If the device on which the command is entered is the normal terminal device, then the "other" console is the alternate console device. Likewise, if the command is entered from the alternate console device, the "other" console is the normal terminal device. The device on which the command was entered is then placed into either MMSC mode (when the device is the terminal) or RAT mode (when the device is the alternate console).

A user must have supervisor authority to execute this command.

ver — Returns a string that indicates the MMSC firmware revision. This command returns a different string than the MSC *ver* command, and so can be used from a tty device to determine whether it is attached to an MSC or an MMSC. The format of the response is MMSC *major.minor*, for example:

MMSC 1.2

This command accepts an optional destination specification. A user must have service authority to execute this command.

carriage return — If neither a destination nor a command follows an MMSC escape character (that is, the first character of the command is carriage return, essentially an empty command), the command is ignored.

If a destination but no command is specified, the destination becomes the new default destination (see Section C.1.3).

C.4 Intercepted MSC Commands

Certain MSC commands are intercepted by the MMSC before being passed to the MSC itself. This interception is not performed when the commands are sent to an MSC using the msc args command. Avoid sending commands to an MSC with the MSC args command whenever possible, because the additional actions performed by the MMSC are usually beneficial. The intercepted commands and the actions associated with them include:

Table C-6 Intercepted MSC Commands

Command	Actions
pas s pw	The command is converted into the MMSC command "password set msc pw", which should have the same effect on the MSC password, plus the added benefit of setting the password on the MMSC.
pwr u pwr c	The command is forwarded to the addressed MSCs, but one rack at a time with a time delay between each. This prevents all of the modules from being powered on nearly simultaneously, which could cause a power surge.
ver	The MMSC version string (rather than the MSC version) displays as documented above.

C.5 Firmware Revisions

Origin2000 and Onyx2 MMSCs have flash-upgradable PROMS. These PROMs control the way the MMSC interacts with other MMSCs, MSC, the user, and the display. To determine the revision level of the MMSCs in a system, type the following at the MMSC prompt:

`MMSC>` **`r * b * ver`**

Note: MMSCs with attached displays will show their PROM revision level in the lower-left corner of their display. This only applies, however, to the MMSC to which the display is directly attached. Other MMSCs in the system may be at different revision levels.

For information on flash-upgrading MMSC PROMs, refer to "Flashing the PROMs" in Chapter 7 and to the *Origin2000 and Onyx2 Diagnostic Roadmap,* 108-0164-00x.

Appendix D

MSC Commands

The Origin2000 and Onyx2 module system controller (MSC) front panel has a DIN-type RS-232 serial port, labeled the Diagnostic Port (see Figure D-1). This port is also known as the alternate console port (ACP). The diagnostic port is primarily used for debugging during manufacturing and when the regular serial-port console is not available.

An RS-232 dumb terminal connects to the diagnostic port and can issue commands to the individual CPUs in the module. All of the CPUs in a single module share this console. As a result, output from multiple CPUs appears on the console simultaneously, interleaving on a line-by-line basis. The CPU is identified at the start of each line in the form of a slot (1 to 4) and slice (A or B). The MSC can be directed to send diagnostic port keyboard input to a particular CPU. It can also be directed to show only the output from a particular CPU.

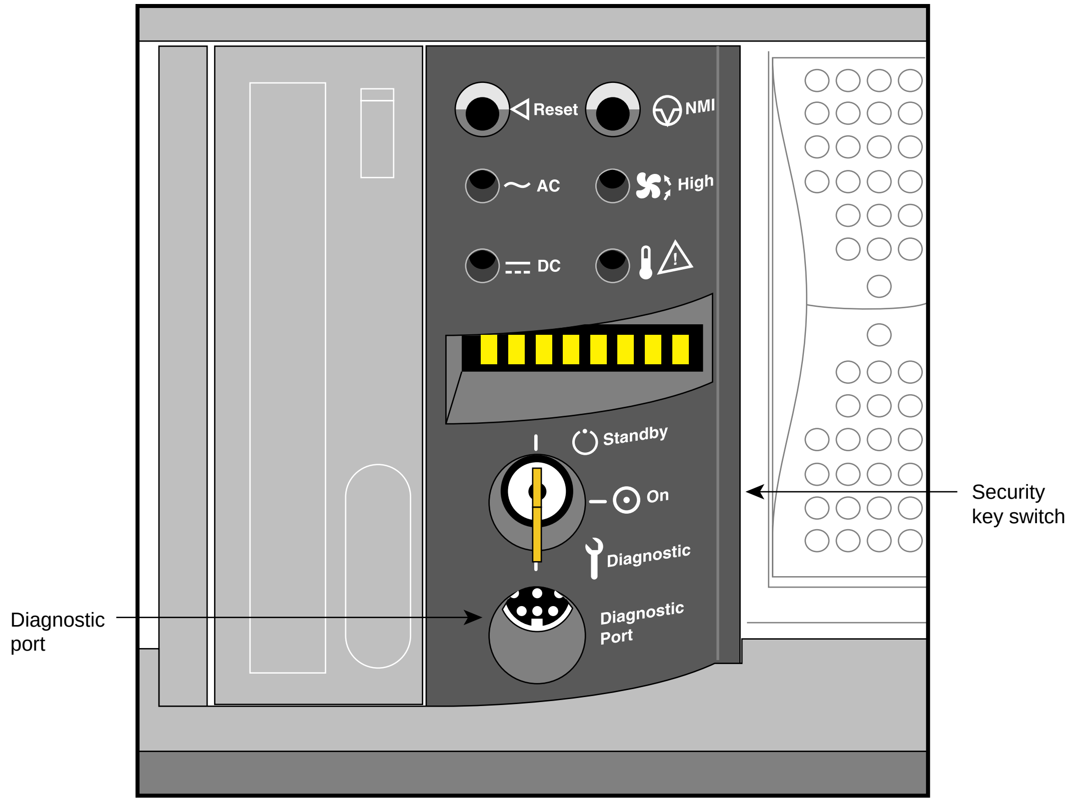

Figure D-1 Diagnostic Port on MSC

D.1 Command Set

Table D-1 describes the MSC command set. Invoke these commands at the diagnostic port terminal by

- typing the escape character ^T (Control-T),
- entering the command line text
- pressing <Enter>

If echoing is enabled (default), the prompt MSC> displays at ^T, and characters that you type are visible. If echoing is not enabled (see the *ech* command), nothing is visible as you type the command.

Each command consists of three letters. Some of the commands take parameters; numeric parameters are always in hexadecimal. All commands return responses that consist of ok, that consist of ok and some hexadecimal value, or one of three error responses (described in Table D-2).

Table D-1 MSC Command Set

Command	Function
clr	Resets all MSC options to their power-up defaults. This includes echo mode, no heartbeats, etc.
dbg	Displays the virtual and physical Debug Switch bytes.
dbg *V P*	Sets the virtual Debug Switch byte to *V*, and the physical Debug Switch byte to *P*.
dsp *M*	Displays message *M* on the 8-character alphanumeric display. *M* may contain up to 8 ASCII characters other than NUL. For example, dsp testing would overwrite the first seven characters with “testing,” and leave the eighth character.
dsc *N* C	Replaces the Nth character on the alphanumeric display with character C. N is a digit from 0 to 7, and C is any ASCII character other than NUL.
ech	Toggles the echoing. By default, echoing is on after reset.
ech 0	Turns off echoing as MSC commands are entered.
ech 1	Turns on echoing.
fan	Displays the fan status. If no fan has failed, returns n or h, according to whether the fans are currently operating at Normal or High speed. If a fan fails on an Origin2000 system, returns f xyz, where xyz are bit maps of the failed fans by row (fan 1 = 1, fan 2= 2, fan 3 = 4). When the MSC detects that a fan failed, it speeds up the remaining fans to maintain cooling levels. The failed fan should be replaced because the increased speed reduces the life of the remaining fans. If more than one fan fails, the system powers down.
fan n	Sets the fans to normal speed.
fan h	Sets the fans to high speed.
key	Returns the status of the key switch as Off, On, or Diagnostic.
mod	Displays the ID number of the module that contains the MSC.
mod xx	Sets the ID number of the module that contains the MSC to xx, where xx is a hexadecimal module number from 01 to ff. (A module number of 00 indicates that no module number is yet assigned.)
nmi	Sends a hardware nonmaskable interrupt (NMI) to all node cards in the MSC’s module. Note: On multimodule systems, an NMI is best performed from the MMSC.

Table D-1 (continued) MSC Command Set

Command	Function
pas xxxx	When xxxx is the correct 4-character password, places the MSC in supervisor mode, where various restricted commands may be used. If the MSC is not in supervisor mode, restricted commands results in a permission error (`err perm`). The MSC is automatically in supervisor mode if the key switch is in the Diagnostic position.
pas s xxxx	Sets the password to xxxx. This is a restricted command. The password is stored in NVRAM and defaults to none.
pwr	Returns u or d, according to whether the system is powered up or down.
pwr u	Powers the system up.
pwr d	Powers the system down.
pwr d *N*	Powers the system down after a delay of N seconds. N is a hex value from 5 to 258 (5 to 600 decimal).
pwr c *N*	Powers the system down, waits N seconds, then powers the system back up; c stands for cycle, and the variable N is a hex value from 5 to 258 (5 to 600 decimal).
rst	Sends a hardware reset to all node cards in the MSC's module. Note: On multimodule systems, a system reset is best performed from the MMSC.
rsw	Returns the current settings of the MSC hardware debug switches. The output of the rsw command is an inverted hexadecimal byte. For example: if switches 1, 2, and 6 are on, the rsw command returns `dc`. The hexadecimal value of `dc` translates to 11011100, which is the inverse of 00111011 (3b).
see	Displays which CPU's output is currently shown on the MSC debut port. This command displays a specific CPU by slot and slice (for example, 2a refers to slot N2, CPUA), or it displays `all` if the output of all CPUs in the module is shown, intermixed by line.
see *cpu*	Displays the output from a specific CPU and discards output from all other CPUs. (Ordinarily, the output from all CPUs is displayed intermixed by line.) CPUs are named by slot and slice (for example, see 2a displays only CPU 2A's input).
see all	Displays the output from all CPUs (this is the power-up default).
sel	Displays the CPU that is currently selected to receive ACP input (MSC debug port input). This command displays a CPU by slot and slice (for example, 2a refers to slot N2, CPUA), or displays `none` if no CPU is selected to receive input, or displays `auto` if the CPU with the last output is selected to receive input.

Table D-1 (continued) MSC Command Set

Command	Function
sel *cpu*	Selects the CPU that receives input from the ACP. Anything typed on the ACP that is not an escaped command is sent through to the selected CPU. CPUs are named by slot and slice as described in Slot and CPU Numbering. For example, **`sel 2a`** selects CPU A in slot N2 for input.
sel auto	Automatically selects the CPU with the last output to receive ACP input (this is the power-up default).
sel none	Causes no CPU to receive ACP input.
tmp	Displays the temperature of the module that contains the MSC. Returns `o` for overtemperature (automatic shutdown pending), `h` for high operating temperature, or `n` for normal operating temperature.
ver	Reports the MSC firmware revision number.
vlm *x y*	Margins the power supply voltages or the midplane termination voltage. Where *x* is 3, 5, or v: 3 = the power supply 3.45 Vdc 5 = the power supply 5.0 Vdc v = the midplane termination Vdc and *y* is h, l, or n: h = margin the specified voltage high (5%) l = margin the specified voltage low (5%) n = margin the specified voltage normal If {l | h |n} is omitted, the specified voltage is returned to the normal setting.

The following MSC commands are not documented in Table D-1: get, hbt, rcf, scf, and tas.

Table D-2 MSC Responses

Response	Reason
err perm	Permission denied. The key switch must be in the Diagnostic position, or you must enter a password via the *pas* command.
err cmd	Unrecognized command mnemonic.
err arg	Invalid command argument(s).

D.2 Security

MSC commands that may be destructive to system operation may only be executed when the MSC is in supervisor (Diagnostics) mode. If the MSC is not in supervisor mode, these commands result in the following error:

```
err perm
```

The MSC is automatically in supervisor mode when the front panel key switch is in the Diagnostics position. It may also be placed in supervisor mode by issuing the MSC command *pas none* to enter the four-character MSC password, where *none* is the default MSC password. The command *pas s* “abcd” would change the MSC password to abcd.

For more information on the MSC, refer to the *Origin2000 and* Onyx2 *Diagnostic Roadmap* document, 108-0164-00x.

D.3 Firmware Revisions

Origin2000 and Onyx2 MSC firmware is contained in a 44-pin PLCC that must be physically replaced to upgrade the MSC firmware revision. Complete the following tasks to determine the revision level of an MSC’s PROM:

- At the MSC console, type **`^T`**
- At the MSC> prompt, type:
  `MSC>` **`version`**

On rack systems, use the following command to simultaneously determine the revision levels of all the MSCs in the system:

- At the MMSC command prompt, type:
  `MMSC>` **`r * b * msc ver`**

For information on upgrading the firmware, refer to the *Origin2000 and Onyx2 Diagnostic Roadmap*, 108-0164-00x.

Index

N

O

P