Netra™ CT 820 Server Product Overview

Sun Microsystems, Inc.
www.sun.com

Submit comments about this document at: http://www.sun.com/hwdocs/feedback

Adobe PostScript

Contents

Figures

Tables

Preface

The *Netra CT 820 Server Product Overview* describes the basic hardware components of the Netra™ CT 820 server. It is a companion to the *Netra CT 820 Server Installation Guide*, which describes how to install the Netra CT 820 server, and the *Netra CT 820 Server Service Manual*, which describes the removing and replacing of the server's field-replaceable units (FRUs).

The intended reader of this manual is an experienced system administrator who has experience with the Solaris operating environment. The reader should be comfortable with LAN fundamentals and with networking in general.

Safety and Compliance

All Netra CT 820 servers are shipped with the *Netra CT 820 Server Safety and Compliance Manual*, which specifies the environmental and electrical safety requirements for the product and contains compliance certification for various countries.

How This Book Is Organized

Chapter 1 introduces the Netra CT 820 server.

Chapter 2 describes the chassis.

Chapter 3 describes the distributed management cards.

Chapter 4 describes the switching fabric boards.

Chapter 5 describes the Netra CP2300 cards.

Glossary is a list of words, phrases, and acronyms and their definitions.

Using UNIX Commands

This document might not contain information on basic UNIX® commands and procedures such as shutting down the system, booting the system, and configuring devices. Refer to the following for this information:

- Software documentation that you received with your system
- Solaris™ Operating System documentation, which is at

  http://docs.sun.com

Typographic Conventions

Typeface or Symbol	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your .login file. Use ls -a to list all files. % You have mail.
AaBbCc123	What you type, when contrasted with on-screen computer output	% **su** Password:
AaBbCc123	Book titles, new words or terms, words to be emphasized	Read Chapter 6 in the *User's Guide*. These are called *class* options. You *must* be superuser to do this.
	Command-line variable; replace with a real name or value	To delete a file, type rm *filename*.

Shell Prompts

Shell	Prompt
C shell	*machine_name*%
C shell superuser	*machine_name*#
Bourne shell and Korn shell	$
Bourne shell and Korn shell superuser	#

Related Documentation

The Netra CT 820 server documentation is listed in the following table:

Title	Part Number
Netra CT 820 Server Product Overview	817-2643
Netra CT 820 Server Installation Guide	817-2641
Netra CT 820 Server Service Manual	817-2642
Netra CT 820 Server System Administration Guide	817-2647
Netra CT 820 Server Safety and Compliance Manual	817-2645
Netra CT 820 Server Product Note	817-2646

You might want to refer to documentation on the following software for additional information: the Solaris Operating System (OS) and the Netra High Availability (HA) Suite.

Third-Party Web Sites

Sun is not responsible for the availability of third-party web sites mentioned in this document. Sun does not endorse and is not responsible or liable for any content, advertising, products, or other materials that are available on or through such sites or resources. Sun will not be responsible or liable for any actual or alleged damage or loss caused by or in connection with the use of or reliance on any such content, goods, or services that are available on or through such sites or resources.

Contacting Sun Technical Support

If you have technical questions about this product that are not answered in this document, go to:

http://www.sun.com/service/contacting

Sun Welcomes Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions. You can submit your comments by going to:

http://www.sun.com/hwdocs/feedback

Please include the title and part number of your document with your feedback:

Netra CT 820 Server Product Overview, part number 817-2643-11

CHAPTER **1**

Introduction

This chapter provides an overview of the Netra CT 820 server's basic hardware components. The Netra CT 820 server is a CompactPCI (cPCI) packet switching backplane-based, NEBS level 3-certified, ETSI-compliant, rack-mountable server.

Note – Reliability, Availability and Serviceability metrics for the Netra CT 820 server are available through the Sun Sales office under a Non-Disclosure Agreement.

The Netra CT 820 server complies to the following specifications:

■ PICMG 2.16 R1.0 CompactPCI specifications
■ PICMG 2.1 R2.0 CompactPCI Hot-Swap specifications
■ PICMG 2.9 R1.0 CompactPCI IPMI specifications
■ PICMG 2.0 R3.0 CompactPCI Mechanical and Power specifications

The hardware components for the Netra CT 820 server can be broken down into four sections:

■ The chassis—Chapter 2
■ The distributed management card—Chapter 3
■ The switching fabric board—Chapter 4
■ The Netra CP2300 card—Chapter 5

FIGURE 1-1 shows the components in a Netra CT 820 server from the front. FIGURE 1-2 shows the components in a Netra CT 820 server from the rear.

Fan trays (Chapter 2)

Node boards
(Chapter 5)

Distributed
management
cards (Chapter 3)

Switching
fabric boards
(Chapter 4)

ESD ground jack
(Chapter 2)

Removable grill
with filter (Chapter 2)

Power supplies (Chapter 2)
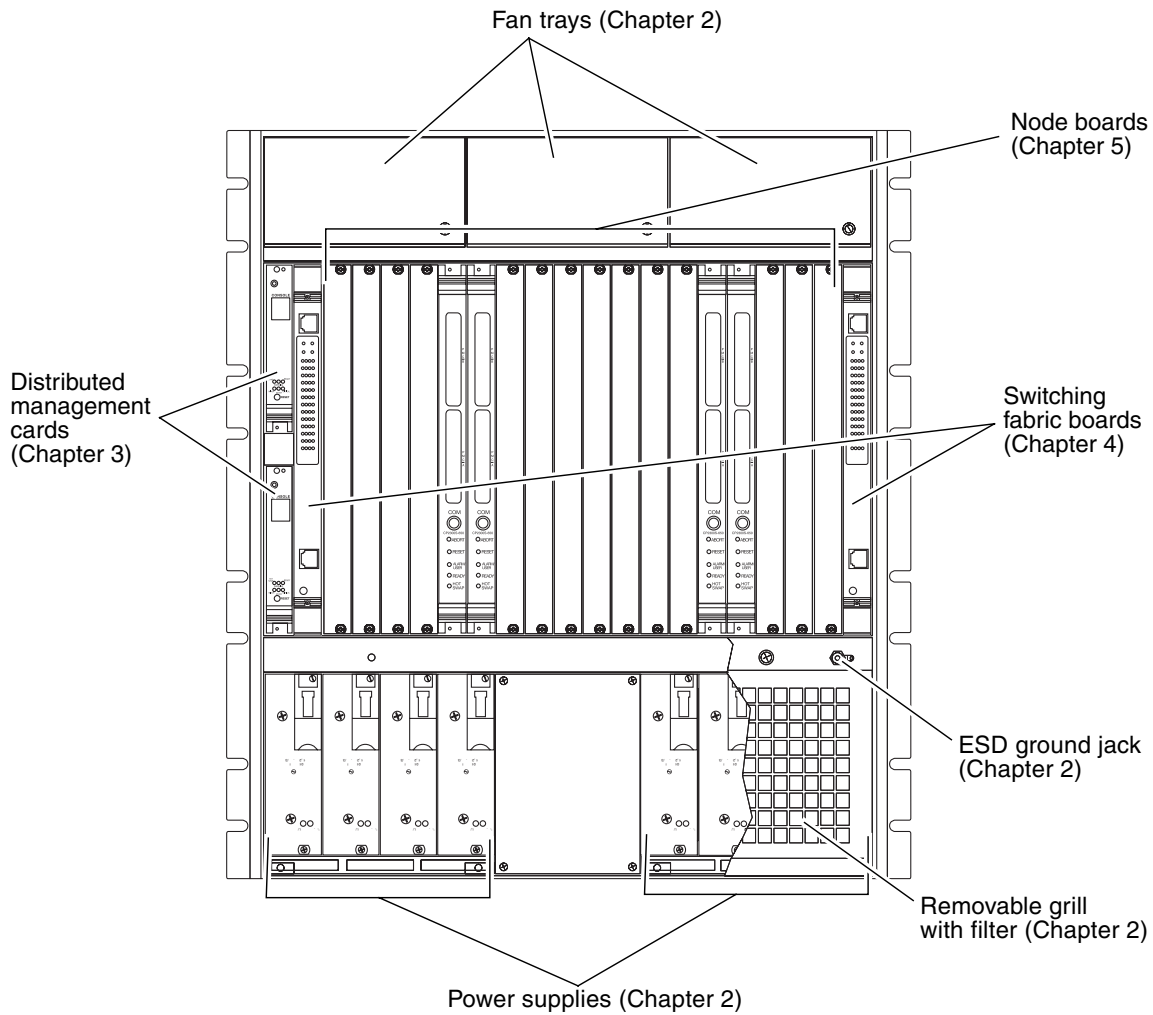
FIGURE 1-1 Netra CT 820 Server Components (Front)
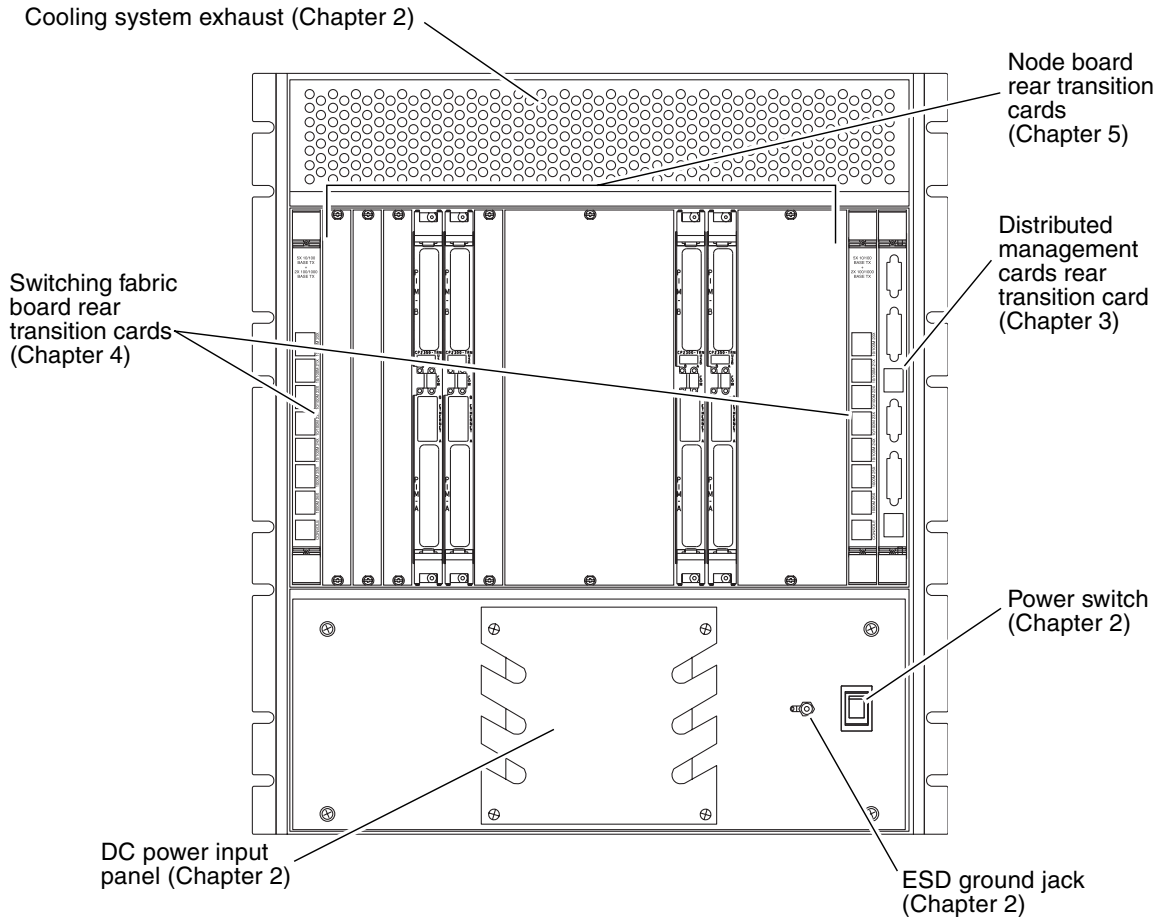
FIGURE 1-2   Netra CT 820 Server Components (Rear)

Chassis Description

The Netra CT 820 server provides OEM equipment designers with carrier-grade, standards-based, high availability solutions built on the PCI Industrial Computer Manufacturer's Group (PICMG) 2.16 specification. This high-capacity platform features 18 node slots and a redundant infrastructure (switching fabric board, management, power, and cooling), making it ideal for carrier-grade telecom and Internet applications. Beyond its high availability features, the Netra CT 820 server is highly modular, scalable, and serviceable.

Hot-swappable system components provide built-in redundancy to simplify replacement and minimize service time. Redundant distributed management cards enable customers to manage multiple processor boards and conduct chassis diagnostics remotely for enhanced system reliability. Two 6U slots are reserved for PICMG 2.16 switching fabric boards. The Netra CT 820 server routes Ethernet signals across the midplane without the use of cables, saving time in setup, maintenance, and repair, and eliminating the thermal challenges of traditional cabling methods.

Chassis Features

- **High Capacity** – 12U, 19-inch rackmount enclosure
    - 21 slots (18 node slots, 2 switching fabric board slots, and 2 3U distributed management card slots)
    - Up to 40W per node slot, 60W per switching fabric board slot power and cooling
    - Efficient front-to-rear cooling
- **High Availability** – Delivers "five-nines" availability using redundant modules for fault recovery:
    - Two PICMG 2.16 extended switching fabric board slots supporting hot-swappable boards
    - Two hot-swappable Intelligent Platform Management Interface (IPMI)-based distributed management cards
    - 2N hot-swappable fan trays for cooling
    - 2N hot-swappable, load-sharing 250W DC power supplies
    - N+N power input
    - Dual power domain midplane, which isolates catastrophic power failures
- **Highly Manageable** – IPMI-based redundant distributed management cards
    - Unique star topology for management buses increases reliability and security
- **Highly Serviceable** – All field-replaceable units serviceable from the front (except rear transition cards [RTCs])
- **Highly Scalable** – flexibility to scale processors, I/O boards, power, and cooling
- **Standards and Compatibility** – Interoperable with processor boards, peripherals, switches, and subcomponents from multiple manufacturers
    - PICMG 2.16-compliant midplane
    - Designed for NEBS Level 3 and FCC Class A
    - PICMG 2.9 IPMI-based management

CompactPCI Packet Switching Backplane (CompactPCI/PSB)

The PICMG 2.16 specification defines the CompactPCI Packet Switching Backplane. The Netra CT 820 server sends all Ethernet signals across the midplane. By moving system traffic from the shared bus architecture to a fault-tolerant switched midplane, overall system throughput can be vastly increased while retaining the reliability and hot-swap capacity of cPCI.

The Netra CT 820 server accommodates dual redundant switching fabric boards and node boards.

- A switching fabric board links to each node board in a packet switching chassis. In this way, every node board can communicate with every other node board, thereby forming a switching fabric. The switching fabric boards are linked to each other in a Netra CT 820 server. A switching fabric board can only be used in a switching fabric board slot.

- A node board links to the switching fabric board in a Netra CT 820 server. Each node board is linked to both switching fabric boards, thereby providing redundant fabric. Node boards can only be used in node slots.

Midplane Slot Assignments

The Netra CT 820 server midplane features 21 6U slots. Slot 1 is dedicated to two 3U distributed management cards (slots 1A and 1B). Slots 2 and 21 are dedicated to 6U PICMG 2.16-compliant switching fabric boards. Slots 3 through 20 are PICMG 2.16-compliant node slots.

All slots support IEEE 1101.11-style, 80 mm–deep transition cards in the rear-panel I/O section directly behind the midplane.
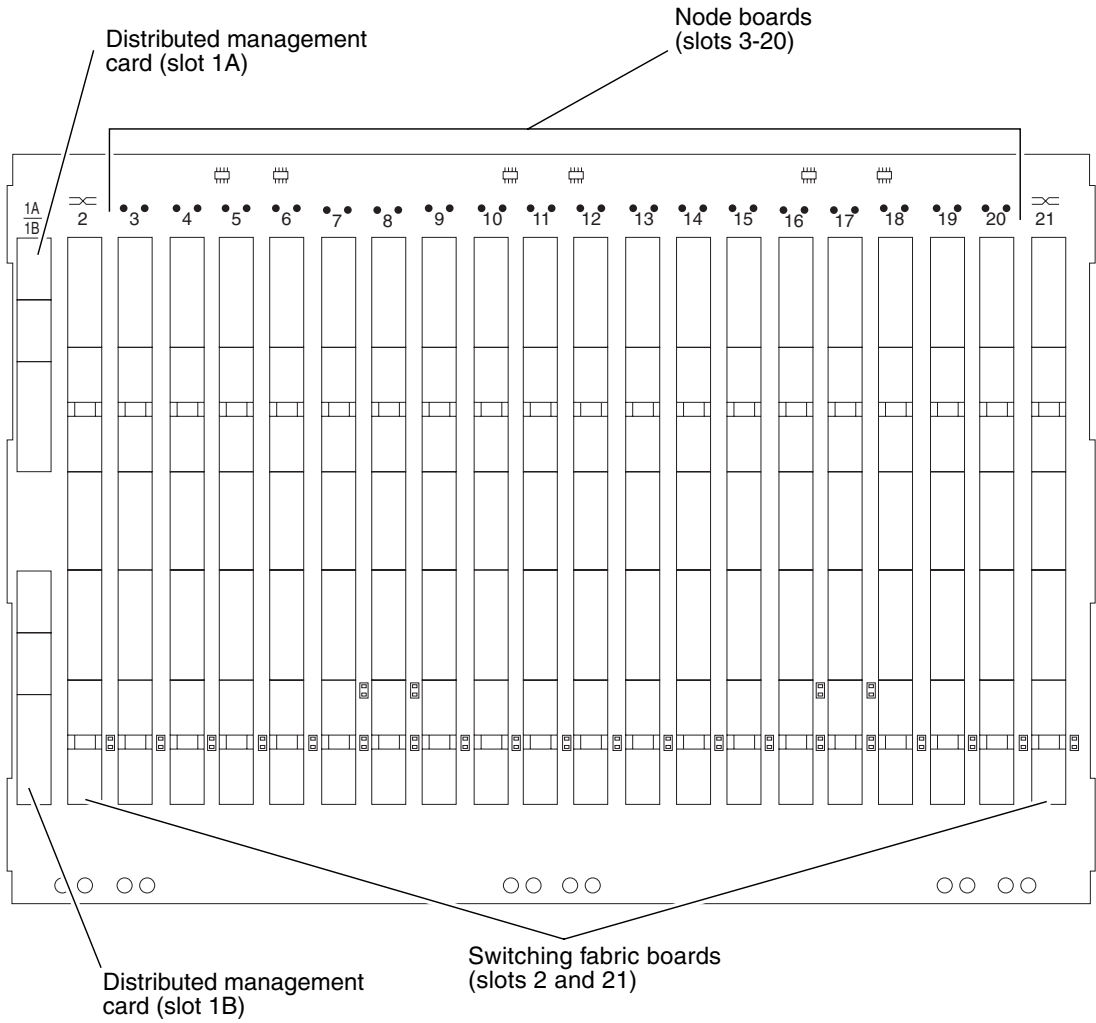
FIGURE 2-1 Midplane Slot Assignments (Front)

Power Subsystem

Central office power plants are based on redundant high-availability -48V DC power. AC power delivered by the public utility is rectified to DC and stored in battery systems, which in turn supply the equipment. Equipment is fed from two sets of power feeds.

Increasingly, these feeds are being provided by independent power plants. In environments with independent power plants, the two feeds must be completely isolated from one another.

Caution – It is imperative that these feeds *never* join. Failure to properly design a chassis's power input subsystem could result in voltage transference, system damage, or safety issues.

The Netra CT 820 server is designed specifically to meet the increasing requirements of the telco environment and central office. The Netra CT 820 server delivers a highly redundant 2N power system, including dual power domains on the midplane. The system minimizes the impact caused by failures of power feeds, power supplies, the midplane, or the cooling subsystems.

Power Distribution Overview

The Netra CT 820 server implements dual power domains. Each power domain is powered from 2N redundant power supplies with 2N redundant power inputs. FIGURE 2-2 illustrates how power is distributed within the Netra CT 820 server.
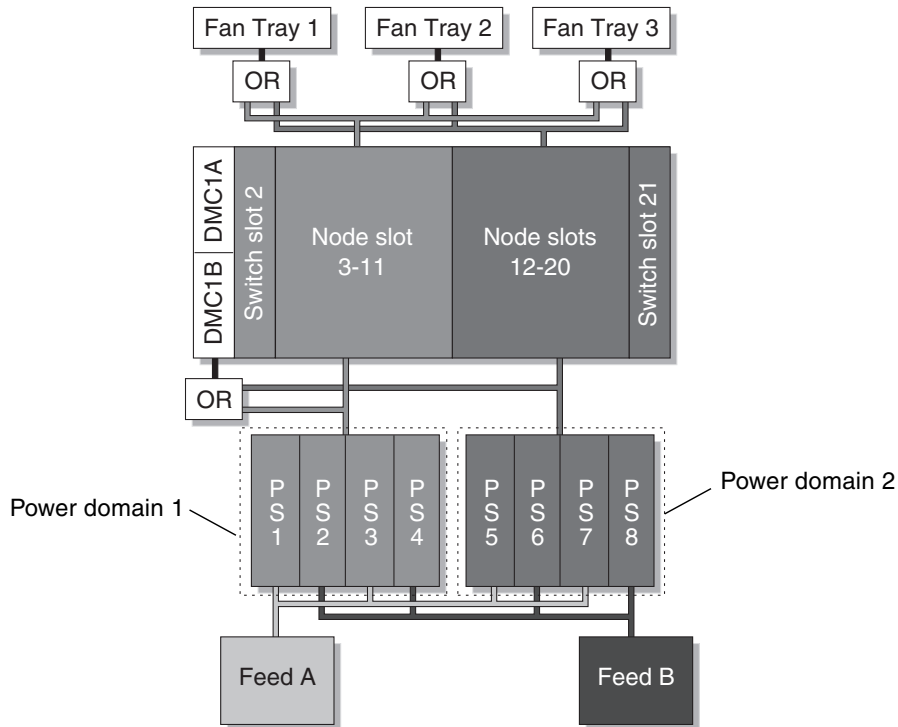
FIGURE 2-2 Power Distribution Overview

There are two redundant power feeds: A and B. Each power feed provides input to two power supplies in each of the two power trays. As shown in FIGURE 2-2, feed A provides power to the odd–numbered power supplies while feed B provides power to the even–numbered power supplies.

This configuration provides 2+2 redundant power to each power domain. The system is designed to continue operation with any single fault. Examples of faults include:

- Power feed failure
- Power supply failure
- Power domain failure

If a single power feed fails, each domain continues to receive power from the two power supplies powered from the healthy feed. If a single power supply fails, each domain continues to receive adequate power from the healthy power supplies. If a power domain fails, the other power domain is unaffected. A power domain might fail if pins get bent, creating a short between two voltage rails or between a voltage rail and ground.

TABLE 2-1 shows the effects of various types of power faults with 250W power supplies.

TABLE 2-1   Power Fault Tolerance Matrix

Failure	Healthy Feeds	Healthy Supplies		Healthy Domains	Watts	
		Domain 1	Domain 2		Domain 1	Domain 2
No faults	2	4	4	2	500 (Redundant)	500 (Redundant)
Power supply fault in domain 1	2	3	4	2	500 (Partial redundancy)	500 (Redundant)
Power supply fault in domain 2	2	4	3	2	500 (Redundant)	500 (Partial redundancy)
Feed A fault	1	2	2	2	500 (Non-redundant)	500 (Non-redundant)
Feed B fault	1	2	2	2	500 (Non-redundant)	500 (Non-redundant)
Domain power fault in domain 1	2	0 (Supplies fold back)	4	1	0 (Supplies fold back)	500 (Redundant)
Domain power fault in domain 2	2	4	0 (Supplies fold back)	1	500 (Redundant)	0 (Supplies fold back)

Note that the redundant power feeds are kept isolated from one another. The feeds are not diode-ORed together. This is important for telco facilities that provide redundant power plants, because a diode failure could allow faults in one power plant to affect the other power plant.

A power switch is provided on the back of the chassis. This switch disables the *output* from all eight power supplies.

Power Budgeting

To properly budget for power consumption within the Netra CT 820 server, many factors need to be considered. These factors include:

- Board and component power consumption
- Dual power domains
- Redundant power supplies
- Redundant power feeds
- Power supply capacity
- Multiple voltages
- Fan trays and distributed management cards
- Cooling capacity

Board and Component Power Consumption

See the information in this section for the power consumption of the switching fabric boards, distributed management cards, and fan trays.

TABLE 2-2   Power Consumption for a Switching Fabric Board

Power Consumption	Voltage	Maximum Current	Maximum Power
38.5 watts maximum	3.3V	4.5A	14.8W
32.4 watts typical	5V	4.5A	22.5W
	+12V	0.1 A	1.2W
	-12V	0A	0W

TABLE 2-3   Power Consumption for a Fan Tray (Three Blowers per Fan Tray)

	Nominal	Maximum
Current (@12V)	3.75A	5.64A
Power	45.00W	67.68W

TABLE 2-4   Power Consumption for a Distributed Management Card

	Nominal
Current	2A @ 5V
Power	10W

Dual Power Domains

The Netra CT 820 server has two power domains. Each power domain is powered from four power supplies in a 2N redundant configuration.

Redundant Power Feeds and Power Supplies

Each power domain receives input power from redundant power feeds. Each redundant power feed provides input power to half the power supplies within the power domain. Power budgets should assume that only *half* of the provided power supplies in each domain are operational when faults are present. For example, a system with four power supplies per power domain only has two operation supplies in each power domain if a DC input feed has failed. See "Power Distribution Overview" on page 10 for more information about the effects of faults on the power subsystem.

Power Supply Capacity and Multiple Voltages

Power budgeting must account for the power supply capabilities for each of the four voltage levels (3.3V, 5V, +12V, and -12V). For a system using eight 250W hot-swap DC power supplies, the power supplies can provide the following currents for each voltage level:

TABLE 2-5   Budget by Voltage Using 250W Power Supplies

Voltage	Max Current/Supply	Max Current/Domain[2]	Redundant Max Current/Domain	Max Current/Chassis[2]	Max Redundant Current/Chassis
3.3V[1]	40A/supply	160A/domain	80A/domain	320A	160A
5V[1]	40A/supply	160A/domain	80A/domain	320A	160A
+12V	5.5A/supply	22A/domain	11A/domain	44A	22A
-12V	1.5A/supply	6A/domain	3A/domain	12A	6A

1 Current on the 3.3V and 5V combined current cannot exceed 40A per power supply.

2 The Max Current/Domain or Chassis should only be used in a development environment, or an environment that does not require power redundancy.

As new power supplies are introduced with higher power outputs, more redundant power and current can be delivered to the chassis.

Distributed Management Cards and Fan Trays

The distributed management cards and fan trays draw power from both power domains. Under normal conditions, each domain provides approximately 50% of the power for these shared components. Power budgets for each domain should provide for *at least* 50% of the necessary distributed management card and fan tray power. Due to tolerances of the ORing diodes and the voltage levels, it is likely that slightly more than 50% of the power will be drawn from the domain with the higher voltage. If this causes a slight over-current situation, that voltage will drop until the ORing diodes automatically balance the load between the two domains.

Note that in the event of a power domain failure, 100% of the distributed management card and fan tray power is drawn from a single domain. However, the healthy domain has a complete set of power supplies (assuming only single failures); that is, 50% is drawn from the "budget" supplies and 50% from the "redundant spare" supplies.

Cooling Capacity

The Netra CT 820 server is designed to cool a maximum of 20W per distributed management card slot, 40W per node slot, and 60W per switching fabric board slot. Unused slots must be covered with a filler panel to properly maintain the cooling in other slots.

Power Input

The following topics specify the various interfaces for the chassis's power input subsystem.

Power Feed Interface

DC input power is provided through two DC line filters. These filters are labeled A and B. Each filter has three terminals: -48V DC, Return, and chassis ground. The -48V DC and Return terminals are designed for connection with #10 ring terminals, and the chassis ground terminal is designed for connection with a 1/4-inch ring terminal.

The A and B connections provide redundant input power to both power trays. For full redundancy, these connections *must*:

- Have independent circuit breakers
- Come from different sources

Circuit breakers should be rated based on maximum load at minimum voltage with one failed feed. The circuit breakers *cannot* be rated for more than 60A under any circumstance.

Each DC power input includes an EMI filter to reduce line-conducted emissions.

Refer to the *Netra CT 820 Server Installation Guide* for more information about connecting DC power to the Netra CT 820 server.

Power Trays

Two power trays occupy the bottom of the Netra CT 820 server. Each tray has its own backplane that provides slots for four power supplies.
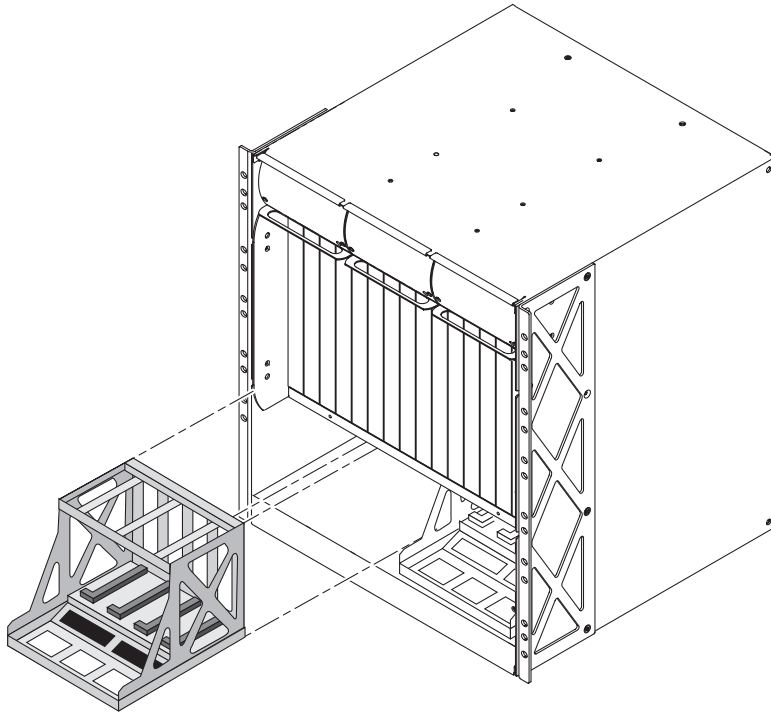
FIGURE 2-3   Power Trays

Power Supplies

The Netra CT 820 server is designed for use with the 3U by 8HP CompactPCI power supplies that comply with the PICMG 2.11 Power Interface Specification R1.0 that are designed specifically for use with the Netra CT 820 server.

The Netra CT 820 server uses 250W DC power supplies. These supplies provide non-IPMI management features (DEG#, FAIL#, and INH#). For information on removing and replacing the power supplies, refer to the *Netra CT 820 Server Service Manual*.

The Netra CT 820 server power supply is a hot-swappable, CompactPCI power supply using a standard Positronic 47-pin connector that is fully compliant with the PICMG 2.11 Power Interface Specification. High current density enables this unit to deliver up to 40A on both the +5V and +3.3V outputs combined. Remote sense and active current share on the +5V, +3.3V, and +12V outputs, along with ORing FETs, allow these units to be used in redundant, hot-swap applications.

TABLE 2-6 DC Input Specifications

Parameter	Conditions/Description	MIN	NOM	MAX	Units
Input Voltage –DC	Continuous input range	36		75	VDC
Hold-up Time	From 48 VDC input	4			ms
Input Current	At full rated load; 36 VDC, 48VDC		7.6, 5.4		A
Input Protection	Non-user serviceable, internally-located input line fuse				
Inrush Surge Current	Internally limited by thermistor and electronic switch			12	A
Operating Frequency	Switching frequency of main output transformer	125		145	KHz
Input Transient Protection	Varistor				

Midplane Power Distribution

The Netra CT 820 server's power distribution architecture includes dual power domains and ORed management power to central resources. FIGURE 2-4 illustrates power distribution on the midplane.
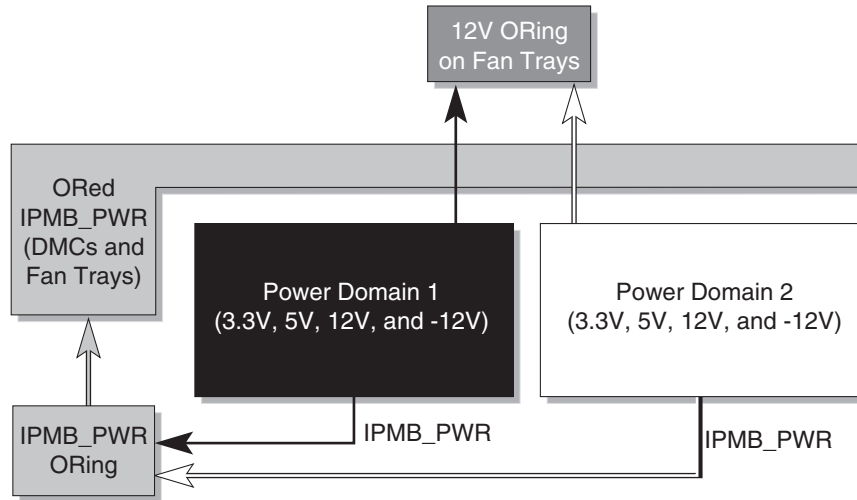
FIGURE 2-4   Midplane Power Distribution

Dual Power Domains

The Netra CT 820 server's midplane implements two independent power domains for the CompactPCI slots. FIGURE 2-2 shows the power domains.

Power domain 1 includes slots 2 through 11. One switching fabric board is installed in slot 2 and node boards can be installed in slots 3 through 11. Power domain 1 is powered from power tray 1.

Power domain 2 includes slots 12 through 21. One switching fabric board is installed in slot 21 and node boards can be installed in slots 12 through 20. Power domain 2 is powered from power tray 2.

Both power domains are cooled and managed by common, redundant fan trays and chassis management modules. These common components receive management power by ORing the two domains' IPMB_PWRs together. The IPMB_PWR ORing diodes are on the midplane. The fan tray blowers are powered by ORing +12V from both power domains. These ORing diodes are physically located in the fan trays.

Cooling Subsystem

The Netra CT 820 server uses three hot-swappable fan trays to cool the chassis. Each fan tray includes three blowers. FIGURE 2-5 illustrates the airflow through the platform.

Rear of chassis

Front of chassis
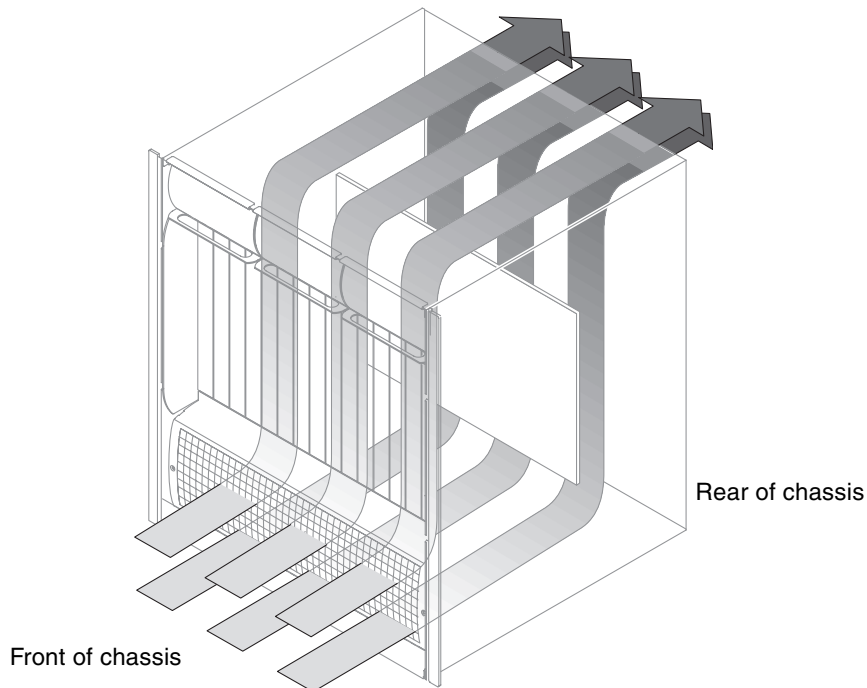
FIGURE 2-5   Airflow

Fan Trays

Each of the Netra CT 820 server's three fan trays is a modular, field–replaceable unit. Each tray contains three blowers, a control board, and a status LED. These features are discussed in the following sections.

Blowers

Each fan tray contains three DC brushless blowers for a total of nine blowers in a Netra CT 820 server. Each blower has a two-phase, four-pole motor. The following table lists blower specifications:

TABLE 2-7   Blower Specifications

Parameter	Description
Rated voltage	12 VDC
Operating voltage	7.0–13.8 VDC
Input current	1.15A, 1.73A max
Input power	13.8W, 20.6W max
Max speed	3100 rpm
Max air flow (at Zero Static Pressure)	1.12 m^3/min, 1.00 m^3/min 39.55 CFM, 31.5 CFM min
Max air pressure (at Zero Air Flow)	33.0 mmH$_2$O, 25.0 mmH$_2$O min 1.299 inchH$_2$O, 0.984 inchH$_2$O min
Acoustical noise (avg.)	56.5 dB-A, 59.0 dB-A max
Life expectancy	70,000 hours continuous operation at 40°C with 15-65% RH

Control Board

The fan tray includes a control board for fan tray management. The control board provides the following features:

- Power sequencing/inrush limiting
- Over-current protection
- Backfeed protection
- Health LED control
- FRU information
- Temperature monitoring
- Voltage monitoring

On insertion or initial power–up, the fan control board stages power to the blowers. The first blower is enabled immediately. The second blower is enabled about 300 ms after the first, and the third blower is enabled about 300 ms after the second. Each blower is independently protected using a self-resetting positive temperature co-efficient (PTC) component.

The control board also includes logic to set the LED to the "unhealthy" state if any of the blowers drop below approximately 1000 rpm, or if onboard voltage is invalid. The distributed management card can also force the LED to the "unhealthy" state.

Fan Tray Status LED

Each fan tray provides a health/status LED on its front panel. TABLE 2-8 lists the different LED states:

TABLE 2-8 Fan Tray Health/Status LED States

Color	Description
Off	Not Powered
Green	Powered and Healthy
Amber	Unhealthy

Intake Grill With Air Filter

An air intake grill with a replaceable air filter is attached to the bottom front of the Netra CT 820 server. The grill must be removed to replace the filter. The grill must also be removed to access the system's power supplies.

Distributed Management Card Description

The distributed management cards are 3U, single-slot cards that plug into slots 1A and 1B in the Netra CT 820 server. They provide centralized management and alarming for up to 21 slots, as well as for system power supplies and fans.

The distributed management card is essentially a special–purpose single–board computer with its own CPU, memory, operating system, and peripherals. The distributed management card monitors and configures IPMI-based components in the chassis. When thresholds (such as temperature and voltage) are crossed or a failure occurs, the distributed management card captures these events, stores them in an event log, sends SNMP traps, and drives the telco alarm relays and alarm LEDs. The distributed management card can query FRU information (such as serial number, model number, and manufacture date), detect presence of components (such as fan tray and CPU board), and perform health monitoring of each component. In addition, the distributed management card controls the power-up sequencing of each device (except the fan trays), and the power to each slot through the BD_SEL# signal.

Distributed Management Card Features

- High–density 3U x 1-slot form factor
- Compatible with PICMG 2.1, 2.16, and 2.9-compliant components
- Manages through the Intelligent Platform Management Bus (IPMB) protocol
- Supports failover from the active distributed management card to the standby card
- Provides isolated IPMI signals for each slot for maximum security and reliability
- Out-of band management interface through the rear transition card
- Hot-swap support for IPMI–based field–replaceable components
- DB15 alarm interface through the rear transition card
- Critical, Major, and Minor alarm LEDs at the front panel
- Status and hot–swap LEDs at the front panel
- Monitors backplane voltages and status for up to eight power supplies
- Monitors system temperature sensors
- Monitors system fan tray presence
- Monitors tachometers for up to 16 system fans

Distributed Management Card Ports and LEDs

The Netra CT 820 server is available only as a rear-access server, so most of the cabling for the distributed management card will be to the rear transition card at the rear of the server. The only exception would be the console port on the distributed management card at the front of the server. You can administer the Netra CT 820 server either from the front of the server, through the console port on the distributed management card, or from the rear of the server, through one of the two console ports on the rear transition card.

⚠ **Caution –** Do *not* attempt to administer the Netra CT 820 server through the console ports on both the distributed management card and the rear transition card at the same time. No mechanism is provided to disable simultaneous input from the front and rear cards, so attempts to input data from both ports will place the boards in an unknown state.

FIGURE 3-1 shows the console port and LEDs on the distributed management card. TABLE 3-1 describes the functions of the console port, and TABLE 3-2 describes the functions of each LED.
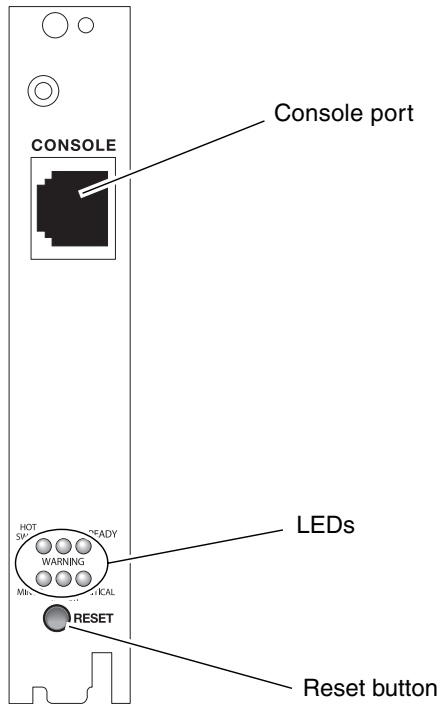
FIGURE 3-1 Distributed Management Card Ports

TABLE 3-1 Distributed Management Card Port Functions

Port	Function
Console	This port is used to remotely administer the Netra CT 820 server using the Netra CT 820 server software. Refer to the *Netra CT 820 Server Software Administration Manual* for more information.

TABLE 3-2   Distributed Management Card LED Functions

LED	Function
HOT SWAP	• Off – The card is up and running. Do not remove the distributed management card in this state. • On – The card is powered off. It is safe to remove the distributed management card in this state.
WARNING	Reserved for future use.
READY	• Off – The card is not ready. • On – On the *active* distributed management card, this signals that the card has booted and is fully powered up and running. • Blinking – On the *standby* distributed management card, this signals that the card has booted and is fully powered up and running.
MINOR	• Off – No alarm has been triggered. • On – A minor event has been triggered.
MAJOR	• Off – No alarm has been triggered. • On – A major event has been triggered.
CRITICAL	• Off – No alarm has been triggered. • On – A critical event has been triggered.

There is also a Reset button at the base of the distributed management card. This button is used to perform a hard reset of the distributed management card. You should use this button to perform a hard reset of the distributed management card only if you cannot go through a graceful shutdown of the system.

Note – By default, the top distributed management card (labeled 1A) will be the *active* management card after the initial chassis power-on, and the bottom card (labeled 1B) will be the *standby* management card. After a hardware or software failover event, the standby distributed management card will become the active card and take over all of the management services in the chassis. For more information about distributed management card failover, and what hardware and software events cause failovers, refer to the *Netra CT 820 Server System Administration Guide* (817-2647).

Rear Transition Card for the Distributed Management Card

The rear transition card is used to route the signals from the distributed management cards installed in slots 1A and 1B at the front of the server to ports in the rear transition card installed at the rear of the server. FIGURE 3-2 shows the ports on the rear transition card for the distributed management cards.
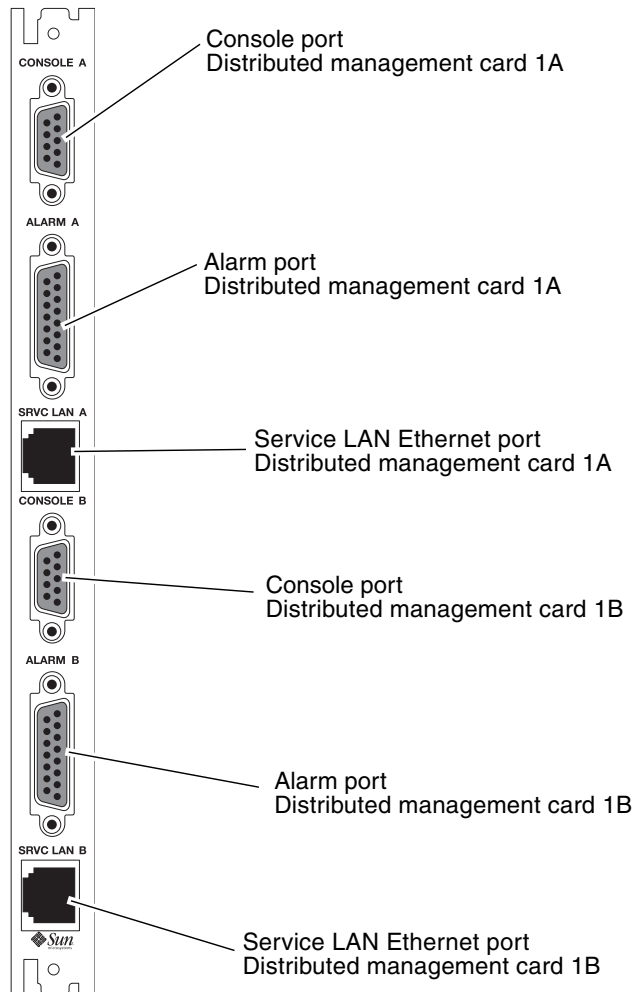
FIGURE 3-2 Ports on the Rear Transition Card for the Distributed Management Card

See TABLE 3-3 for a description of the functions of each of the ports used in the rear transition card.

TABLE 3-3 Port Functions of the Rear Transition Card for the Distributed Management Card

Port	Function
Console	This port is used to remotely administer the Netra CT 820 server using the Netra CT 820 server software. Refer to the *Netra CT 820 Server Software Administration Manual* for more information.
Alarm	This port transmits alarm signals to external systems (dry–contact relay).
Service LAN Ethernet	This port provides a Service LAN Ethernet connection.

Switching Fabric Board Description

The Netra CT 820 server is available only as a rear-access server, so all of the cabling for the switching fabric board will be to the rear transition card at the rear of the server. There are no active components on the rear transition card, so you must have the accompanying switching fabric board installed in the same slot at the front of the server, even though no cables will be connected to that front card. However, you will use the LEDs on the switching fabric board installed in the front of the system to check the status and fault conditions on the board.

The rear transition card is connected to the switching fabric board through the midplane through the J3 and J5 CompactPCI connectors. This card is used to route Ethernet traffic through four 10/100–megabit switched Ethernet MDIX ports through shielded RJ-45 connectors (ports 20–23), two 1000–megabit Ethernet MDIX ports (ports 25–26), and an RJ-11 console port.

Note – Port 24 on the rear transition card for the switching fabric board is reserved for future use.

FIGURE 4-1 shows the location of the LEDs on the switching fabric board. FIGURE 4-2 shows the ports on the rear transition card for the switching fabric board.
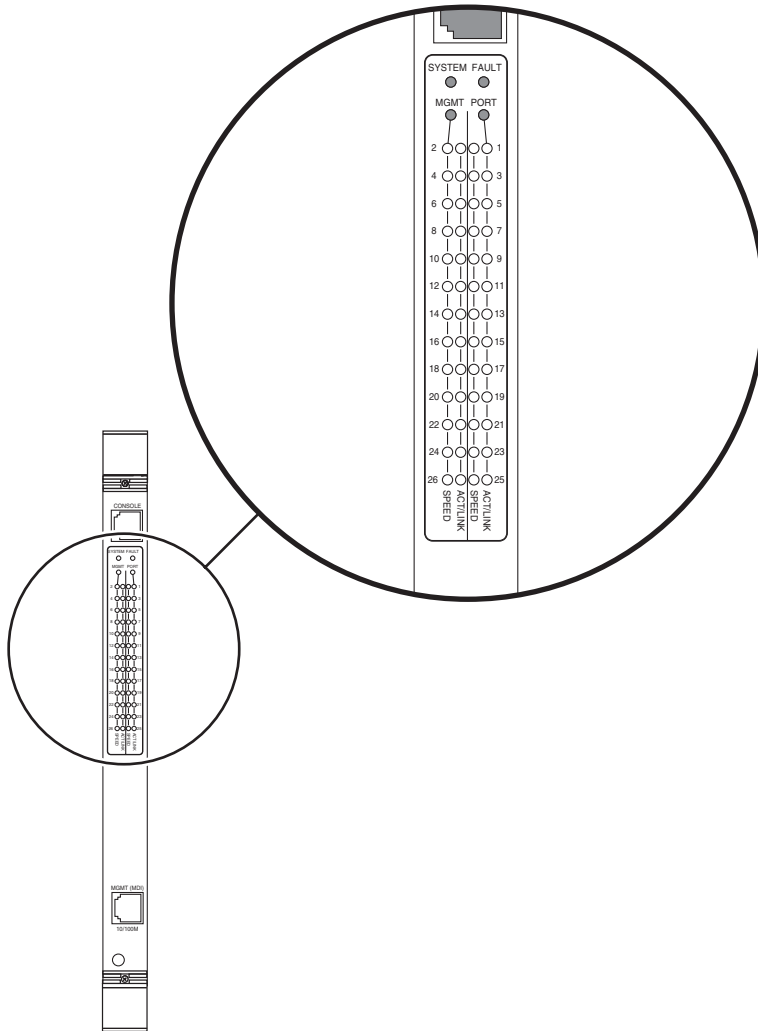
FIGURE 4-1 Switching Fabric Board LEDs

10/100 Mbps RJ-45
Ethernet ports
(20 - 23)

Reserved for
future use

1000 Mbps MDIX
Ethernet ports
(25 and 26)

RJ-11 Console port
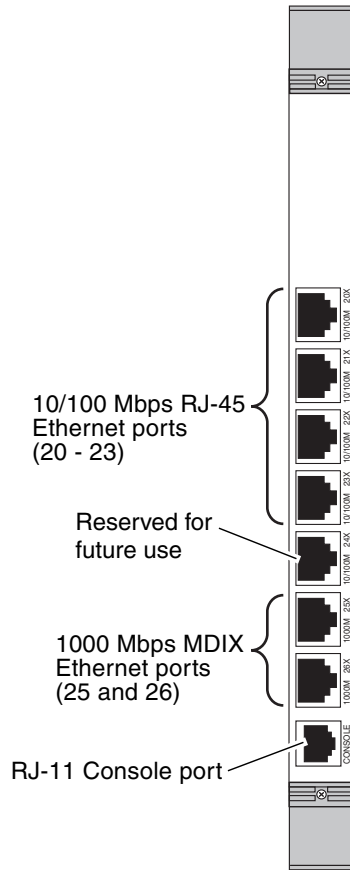
FIGURE 4-2   Ports on the Rear Transition Card for the Switching Fabric Board

Front Panel LEDs

The front panel LEDs are:

- Board-level System Status LEDs
- Port LEDs
- Hot–swap LED

System and Fault LEDs

Management (Out-of-Band) port LEDs

Port LEDs
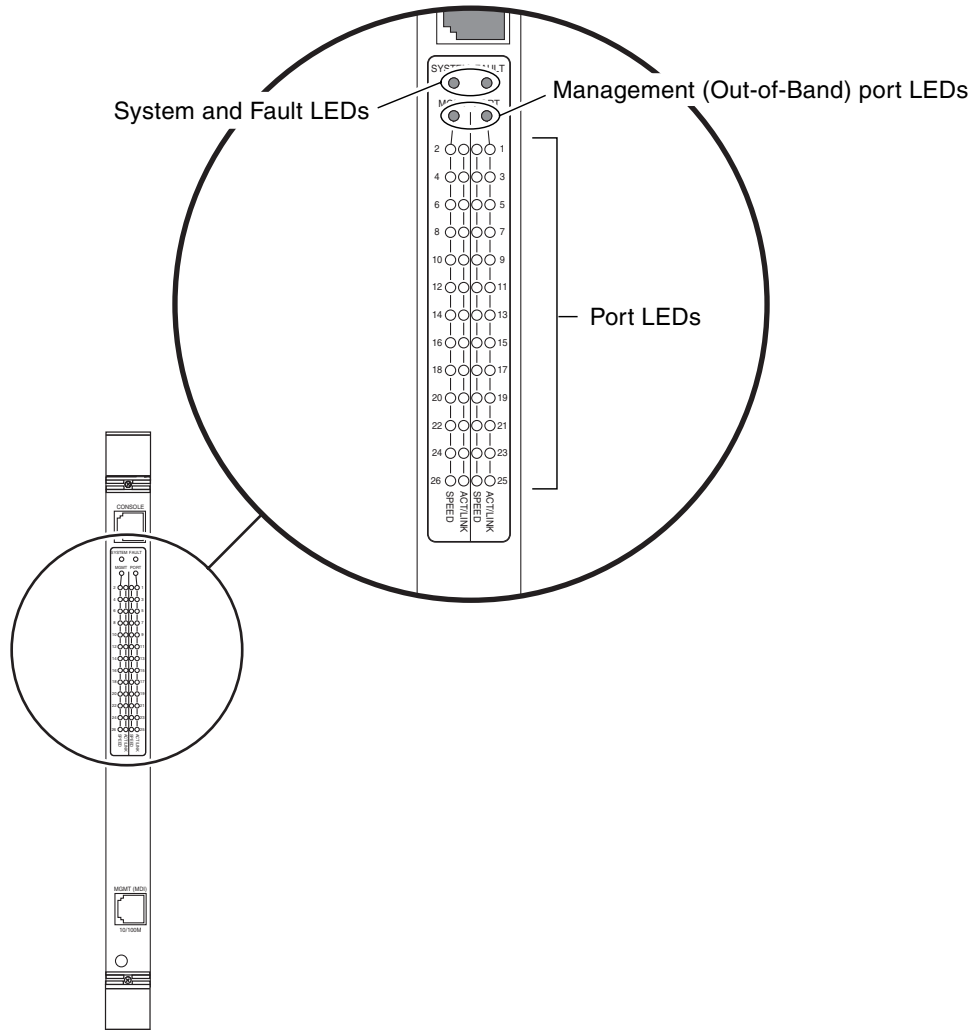
FIGURE 4-3 Switching Fabric Board Front Panel LEDs

System Status LEDs

Two LEDs provide board-level system status information.

- A bi-color (green-yellow) LED provides *system* status on the front panel.
- A red LED provides *fault* status on the front panel.

On power-up, the System LED is yellow. After the basic board integrity checks are complete, the System LED changes to green.

The System LED normally remains green unless a board failure is detected. A board failure turns the System LED off. The System LED could also be off if the board is in a fault state. See TABLE 4-1.

TABLE 4-1 System LED

LED State	State Indication
Green	Board is operational
Yellow	Board is booting
Off	No power, or serious fault

The Fault LED is off unless a serious internal error is detected. If a serious error is detected the red Fault LED turns on. To turn off the Fault LED once it is illuminated, the switching fabric board must be powered off. Refer to the *Netra CT 820 Server Service Manual* for instructions on powering off a single slot in the server. See TABLE 4-2 for more detail on the Fault LED.

TABLE 4-2 Fault LED

LED State	State Indication
Red	Fault
Off	No fault or no power

Port Status LEDs

The status of all Ethernet ports (including the out-of-band management port) is indicated by LEDs visible from the front panel. A green LED is used to indicate link and activity (transmit or receive), and a yellow LED is used to indicate the current port speed. See TABLE 4-3 and TABLE 4-4.

TABLE 4-3 Port Link/Activity LED

LED State	State Indication
Green On	Link
Flashing Green	Activity (transmit or receive)
Off	No Link

Note – The 10/100–megabit ports and the 100/1000–megabit twisted pair ports flash the Activity LED on both transmit and receive activity. The Gigabit (FX) ports flash the Activity LED on receive activity only.

TABLE 4-4 Port Speed LED

LED State	10/100 Mb Port Indication	100/1000 Mb TX Port Indication
Yellow On	100 Mb	1000 Mb
Off	10 Mb	100 Mb

TABLE 4-5 explains how the LED port numbers on each switching fabric board map out to the physical card slots in the Netra CT 820 server.

TABLE 4-5 LED-to-Slot Mapping

LED Port Number(s)	Physical Card Slot	LED Port Number(s)	Physical Card Slot
1	Node card slot 3	14	Node card slot 16
2	Node card slot 4	15	Node card slot 17
3	Node card slot 5	16	Node card slot 18
4	Node card slot 6	17	Node card slot 19
5	Node card slot 7	18	Node card slot 20
6	Node card slot 8	19	• LED 19 on switching fabric board in slot *2* maps to distributed management card *1A* • LED 19 on switching fabric board in slot *21* maps to distributed management card *1B*
7	Node card slot 9	20	Port 20 on the rear transition card for that switching fabric board
8	Node card slot 10	21	Port 21 on the rear transition card for that switching fabric board
9	Node card slot 11	22	Port 22 on the rear transition card for that switching fabric board
10	Node card slot 12	23	Port 23 on the rear transition card for that switching fabric board
11	Node card slot 13	24	Port 24 on the rear transition card for that switching fabric board
12	Node card slot 14	25	Port 25 on the rear transition card for that switching fabric board
13	Node card slot 15	26	Port 26 on the rear transition card for that switching fabric board

Switching Fabric Board Features

- Highly integrated 10/100 ports and Gigabit port switching fabric
- Low latency wire-speed switching
- 9–Gbps switching fabric
- 8K MAC addresses
- Out-of-band management port
- Spanning Tree extensions for improved reconvergence
- Switching fabric board redundancy
- Broadcast storm suppression
- Embedded online diagnostics

Ethernet Ports

Switched Ethernet Ports

The switching fabric board provides 10/100–megabit Ethernet ports and 1000–megabit Ethernet ports interconnected through a fully meshed 9–Gbps switching fabric. An 802.1D bridge relay entity exists between any pair of switched ports. The switched Ethernet ports offer low–latency, wire-speed switching performance on all ports simultaneously. The switching fabric board supports 8192 (8K) MAC address entries (simultaneously) in its filtering database.

The switching fabric board provides a logical link control (LLC) interface between the IP stack of the embedded onboard operating system and the switched ports. This interface allows for sending and receiving frames to and from the onboard CPU and any of the switched Ethernet ports. This capability allows the onboard CPU to be a presence on the attached network and allows external clients to manage the operation of the switch through any of the switched Ethernet ports.

Out-of-Band (Management) Ethernet Port

One 10BASE-T/100BASE-TX unswitched Ethernet port is provided on the switching fabric board for switch management purposes. This port conforms to all requirements of the IEEE 802.3 (CSMA/CD) MAC interface.

This port does *not* have a MAC relay entity between itself and any of the switched Ethernet ports. This separation makes it possible to allocate all switched ports to users while always maintaining a secure and available path to manage the switch.

The switching fabric board provides an LLC interface between the IP stack of the embedded onboard operating system and the out-of-band Ethernet port. This interface allows sending and receiving frames to and from the onboard CPU and the out-of-band Ethernet port. This capability allows the onboard CPU to be a presence on the attached network and allows external clients to manage the operation of the switch through the out-of-band Ethernet port.

The out-of-band Management port provides access to a separate subnet from that of the switched ports. The switching fabric board must be provided with an IP address for each LLC interface that is to be enabled. If the out-of-band Ethernet interface is to be used, it must be assigned an IP address and subnet mask combination that place the interface in a separate subnet from that of the switched ports.

Spanning Tree Extensions

The switching fabric board supports a pair of extensions to the standard Spanning Tree protocol that enhance its utility in building fault–tolerant networks. The standard Spanning Tree protocol supports redundant links for resiliency, but the algorithm that selects the active link can take anywhere from 30 to 50 seconds to converge. This delay often causes application layer timeouts.

The switching fabric board addresses this problem with two extensions to the Spanning Tree protocol: Fast Uplink and Fast Port.

Fast Uplink

Fast Uplink enables a redundant pathway toward the root bridge to be quickly selected upon failure of link integrity on the current path to the root bridge. The switchover time is typically less than one second.

Fast Port

Fast Port enables the user to select ports that do not need to participate in the Spanning Tree protocol. These ports must have simple client stations attached to them and must not provide a path to any device that might act as a bridge between LAN segments. That is, the user must ensure that no bridges or routers can be accessed through a port that has Fast Port turned on. When Fast Port is turned on, the port does not participate in Spanning Tree and does not undergo or cause any Spanning Tree topology changes.

Broadcast Storm Suppression

The switching fabric board offers the ability to limit forwarding of broadcasts, multicasts, and frames with unknown destination MAC addresses (flooding). The user can set a maximum frame rate forwarded by the switch for these traffic types. This feature prevents malfunctioning network equipment from flooding the network.

Switching Fabric Board Embedded Diagnostics

The embedded diagnostic functions of the switching fabric board ensure virtually faultless operation both at initialization and continuously during switch operation. The In-System verification includes Power-On Confidence (POC) tests and Online Integrity tests.

Power-On Confidence (POC)

Power-On Confidence (POC) tests are diagnostic tests run when the switching fabric board is booted. The POC test suite includes only those diagnostic tests that do not affect externally attached devices. When running POC tests, the switching fabric board becomes available to the user within 150 seconds of a power-on or reset. All subsystems are included in the POC suite.

Online Integrity Tests

These are tests that are run on a continuous basis while the switching fabric board is in service. Online Integrity tests are enabled by default, but the switching fabric board also provides a method to disable them.

The integrity tests operate continuously during execution of application code to detect hardware faults. Upon detecting a fault, the switching fabric board can take itself out of service. All attached clients will lose link.

Watchdog Timer

A hardware watchdog timer removes the switching fabric board from service if it is not refreshed periodically by the CPU. This guards against CPU failure. If the watchdog timer expires, link will be deasserted on all external ports, and the switching fabric board will be reset. The user can configure the board to reboot after a watchdog event or to remain offline.

System Management Bus (PICMG 2.9)

The switching fabric board's implementation of the System Management Bus (SMB) conforms to PICMG 2.9, the specification that defines SMB in the CompactPCI environment. The SMB utilizes the I^{2}C bus hardware layer coupled with Intelligent Platform Management software. The Intelligent Platform Management Interface specification enables system components on the I^{2}C hardware layer bus to communicate with the distributed management card, thereby making possible "out-of-band" management of a system's resources, with an emphasis on controlling hot-swap activities.

Netra CP2300 cPSB Board Description

> **Note –** This chapter gives general information on the Netra CP2300 cPSB board. For more detailed information, refer to the Netra CP2300 cPSB board documentation.

The Netra CP2300 cPSB board is a crucial building block that network equipment providers (NEPs) and carriers can use when scaling and improving the availability of carrier-grade systems. Based on industry standards, the Netra CP2300 board provides high performance and is powered by a 650–MHz UltraSPARC® IIi processor.

The Netra CP2300 board enables customers to mix and match third-party PCI mezzanine cards (PMCs), making it easier for them to tailor solutions to their specific application needs. The Netra CP2300 board provides PCI Industrial Computers Manufacturers' Group (PICMG) CompactPCI packet-switched backplane compliance (for details, see TABLE 5-1) and is NEBS Level 3–certified to meet the requirements of the communications and service provider environments.

Netra CP2300 cPSB Board Features

The Netra CP2300 board is a cPSB single-board computer designed for high-performance embedded, compute density applications. The Netra CP2300 board has system management controller (SMC) capability that supports hot-swap operations, system management, and environmental monitoring. Powered by a 650–MHz UltraSPARC IIi processor, with full utilization of the PMC slots, the Netra CP2300 board is an ideal platform for NEPs to use for a wide variety of Solaris applications.

FIGURE 5-1 displays an illustration of a typical Netra CP2300 board, and TABLE 5-1 lists a summary of features of the board.
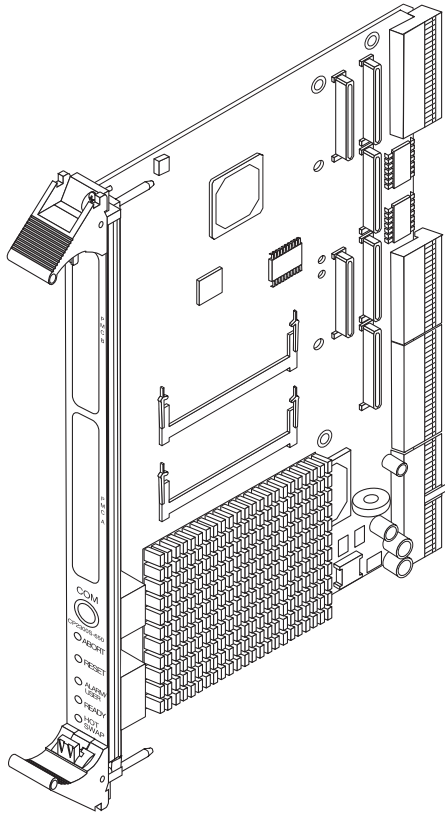
FIGURE 5-1 Typical Netra CP2300 cPSB Board

TABLE 5-1 Feature Summary

Feature	Description
CPU	UltraSPARC IIi 650–MHz processor with internal L2 cache (2:2 mode, 512–Kbyte, 4-way set association).
Memory	512 MB on-board memory, with two PC133–compliant DRAM EEC SODIMM slots available for additional memory expansion.
Power requirement	Estimated at 22W (typical) and 26W (peak maximum) at 650 MHz, and 30W (peak maximum) with two 512–MB SODIMMs installed. (The power requirements exclude PMC, SODIMM memory, and rear transition module (RTM) power.)

TABLE 5-1 Feature Summary *(Continued)*

Feature	Description
PICMG and PCI compliance	• PICMG 2.0 R3.0 CompactPCI bus specification for 33 MHz PCI speed • PICMG 2.1 R2.0 Hot-Swap Specification • PICMG 2.10 R1.0 Keying of cPCI boards and backplanes • PICMG 2.15 R1.0 PCI Telecom Mezzanine/Carrier Cards (PTMC) support • PICMG 2.16 R1.0 CompactPCI Packet Switching Backplane • PICMG 2.3 R1.0 PMC on CompactPCI Tables 1 and 3 • PICMG 2.9 R1.0 System Management Specification • PMC Specification P1386 Draft 2.4 • CMC P1386 Draft 2.4 Standard for CMC
Node Board support	The board functions as a cPSB node board with the Solaris software package.
IPMI system management	Uses IPMI communications with Baseboard Management Controller (BMC); performs Advanced System Monitoring (ASM) on local board interface–for example, temperature sense, FRU ID, and control.
Hot-swap support	Basic, Full, and High-Availability (HA) hot-swap support.
Operating system	Solaris 8 OE, Release 2/02, and subsequent compatible versions.
Front I/O and connectors	• Two PMC slots • One serial port (can only be used if rear serial port COM A is not used)
Connectors on rear transition card (optional)	• Two 10/100 Ethernet ports (switch selectable and mutually exclusive with cPSB networks) • Two serial ports • One USB port • Two PCI Interface Module (PIM) slots • 40-pin EIDE connector on board (not on panel); EIDE connector is switch selectable and mutually exclusive with PIM A I/O 1-29.
PMC I/O	Provision for adding up to two independent hardware vendor (IHV)–supplied PMC expansion ports on front panel.
Backplane PMC I/O	One USB port; also provision for adding two IHV–supplied PIM I/O ports when used with transition card
Watchdog timer	Two-level watchdog timer.
NVRAM	8 Kbyte non-volatile I^{2}C SEEPROM to save OpenBoot PROM configuration. TOD has no battery backup, as battery backup is not needed.
System flash	1 Mbyte on board.

TABLE 5-1   Feature Summary *(Continued)*

Feature	Description
User flash	7 Mbyte on board.
Building compliance	NEBS Level 3.
Flash update	Supported from downloaded file.

Note – For EMI compliance of front access ports, use shielded cables on all I/O ports. The shields for all shielded cables must be terminated on both ends.

Netra CP2300 System Configurations

The boards can be deployed in various electrical configurations to suit each end-user requirement. For example, a board can be configured to boot from a network as a diskless client with either a cPSB or rear transition card network connection. Alternatively, industry-standard PMC and PIM hardware from IHVs can be installed to provide local disk I/O, which can optionally be used as a boot path. The installation procedure is independent of the type of enclosure, whether a floor-mounting rack or a bench-top cabinet is used. The Netra CP2300 board has fixed on-board memory and connectors for additional memory.

PMC and PIM Modules

The Netra CP2300 board has one serial port on the front panel. The IHV-built PMC modules provide additional I/O to the front panels. PMC modules decode their custom I/O from the Netra board's on-board PCI bus A signals.

Rear Transition Card

The optional Netra CP2300 cPSB transition card (Sun part number, XCP2300-TRN) installs into the rear of the Netra CT 820 server, opposite the Netra CP2300 board. The transition card connects with the host CompactPCI connectors through the backplane pins and carries two serial ports and a USB port out to its rear-panel flange. The transition card contains a 40-pin EIDE connector that is switch selectable and mutually exclusive with the PIM A connector I/O pins 1-29. The transition card also provides two switch-selectable RJ-45 10/100 Ethernet connectors, which are

mutually exclusive with cPSB network interfaces. The Netra CP2300 board SW3 DIP switch controls the access to these Ethernet ports. Refer to the *Netra CT 820 Server Service Manual* for more information on setting the Ethernet networking for the Netra CP2300 board.
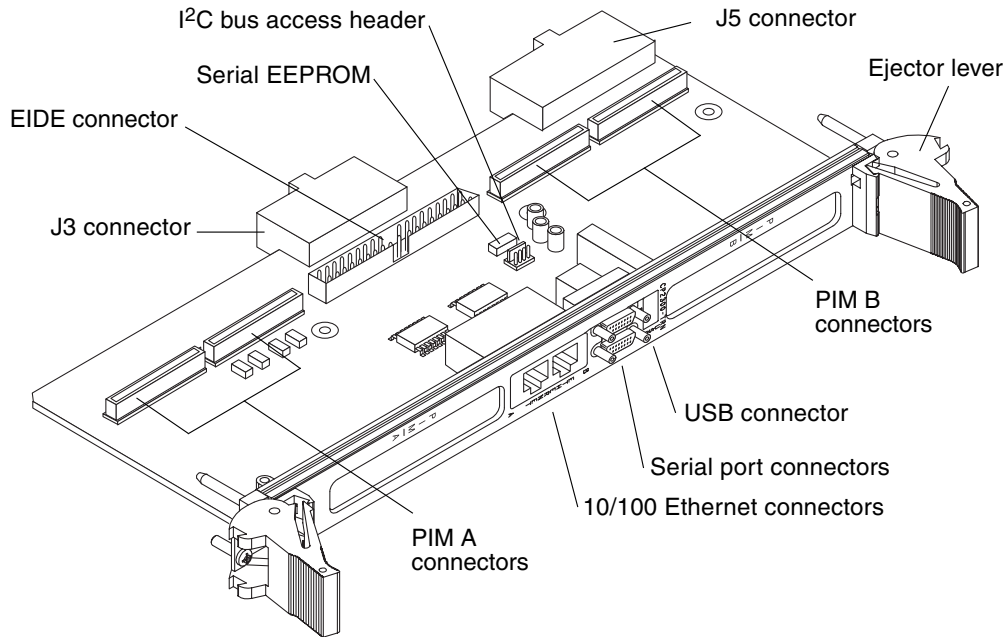

FIGURE 5-2 Netra CP2300 cPSB Board Rear Transition Card

> **Note –** When the transition card is used with the Netra CP2300 board, shielded cables are required for serial and USB I/O ports, and unshielded cables can be used on Ethernet ports in order to satisfy EMI compliance standards. The shields for all shielded cables must be terminated on both ends.

The transition card can also be fitted with IHV PCI Interface Modules (PIMs) that are configured to bring I/O channels to the unit rear panel. A PIM hardware kit includes a card for the PMC slot and a card for the PIM slot on the transition card. A PIM is a rear-panel extension added to a PMC module. When the PIM I/O is configured, the front PMC I/O output is not accessible.

The customer can order the Netra CP2300 transition card, build a custom card, or buy a transition card from an IHV. A minimal set of I/O must provide for a boot path for the host board and for a path for console I/O to deliver commands and to read board and system status.

Possible boot and console configurations are described in TABLE 5-2. Sun Microsystems provides the Netra CP2300 board and a compatible Netra CP2300 transition card. This transition card provides two 10/100 Ethernet RJ45 ports from the host to the rear of the system, which can optionally be used to accomplish a network boot as a diskless client. The other configurations require IHV hardware.

TABLE 5-2   Netra CP2300 Board I/O Configurations

I/O	Hardware Required	Description
Ethernet	Netra CP2300 transition card—supplied as an option for rear access	Default boot path uses cPSB Ethernet port; board runs in diskless client configuration.
SCSI	Netra CP2300 transition card; PMC SCSI I/O	Can be used for local boot; requires optional transition card with PMC SCSI I/O.
Serial data	Netra CP2300 board	Serial port A on front panel provide the path of the default console I/O.
	Netra CP2300 transition card	When optional transition card is installed, the card's Serial port A will become the path of the default console I/O (see FIGURE 5-2 for location).
USB	Netra CP2300 transition card	Can be used for keyboard I/O.

Index