Field Engineer Technical Education Supplement: M2000

AUSPEX

Copyright

Part Number 850XXX, Revision 1.3, December, 1998

Trademarks

FCC Statement

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 in which case the user will be required to correct the interference at his own expense.

Protection Against Electrostatic Discharge

To prevent damage to the system due to electrostatic discharge, always wear the antistatic wrist strap provided with your network server when you come in contact with the system.

Auspex Systems, Inc.
2300 Central Expressway
Santa Clara, California 95050
Phone: (408) 566-2000
Fax: (408) 566-2020
Internet: Info@Auspex.Com
uucp: uunet!auspex!info

Publication Change Record

This page records all revisions to this publication. The first entry is always the publication's initial release. Each entry indicates the date of the revision and a brief description of the area of change.

Part Number	Revision	Date	Description
850XXX	1.0	November 13, 1998	IInitial release
850XXX	1.1	November 27, 1998	Updates and corrections in all chapters.
850XXX	1.2	December 4, 1998	Changed NS 10000 to M2000 in all chapters; added Title, Table of Contents and Chapter 12; incidental corrections and additions in all chapters.
850XXX	1.3	December 30, 1998	Incidental additions and corrections throughout manual; added appendices A, B, C, and D.

Contents

Chapter 7 Processing Subsystem

Chapter 8 Storage Subsystem

Appendix A Glossary

Appendix B Course Evaluation

Appendix C New Products

Appendix D Errata

1 Auspex: Company Overview

Introduction

Auspex occupies a unique position in the server market because of its products and its focus on customer satisfaction. Representing Auspex in the field is challenging. The corporate background provided in this chapter will help you understand Auspex and its views on customer satisfaction.

Objectives

After completing this chapter you should be able to do the following:

▲ State the philosophy behind the Auspex product line.

▲ Define quality from the Auspex perspective.

▲ State the underlying principle behind Auspex's treatment of customers.

Founding Philosophy

Product Need

Fill a product need by building a file server specifically to serve data to desktops

Product Requirements

Reliable network file system

Fast access to network data

Strong, consistent technical support

Founding Philosophy

Product Need

Larry Boucher founded Auspex Systems in 1987. The philosophy behind the product line stems from Boucher’s observation that many businesses used “computing” servers to store and move data. These servers were not designed for this task and had trouble managing and moving data. Boucher believed the solution was to build a true network file server from the ground up; a device optimized to serve data to desktops.

Product Requirements

Offering a product that delivers fast, reliable file service to a customer means meeting three key requirements:

- An extremely reliable network file system. For many businesses, downtime is unacceptable.
- Fast, transparent access to network data. If users encounter a consistently slow network, they will store their data locally.
- Support that is strong and consistent enough to maintain the trust of existing customers and inspire potential customers to use Auspex products.

Meeting these requirements guides Auspex’s product development.

Network Server Marketplace

Products

Auspex builds a range of highly scalable, configurable network servers

Distribution

Sales offices in the U.S. and Europe

Distributors in Europe and Japan

Promotion

Repeat customer business

Customer referrals

Network Server Marketplace

Products

Auspex Systems develops, manufactures, and sells a range of network data server products. The servers are highly configurable and scalable in terms of connectivity, processing power, and storage capacity. In addition to the server hardware, Auspex has software products available to minimize system downtime and increase reliability. Auspex also provides a range of post-sales service contracts to meet varying customer needs.

Distribution

Auspex sells products through approximately 25 domestic sales offices, direct sales offices in the UK, France, and Germany, and distributors in Europe and Asia. Auspex customers are typically Fortune 1000 corporations whose businesses range across semiconductor development, finance and banking, oil and gas, networking, and entertainment.

Promotion

Auspex distinguishes itself from its competitors with its focus on customer satisfaction. As a result, much of Auspex's business comes from repeat customer sales and customer referrals.

Commitment to Quality

Auspex definition of quality: meeting or exceeding the requirements and expectations of the customer

Each department establishes quality goals

All employees evaluated on contribution to quality and customer satisfaction

Customers have access to Auspex executive staff

Full disclosure of all quality statistics

Auspex reviews product plans with customers

Commitment to “Golden Rule” when dealing with customers

Commitment to Quality and Customer Satisfaction

Auspex defines quality as “meeting or exceeding the requirements and expectations of the customer.” As a company, we are committed to total customer satisfaction by providing continuous improvement of our products, services, and support to our customers. To this end, we implement the following:

- Each executive officer of the company establishes quality goals including targets for improved performance.
- Auspex considers each employee’s contribution to customer satisfaction and quality in their performance appraisals, compensation, and in other forms of recognition.
- All members of the Auspex executive staff make themselves available and accessible to customers at all times. They strive to develop, maintain, and enhance relationships with all Auspex customers.
- Auspex releases product quality and customer satisfaction information to our customers on a nondisclosure basis. The information includes quality objectives, progress towards those objectives, and statistical summaries and reports.
- Auspex reviews product plans with appropriate customers prior to final adoption to assure that we address as many of their needs as Auspex resources allow. We announce new products, services, and support only when we are certain that we are able to meet the commitment made in our announcements.
- The fundamental underlying principle for our business is that we always treat our customers, our suppliers, and our associates as we wish to be treated. We employ only those individuals who match the Auspex profile of quality, customer satisfaction, and teamwork.

Student notes

2 Field Support

Introduction

A Field Engineer representing Auspex is accountable for the results of the service call. This chapter familiarizes the student with Auspex's support programs, the service call process, and the Field Change Order (FCO) process.

Objectives

After completing this chapter, you should be able to do the following:

- ▲ State the Field Engineer responsibilities after being dispatched to an Auspex Premier Service customer.
- ▲ State the information needed to close an Auspex service call.
- ▲ Describe, from a high level, the FCO process.

Auspex Support Programs

Auspex Premier Service (APS)

Auspex Advantage Service (AAS)

Auspex Shared Advantage (ASA)

Time and Materials Service (TMS)

Support Programs

Definitions

It is important to understand that unless otherwise specified, warranties and support program contracts cover *individual systems,* not customer sites. There are customer sites with multiple systems, each covered by different support program contracts. Customers make support decisions based upon budget, and the criticality of the system in question.

Auspex has the following support programs available:

- *Auspex Premier Service* (APS) is the premium, fully-comprehensive plan designed to meet the mission-critical needs of most Auspex customers.
- *Auspex Advantage Service (AAS)* is the standard plan for customers who require support only during business hours.
- *Auspex Shared Advantage* (ASA) is for customers who perform their own maintenance using spare parts from a purchased spare parts kit.

A Field Engineer visit to an ASA customer site is rare but does occur under the following circumstances:

- Auspex sells new systems with an optional comprehensive 90-day warranty which entitles a customer to APS service during the warranty period. If an ASA customer elects this warranty option, an FE will be dispatched for service calls on the covered system during the warranty period.
- If an ASA customer does not have the needed replacement part in their spares kit, or the replacement part failed, an FE will be dispatched with parts from the nearest consigned kit.

The details of the Auspex Shared Advantage program are irrelevant to Field Engineers and therefore not included in this manual.

Features of Auspex Premier Service	
Hardware Support	**Software Support**
Committed on-site response within four hours of dispatch (optional two-hour response uplift)	Toll-free telephone software support available 24 hours per day
Continuous service once a call is initiated	Software problem status reports
Management and technical problem escalation	Management and technical problem escalation
Remote diagnosis by a Technical Support Engineer (TSE) that can be contacted 24 hours per day	Distribution of new software maintenance releases
No-cost installation of needed Field Change Orders (FCOs)	Distribution of documentation updates for new software maintenance releases
Assignment of local System Engineer (SE) and Account Executive (AE)	
Spares kit, consigned to a third-party logistics service, in close proximity to the customer site	

Auspex Premier Service

Auspex Premier Service is a fully-integrated hardware and software support program. It is Auspex's most comprehensive support product providing service coverage 24 hours per day, seven days per week, 365 days per year. Refer to the table *Features of Auspex Premier Service* for more information on support coverage. As an FE, you will typically be dispatched to a site whose systems are covered by the APS.

Auspex Advantage Service

The Auspex Advantage Service plan (AAS) is an integrated hardware and software support program. It is designed for customers requiring support during business hours only, Monday through Friday, 8 a.m. to 5 p.m. Very few of Auspex's customers use this type of support, therefore no further details of the program are presented in this manual.

Time and Materials Service (TMS)

Time and Material Service (TMS) is available to customers who choose not to purchase a support program, and for services that are outside the provisions of a support program. Response time to these calls is on a best effort basis and dependent on the availability of local resources.

Features of the Secure Site Support Program	
Secure Site Support Plan Uplift	**Secure Site Addendum**
Special per-incident pricing on the purchase of replacement drives and for the removal of non-volatile memory devices	Special per-incident pricing on the purchase of replacement drives and for the removal of non-volatile memory devices
Supplements the standard Auspex support plans	Supplements the standard Auspex support plans
Customer receives replacement drives automatically, without a purchase order	

Secure Site Support Program

Auspex offers the *Secure Site Support Program* (SSSP) to meet the unique needs of customers who must maintain site security. This program facilitates the economical replacement of parts that cannot be removed from customers' sites.

There are two options to the SSSP: the *Secure Site Addendum*, and the *Secure Site Support Plan Uplift*.

Secure Site Addendum

This addendum to the standard Auspex support plans allows the customer to purchase replacement disk drive assemblies at the then-current list price, less a secure site discount. Purchases are limited to one-for-one replacements and require the return of the drive carrier assembly(ies).

If a customer requires the removal of non-volatile memory devices from Auspex circuit cards, the circuit cards may be returned without those devices for a per-incident charge.

Secure Site Support Plan Uplift

This plan provides for the replacement of a failed drive without the return of the complete failed drive, and without requiring a purchase order. It also allows for the removal of non-volatile memory devices from returning circuit cards for a per-incident charge.

Upon a drive failure, the customer receives a like drive that is new or as-new. If the failed drive is an older technology and a like replacement is not available, the replacement drive will be the next highest capacity available (supported by the system). The customer may be required to return the drive carrier.

Note: The Secure Site Addendum and the Secure Site Support Plan Uplift cannot be mixed at the same site.

Non-Auspex Products

Auspex recommends against adding non-Auspex products in their NetServers

Auspex not responsible for problems from use of non-Auspex products

Customers must identify non-Auspex products to Technical Support

Non-Auspex products may be removed to solve system problem

No analysis or corrective action for non-Auspex products

Non-Auspex Products

Definitions

Non-Auspex means products not listed in the Auspex *U.S. Direct System and Options Price List*.

Products includes, but is not limited to, memory devices, disk and tape products, S-Bus cards, and VME cards.

Policy Statement Highlights

- Auspex cannot prevent customers from adding non-Auspex products to their NetServers, but strongly recommends against doing so.
- Auspex cannot be responsible for support issues arising from the use of NetServers with added non-Auspex products, however, Auspex will provide support to resolve such problems.
- Customers contacting Auspex for support must identify any non-Auspex products to the Technical Support Engineer assigned to the call.
- Resolution of non-Auspex product-related problems may require removal of such products from the system.
- Auspex does not provide root cause analysis or corrective action for such products.

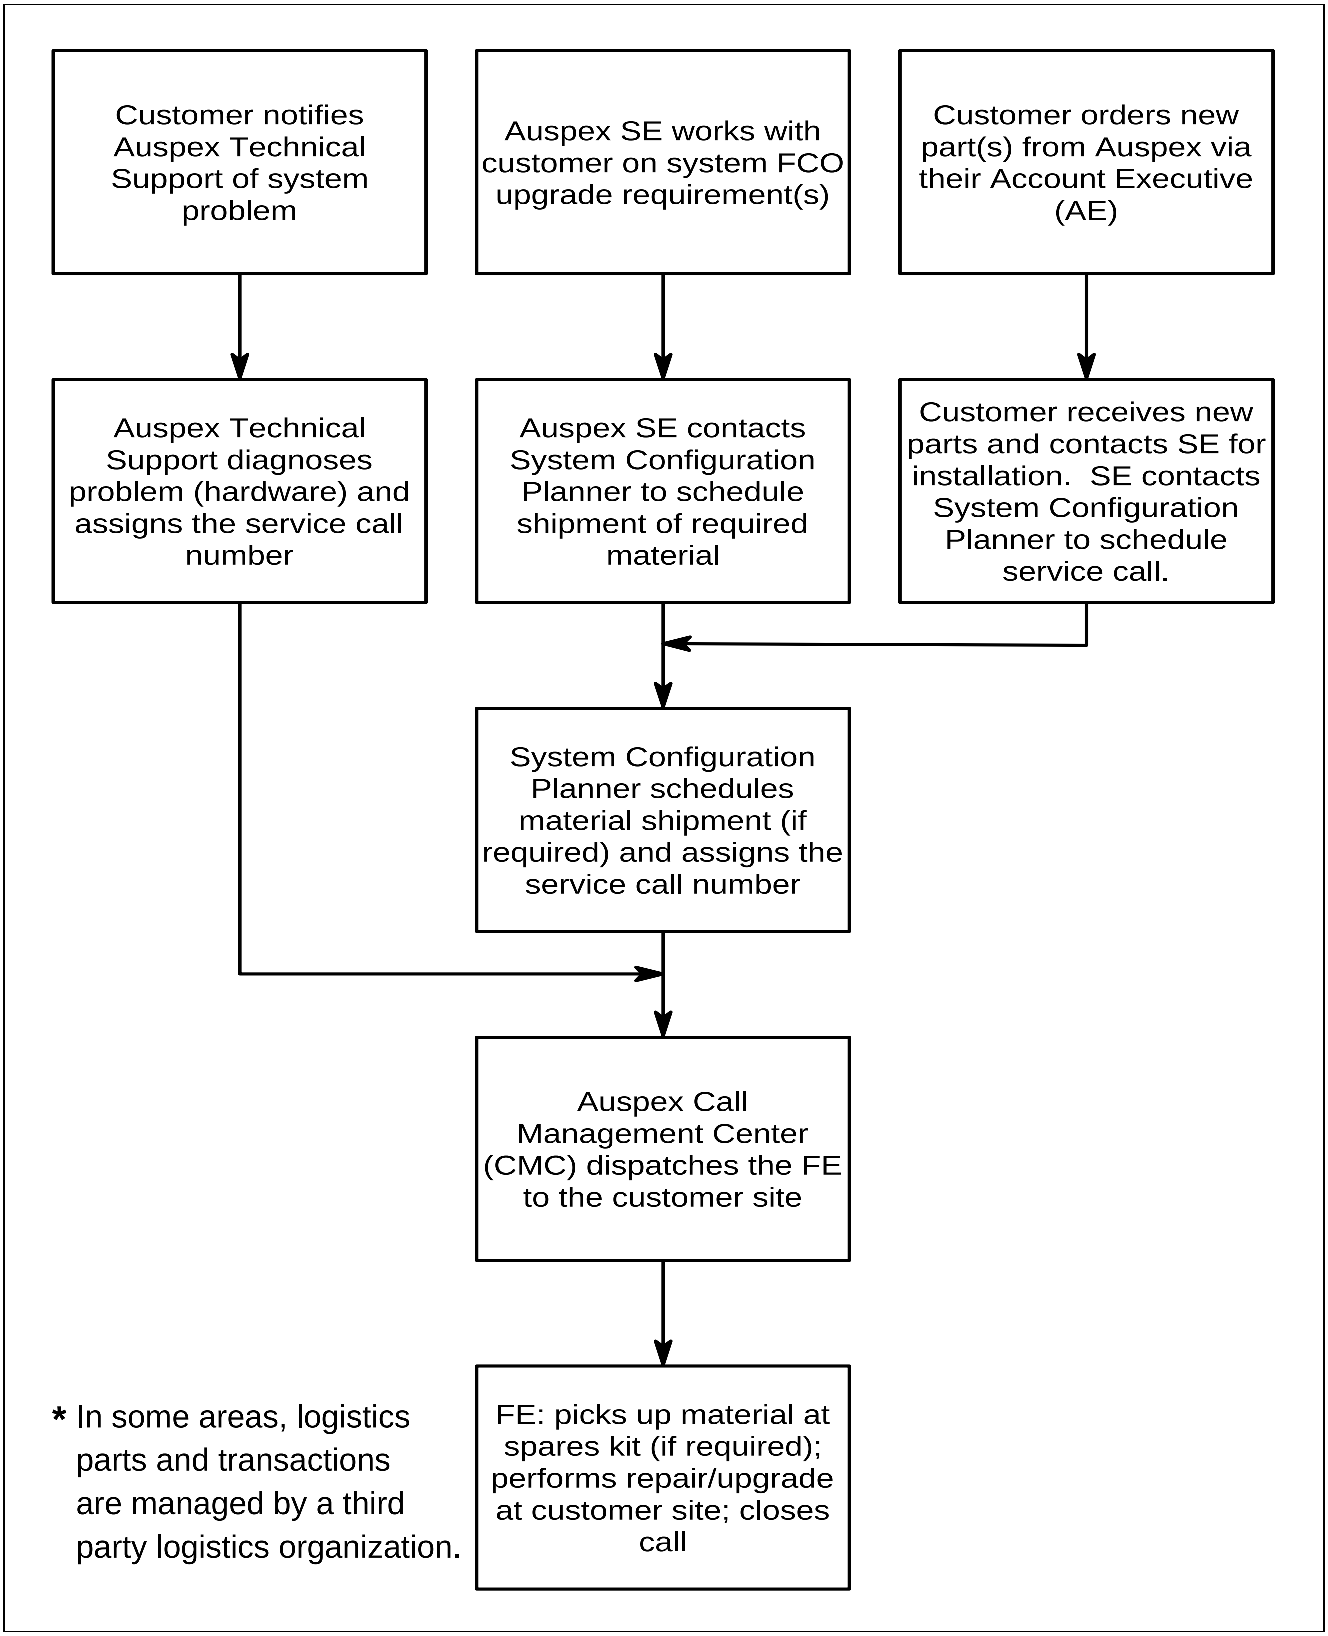

Service Call Process

Service Calls

Generation

Auspex generates service calls based on several different circumstances. The most common are: system failure or sub-optimal performance, installation of a customer-ordered upgrade, and system upgrade based on a Field Change Order (FCO). Auspex generates a service call whenever a Field Engineer participates in a system upgrade, or in resolving a system failure.

Dispatch

The dispatch from Auspex contains the information needed for an FE to make a service call. It is critical that the FE respond to a dispatch within 30 minutes. If there is no response within 30 minutes, same-day service and other critical calls will escalate to the FE’s Dispatch Duty Manager and to the Auspex Service Partner Manager.

Call Closure

Auspex asks that FEs close calls by providing information regarding the service call. Call closure is critical. Without closure information, Auspex cannot make the transactions necessary to maintain database accuracy. Auspex frequently contacts our service providers for help in resolving part-related issues where no closure information was provided. This unplanned work results in extra effort and lost time for both organizations. Auspex feels it is in our mutual interest for FEs to consistently follow through with call closure.

The closure process is simple. The FE prepares the following information prior to calling Auspex:

- Customer’s company name
- Serial number of the affected system(s)
- Part and serial number(s) of the part(s) put into the affected system

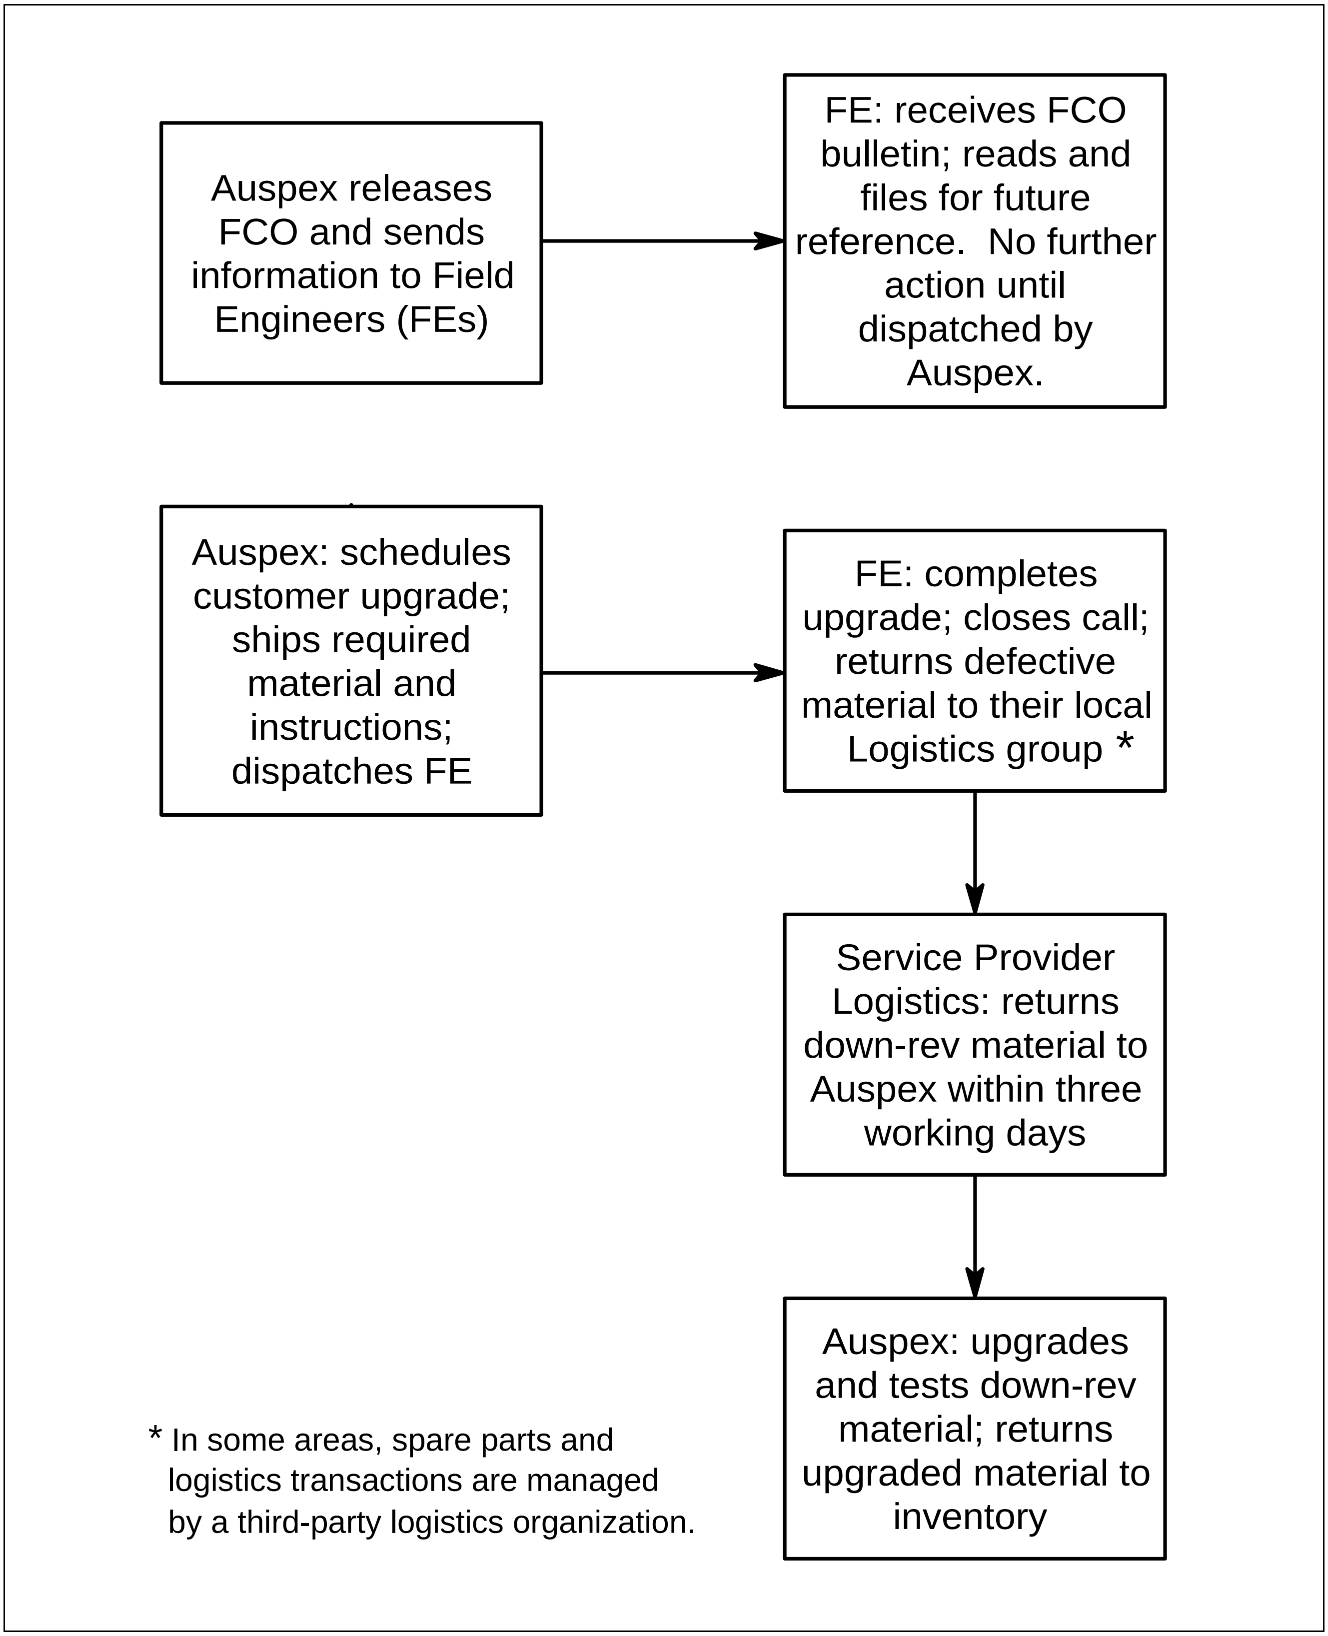
Auspex releases FCO and sends information to Field Engineers (FEs)
FE: receives FCO bulletin; reads and files for future reference. No further action until dispatched by Auspex.
Auspex: schedules customer upgrade; ships required material and instructions; dispatches FE
FE: completes upgrade; closes call; returns defective material to their local Logistics group *
Service Provider Logistics: returns down-rev material to Auspex within three working days
Auspex: upgrades and tests down-rev material; returns upgraded material to inventory
* In some areas, spare parts and logistics transactions are managed by a third-party logistics organization.

FCO Process

- Part and serial number(s) of parts removed from the affected system

After organizing this information, the FE calls 1-800-767-4444 and follows the instructions.

Note: In some locations, spare parts and logistics transactions are managed by a third party logistics company.

Escalation

A call escalation focuses the problem-solving resources of Auspex on a critical customer problem. An FE does not originate an escalation; typically, a Technical Support Engineer (TSE) escalates a call.

Before a call escalates, the responsible TSE works with the customer, and possibly the FE, to solve the problem. If no solution is forthcoming after the initial troubleshooting effort, the TSE will gather appropriate data and escalate the call.

As the escalated call goes to the Advanced Support Management team, key Auspex personnel are automatically notified of the escalation. Auspex utilizes internal resources as necessary to return the failed system to proper operation.

An FE may need to remain on site during an escalation and work with the TSE to complete the task.

Field Change Orders (FCOs)

Field Change Order (FCO) is a term with two definitions. First, it is a document that propagates field change information throughout Auspex, its distributors, ASA customers, and service providers. Included in the bulletin are problem and solution descriptions, a list of parts affected by the change, and the necessary action required of FEs. FEs receive a copy of every FCO field bulletin. Auspex customers without an ASA contract do not receive FCO bulletins.

FCO is also the name of the process of upgrading a customer’s system to comply with a Field Change Order. For more information, see the table *Classes of Field Change Orders (FCOs)*.

Classes of Field Change Orders (FCOs)	
Class of FCO	**Actions Required**
Class III	Affects spares kits only. Auspex upgrades spares kit parts proactively or reactively. Auspex upgrades installed parts of the type affected by the FCO only after they have failed and returned to Auspex.
Class II and Class II-Conditional	Affects spares kits and system-installed parts. This FCO class is often conditional. If so, the FCO bulletin states the condition. An example is a firmware upgrade for the Host Processor *only if* a customer orders the latest generation Network Processor. If the FCO bulletin does not state an implementation condition, the FCO applies to all kits and systems containing the affected part(s).
Class I	The same actions as Class II but implemented with a greater sense of urgency. This class of FCO is very rare and typically deals with safety issues.

System Upgrade Process

Auspex field personnel work with customers to determine how an FCO affects their system(s). After completing the analysis, the responsible Auspex field person establishes the material and time requirements with the customer and with the Auspex System Configuration Planner at Auspex headquarters. The planner arranges the shipment of needed material and rework instructions, and dispatches the FE to the customer site. The rework instructions are comprehensive and allow the FE to work independently. However, if problems arise during an FCO implementation, the FE should contact Auspex Technical Support (1-800-3AUSPEX).

Spares Kit Upgrade Process

Auspex upgrades spares kits using reactive or proactive methods. Reactive upgrades are those implemented in reaction to a new system installation, a customer-purchased upgrade, or an FCO upgrade. The planner responsible for the system upgrade checks the supporting spares kit and arranges any required part replacements.

Proactive upgrade indicates a process of methodically replacing parts in spares kits affected by an FCO. In this case, the responsible planner identifies all kits needing the FCO and prioritizes upgrades based on the level of spares kit activity. Auspex upgrades higher-activity kits ahead of lower-activity kits.

Student notes

3 Documentation Support

Introduction

Auspex designs NetServers to be reliable and serviceable. Because of the simple repair and maintenance requirements, Field Engineers representing Auspex must guard against working from memory.

Note: Take the appropriate documentation with you for every service call.

Objectives

After completing this chapter, you should be able to do the following:

- Name the Auspex documents used for field support and generally describe the information they contain.
- List the documents that support the M2000.

Auspex Documentation

- Field Service Guide
- Field Service Guide Supplement
- Storage Peripherals Manager’s Guide
- Hardware Manuals
- FRU Information Documents
- Hardware Release Notes
- Technical Notes and Tips
- Field Change Order (FCO) Bulletins

Auspex Documentation

Field Service Guide

The *Field Service Guide* is the FE's primary source of reference for information concerning field repair and maintenance. It is organized in indexed chapters focused on the NetServer subsystems: *Processor Boards*, *Drives and Drive Racks*, *Power Systems*, *System Console*, and so forth. Also included are chapters covering system architecture, maintenance, and system configuration.

The *Field Service Guide* is comprehensive, well illustrated, and includes stepped procedures to guide the user through all common field service operations.

Field Service Guide Supplement

Starting with the release of NS 7000/800, Auspex created *Field Service Guide Supplement*s.The supplement has the same design and information as the *Field Service Guide* but only for a specific product or group of products. For example, the *Field Service Guide Supplement for NS 7000/800* provides information and procedures for installing or replacing components of the NS 7000/800 High Density Disk Array (HDDA).

Storage Peripherals Manager's Guide

The Storage Peripherals Manager's Guide (SPMG) describes external tape storage devices used with Auspex NetServers. Also included is detailed information on attachment options, installation, and device numbering. The M2000 version of this document is entitled *Auspex NetServer M2000 Peripheral Devices Manager's Guide*.

Student notes

Hardware Manuals

Hardware Manuals provide procedures for installing specific NetServers and their associated hardware components and peripherals. They guide the user through the entire installation process from unpacking and setting up the components through the power-on sequence. The M2000 version of this document is entitled *Auspex NetServer M2000 Hardware Installation Guide*.

FRU Information Documents

FRU Information documents append the *Field Service Guide* and provide the same type of information. They document the changes in the Auspex product line since the previous *Field Service Guide* release. Customers do not receive the *FRU Information* documents. *FRU Information* documents are tied to software releases and contain the software version number in the document title; for example *1.8.1 FRU Information*.

Hardware Release Notes

Hardware Release Notes are provided to Auspex customers and are similar to the *FRU Information* documents in some respects. They document recent changes in the Auspex product line and they are tied to specific software releases. The important difference is that the *Hardware Release Notes* do not contain repair or replacement procedures. Auspex does not want untrained personnel performing repair or maintenance on NetServers.

Technical Notes and Tips

Technical Notes and Tips (“*Tech Tips*”) are released and distributed as required. They typically describe a specific problem currently being encountered in the field and offer a tested solution. They may also provide updated or more detailed information on a subject already documented in the *Field Service Guide*. Field Engineers should file *Tech Tips* in a binder for future reference.

M2000 documents	
Auspex part number	**Document title**
850566	*Auspex NetServer M2000 Site Planning Guide*
850569	*System Manager's Guide*
850570	*Auspex NetServer M2000 Hardware Installation Guide*
850571	*Auspex NetServer M2000 Field Service Guide*
850572	*Auspex NetServer M2000 Peripheral Devices Manager's Guide*

Field Change Order (FCO) Bulletins

Field Change Order (FCO) bulletins are described in the *Field Support* chapter. Auspex distributes the bulletins to Field Engineers, Auspex System Engineers, distributors, and ASA customers. Auspex recommends that FEs read the FCO bulletins carefully and file them in their *Field Service Guides*. Rework instructions are not distributed with the FCO bulletins; they are shipped with the parts to be used for the FCO upgrade.

M2000 Documents

The table *M2000 documents* provides a list of Auspex documents that support the M2000.

Student notes

4 Product Overview

Introduction

This chapter defines the M2000 hardware platform and generally describes its major components, architecture, interconnection, and packaging.

Objectives

After completing this chapter you should be able to do the following:

- Name and describe the M2000 modular building blocks.
- Describe M2000 system architecture and interconnection.
- Describe the M2000 system packaging.

Product Definition

Scalable file server with stackable, modular physical configuration. The modular building blocks are:

Host Module provides host functions; supports system console; runs Solaris 2.6

E-Box provides network, file, and storage processing

High Density Disk Array (HDDA) provides disk storage

Power Base provides N+1 redundant power for the server

Scalable Coherent Interface (SCI) interconnects the Host module and E-Box(es)

Environmental Monitor Network (EM Net) interconnects the E-Box(es), HDDA(s), and Power Base

Product Definition

System Overview

The M2000 is a scalable file server with a stackable, modular physical configuration. The modular building blocks are:

- the *Host module* that provides host functions and runs Solaris 2.6
- the *E-Box*, a module that provides network, file, and storage processing
- the *High Density Disk Array* (HDDA) that provides disk storage
- the *Power Base* that provides N+1 redundant power for the server

An M2000 system includes a single Host module plus one, two, or three *nodes*. Each node consists of one E-Box, one to three HDDAs, and a Power Base. The Host stacks on top of the E-Box in the base node. A later chapter, *System Configuration and Cabling*, provides configuration detail.

A *Scalable Coherent Interface* (SCI) interconnects the Host Processor and all E-Boxes. An *Environmental Monitor Network* (EM Net) interconnects the E-Boxes, the HDDAs, and the Power Base.

Additional system features:

- The Host Processor supports the system console (17-inch color monitor, mouse, and keyboard), and a modem for dial-up access to the server.
- Tape backup devices are supported but not integrated in the system stack.

Hardware Overview

The hardware components reviewed in this section are described in more detail in later chapters.

Student notes

Host Processor

The Host module contains a Sun Ultra AXi motherboard that runs the Solaris 2.6 operating system. The Host motherboard supports scalable main memory, dual SCSI channels, and standard I/O connectivity. The Host also supports the system console. Field Replaceable Units (FRUs) for the Host are described in the *Processing Subsystem* chapter, later in this manual.

The Host Processor provides standard host functions such as **boot**, **fsck**, and **mount**. It also provides a platform for applications such as backup, GUI administration tools, Auspex DataGuard, and command line utilities.

E-Box

The E-Box contains an Auspex-designed motherboard, optimized for the Auspex Functional Multiprocessing (FMP) architecture. The motherboard contains two CPUs which share a scalable main memory. One processor executes the Network Processor (NP) function, the other executes the File and Storage Processor (FSP) functions. Peripheral Component Interconnect (PCI) cards provide RAID and SCSI control, and network interfaces in various configuration options.

HDDA

The HDDA contains up to four drawers, each with a DC-to-DC converter, and up to seven disk drives.

The SCSI cabling allows some flexibility in configuring the storage subsystem. A three-channel SCSI controller can support four drive drawers by using a daisy-chain cable to connect two drive drawers to one channel (port). This supports more disk drives for each node, but will impact performance in storage-intensive applications.

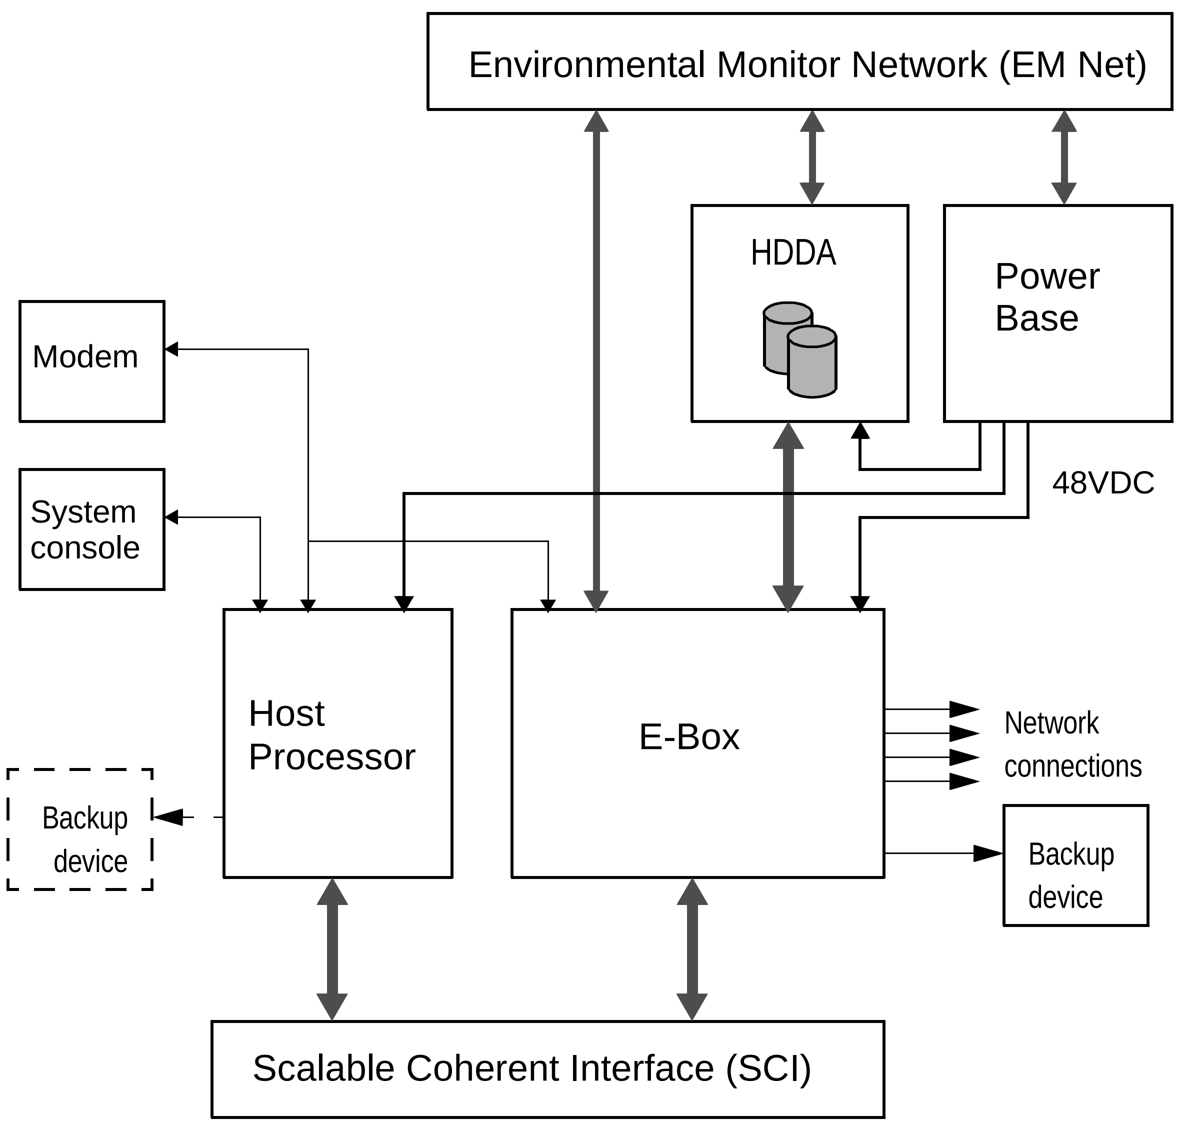

M2000 system hardware architecture

Power Base

The Power Base consists of a crate structure and one or two Power Shelves. Each Power Shelf has a pluggable PDU with its own power cord, and up to three, hot-pluggable, 48VDC bulk power supply modules. A single-HDDA configuration requires only one Power Shelf. Additional HDDAs require a second Power Shelf. Each Power Shelf is independently N+1 redundant.

Cables carry the 48VDC power to each module in the stack where it is converted to the required voltages.

Different AC power cord options allow connection with different wall outlet configurations used around the world.

System Hardware Architecture

The diagram *M2000 system hardware architecture* shows the server design at the most basic level. Later chapters describe each major component in detail.

System Interconnection

The diagram *M2000 system hardware architecture* provides a general look at system interconnection. M2000 interconnection is illustrated and described in more detail in the later chapter *System Configuration and Cabling*.

Processing Subsystem

The Host and E-Box(es) are interconnected with SCI cables. These cables connect to PCI-SCI cards in each module. The PCI and SCI implementations are described in the chapter *Processing Subsystem*, later in this manual.

The system console monitor connects to the Host via a PCI-SVGA card. The keyboard and mouse for the system console connect to the rear panel

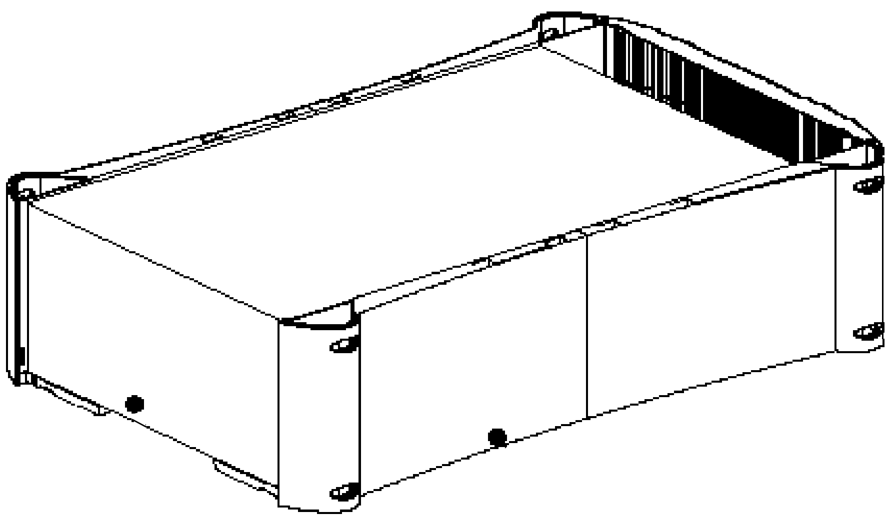

M2000 basic stack element

of the Host. The Host also supports a modem for remote access to the system.

The E-Box supports Network Interface Cards (NICs), so all client network connections are made at the rear panel of the E-Box(es).

Storage Subsystem

The HDDA(s) connects to SCSI controller cards in the E-Box using standard and multi-drop (daisy-chain) SCSI cables. The optional Differential SCSI card provides two channels to connect to a backup device, such as a tape library.

Power Subsystem

Cables carry 48VDC from the Power Shelf to the Host, E-Box(es), and HDDA(s). One Power Shelf with three bulk modules will supply power to a single-HDDA node (Host, E-Box, HDDA). Additional HDDAs require a second Power Shelf with two or three bulk modules.

Environmental Monitor Subsystem

The Environmental Monitor Network (EM Net) cables interconnect the E-Box(es), HDDA(s), and Power Shelf(s). These category-5 cables plug into RJ45 connectors on the rear of each module and are daisy-chained down the stacked configuration. A cable connected from the bottom module to the top module completes a ring topology. For more information, refer to the chapter *Environmental Monitor Subsystem*, later in this manual.

System Packaging

The M2000 uses *virtual* packaging. The frame is integral to the component modules and requires very little conventional mounting hardware. Preformed materials snap onto the stack rails to serve as side skins, and a cover piece snaps onto the top stack element. Locking casters are part of the Power Base crate structure.

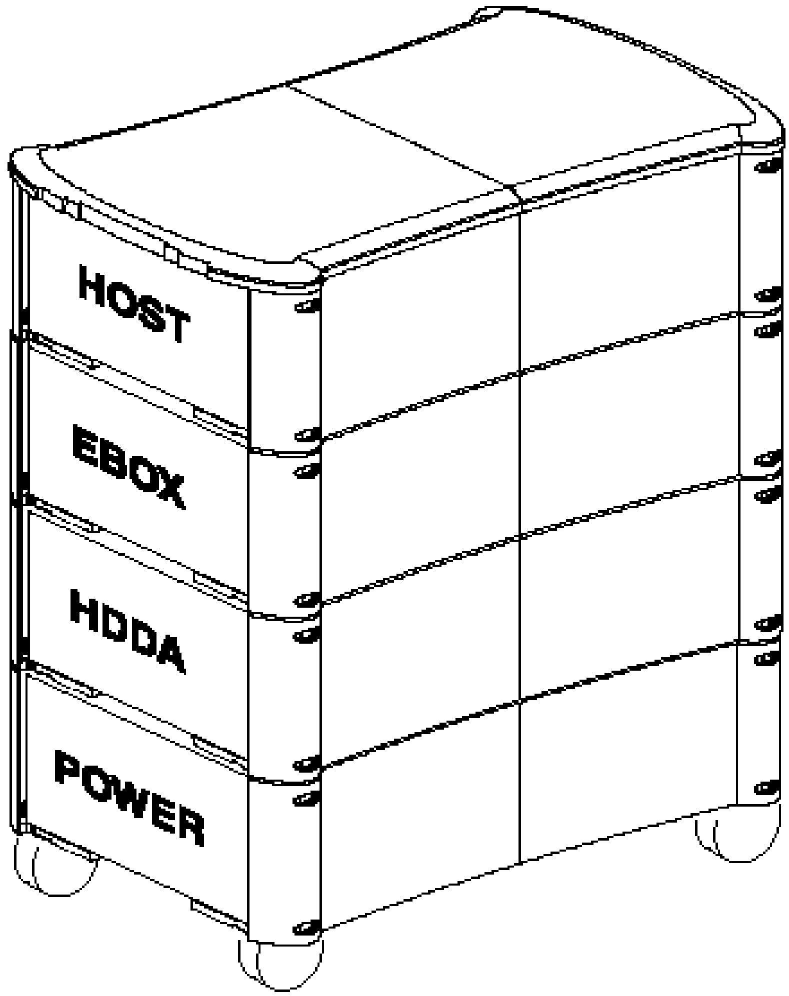

M2000 base node packaging

The Host, E-Box, and HDDA use a common shelf structure. The Power Base uses a separate shelf structure. All modules use mirror-image side skins and, except for the HDDA, all modules have a front bezel. Refer to the figure *M2000 base node packaging* for pictorial detail.

Student notes

5 System Configuration and Cabling

Introduction

This chapter describes and illustrates the M2000 server configuration options, and intermodule and internodal connections.

Objectives

After completing this chapter you should be able to do the following:

- ▲ Generally describe M2000 cable connections.
- ▲ Differentiate between the master and expansion E-Boxes.
- ▲ Distinguish base nodes from expansion nodes.
- ▲ Determine the node configurations achieved by various combinations of NICs and SCSI controllers.

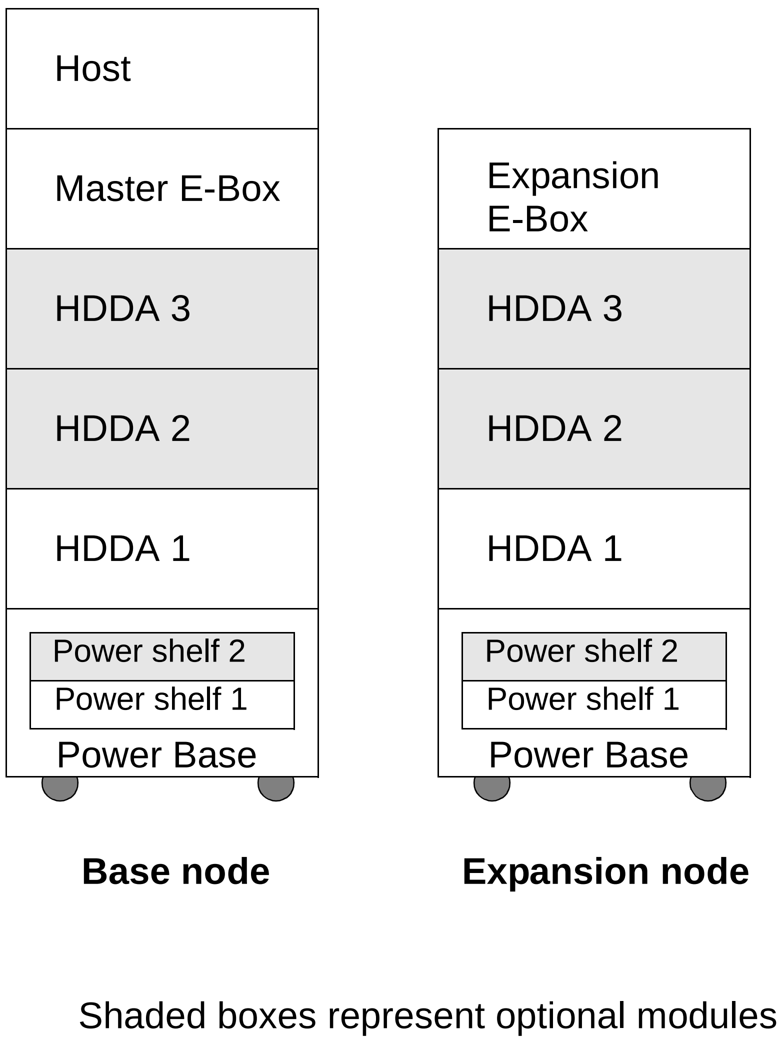

Node configuration stacks

System Configuration

An M2000 system node consists of one E-Box, one to three HDDAs, the Power Base, and all associated intranodal cables.

One E-Box in each system is designated as the *master* E-Box because it contains the Environmental Monitor master board. In a multi-node system, the master E-Box, its HDDA(s), and its Power Base comprise the *base* node. The system Host module is stacked on top of the master E-Box. The M2000 supports only one Host module. Refer to the figure *Node configuration stacks* for pictorial detail.

An *expansion* node consists of an expansion E-Box, one to three HDDAs, the Power Base, and all associated cables. An expansion E-Box differs slightly from the master E-Box:

- ▲ Only the master E-Box contains the Environmental LON Processor (E-LP) board. This board is the master of the EM Net so there is only one per system.
- ▲ The expansion and master E-Boxes have different transceiver boards. These transceivers are used for communication over the EM Net.

An M2000 system supports one base node and, optionally, one or two expansion nodes.

Note: M2000 node and system configuration is flexible and this chapter does not document all possible permutations.

An important factor for determining node configuration is the proportion of NIC and SCSI controller cards in the E-Box. With Auspex's traditional enterprise servers, customers can linearly scale server performance by adding pairs of Storage Processor and Network Processor boards. These large servers have a 14-slot backplane which provides ample expansion capability. The M2000 E-Box has only four available PCI slots for NIC and SCSI cards.

The process of scaling a node beyond 2x2 (two NICs and two SCSI controllers) becomes subtractive: you must give up a network card to gain a SCSI card and vice versa.

M2000 node configuration options			
	Network Interface Cards (NICs)		
SCSI cards	**1**	**2**	**3**
1	1 Enet or 1 FDDI 21 drives; 378 GB, or 28 drives; 504 GB daisy-chained	2 Enet, or 1 Enet + 1 FDDI, or 2 FDDI 21 drives; 378 GB, or 28 drives; 504 GB daisy-chained	3 Enet, or 2 Enet + 1 FDDI, or 1 Enet + 2 FDDI, or 3 FDDI 21 drives; 378 GB, or 28 drives; 504 GB daisy-chained
2	1 Enet or 1 FDDI 42 drives; 756 GB, or 56 drives; 1 TB daisy-chained	2 Enet, or 1 Enet + 1 FDDI, or 2 FDDI 42 drives; 756 GB, or 56 drives; 1 TB daisy-chained	
3	1 Enet or 1 FDDI 63 drives; 1.1 TB or 84 drives; 1.5 TB daisy-chained		

Notes:

1.) The storage capacities listed assume the use of 18 GB drives.
2.) *Enet* can be either Quad 10/100 Base-T, or Gigabit Ethernet. *FDDI* can be either single attachment (SAS) or dual attachment (DAS).
3.) Daisy-chained means HDDA drive drawers are connected together and to one port (channel) on the SCSI controller. Refer to the later chapter Storage subsystem for more information.

One countering advantage is the M2000 SCSI controller has three ports (channels) which can be mapped to the four drawers of an HDDA. This three-to-four mapping requires a daisy-chain cable to connect two of the drive drawers to one SCSI channel. Assuming the use of 18 GB disk drives, two SCSI controllers can support up to 1 TB of storage (56 disk drives, each storing 18 GB). The table *M2000 node configuration options* provides specific data.

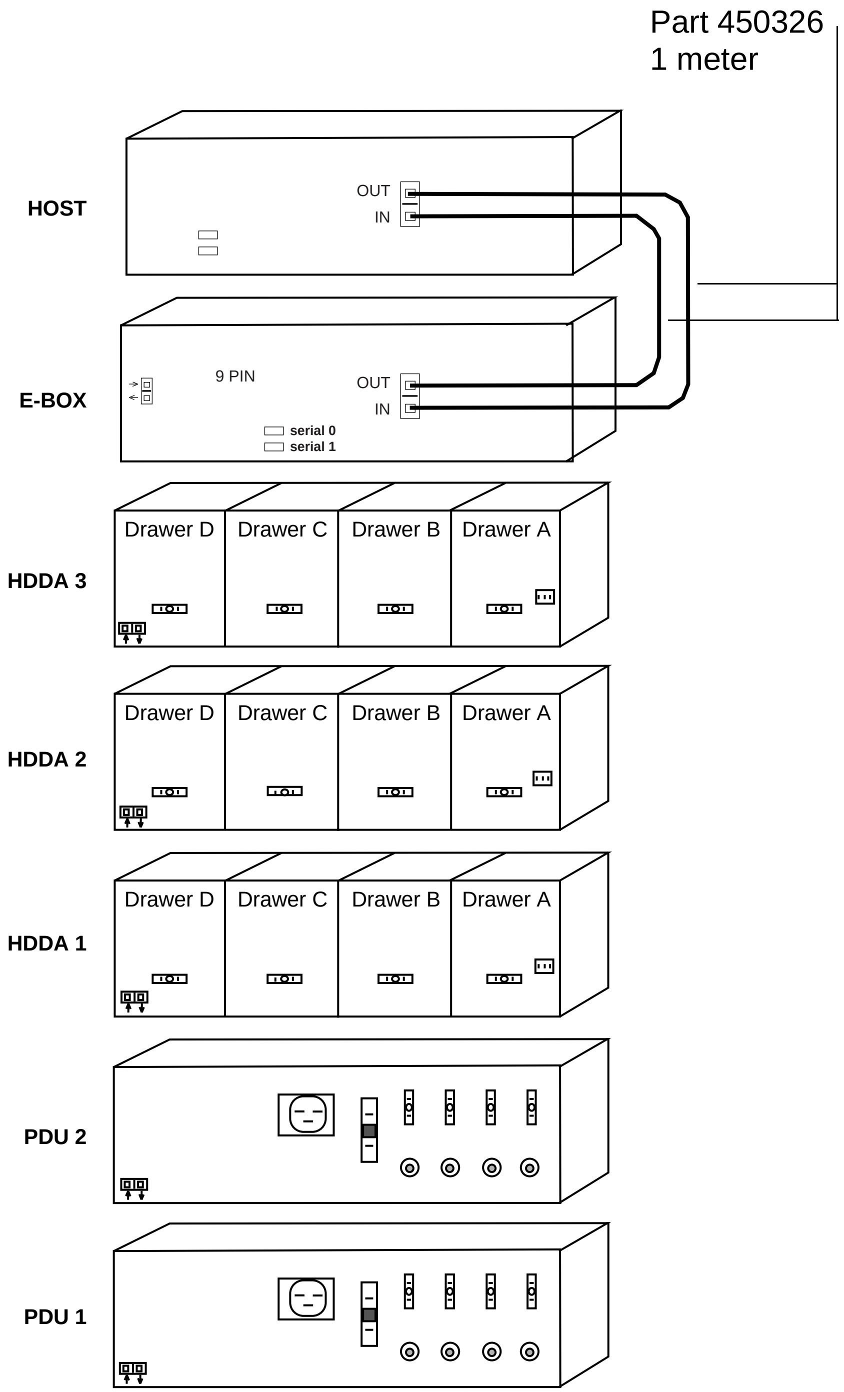

Single node SCI ring cabling

System Cabling

M2000 system cabling is categorized as follows:

- ▲ Processor subsystem cabling (SCI cables)
- ▲ Storage subsystem cabling (SCSI cables)
- ▲ Power subsystem cabling
- ▲ Environmental Monitor subsystem cabling (EM Net)

The illustrations in this chapter are consolidated here to provide a reference and to promote an understanding of system interconnection. Some of the illustrations also appear in later, subsystem-specific chapters.

Processing Subsystem Cabling

SCI cables interconnect the Host and E-Box(es) in an M2000 system. Single- and three-node interconnections are shown in the illustrations *Single node SCI ring cabling*, and *Three node SCI ring cabling*.

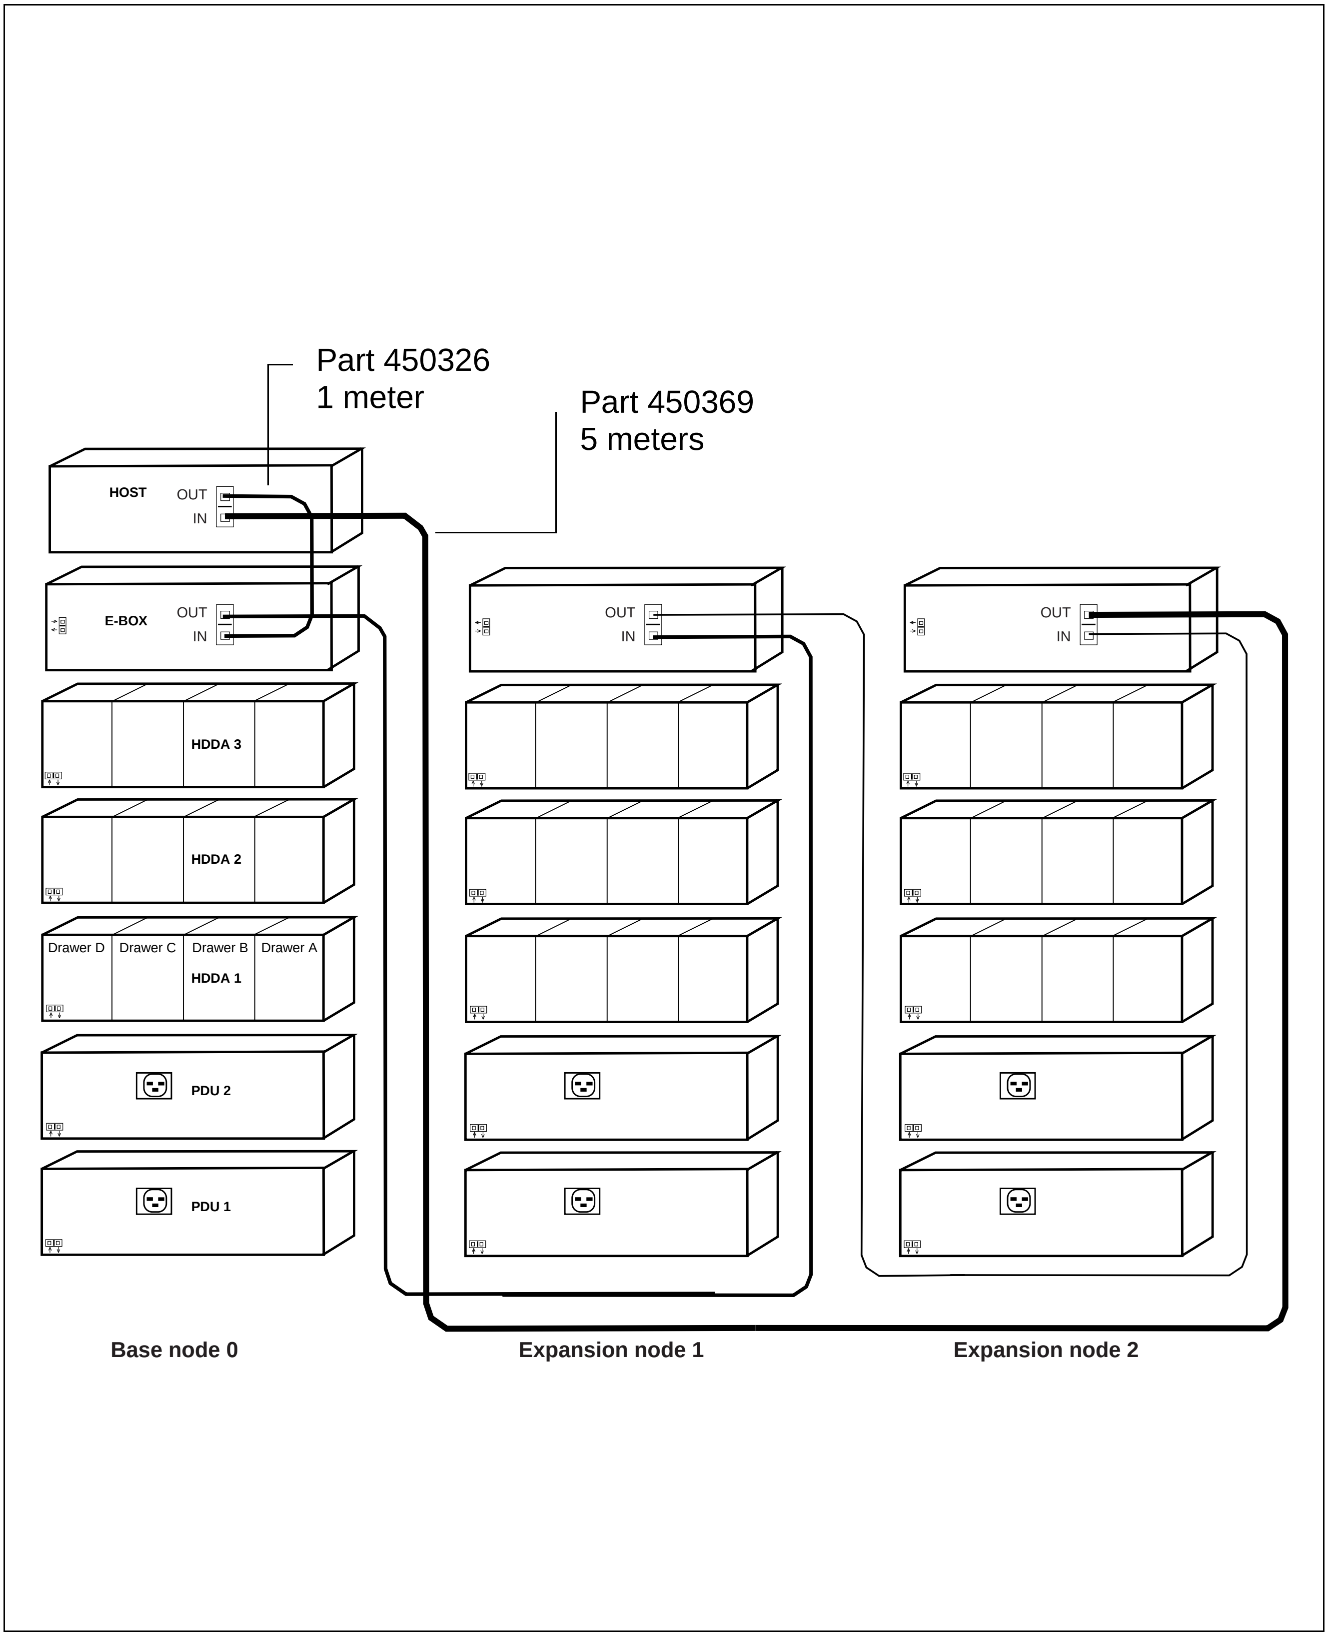

Three node SCI ring cabling

Student notes

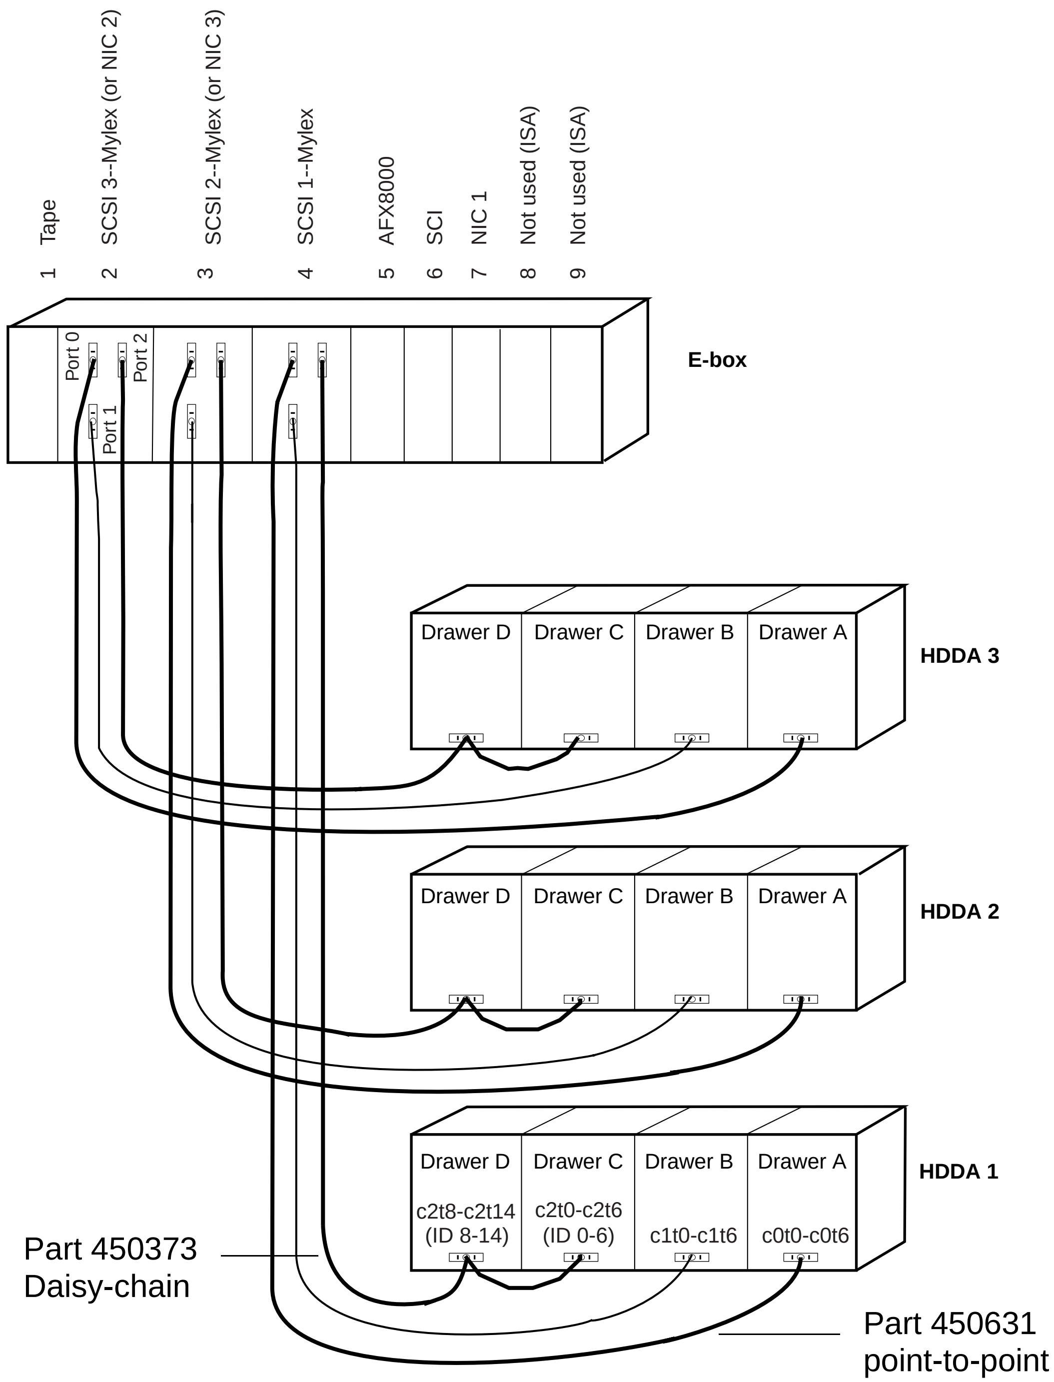

Base node SCSI connections

Storage Subsystem Cabling

The M2000 uses two types of SCSI cables:

- Standard, point-to-point cables that connect a SCSI controller port to an HDDA drawer (up to seven disks).
- Daisy-chain cables that connect a SCSI controller port to two HDDA drawers (up to 14 disks).
  - Daisy chaining two HDDA drawers requires specific termination and SCSI ID settings to the affected drawers. The chapter *Storage Subsystem*, later in this manual, describes these settings.

Note: In the M2000, SCSI busses (cables) do not cross nodes.

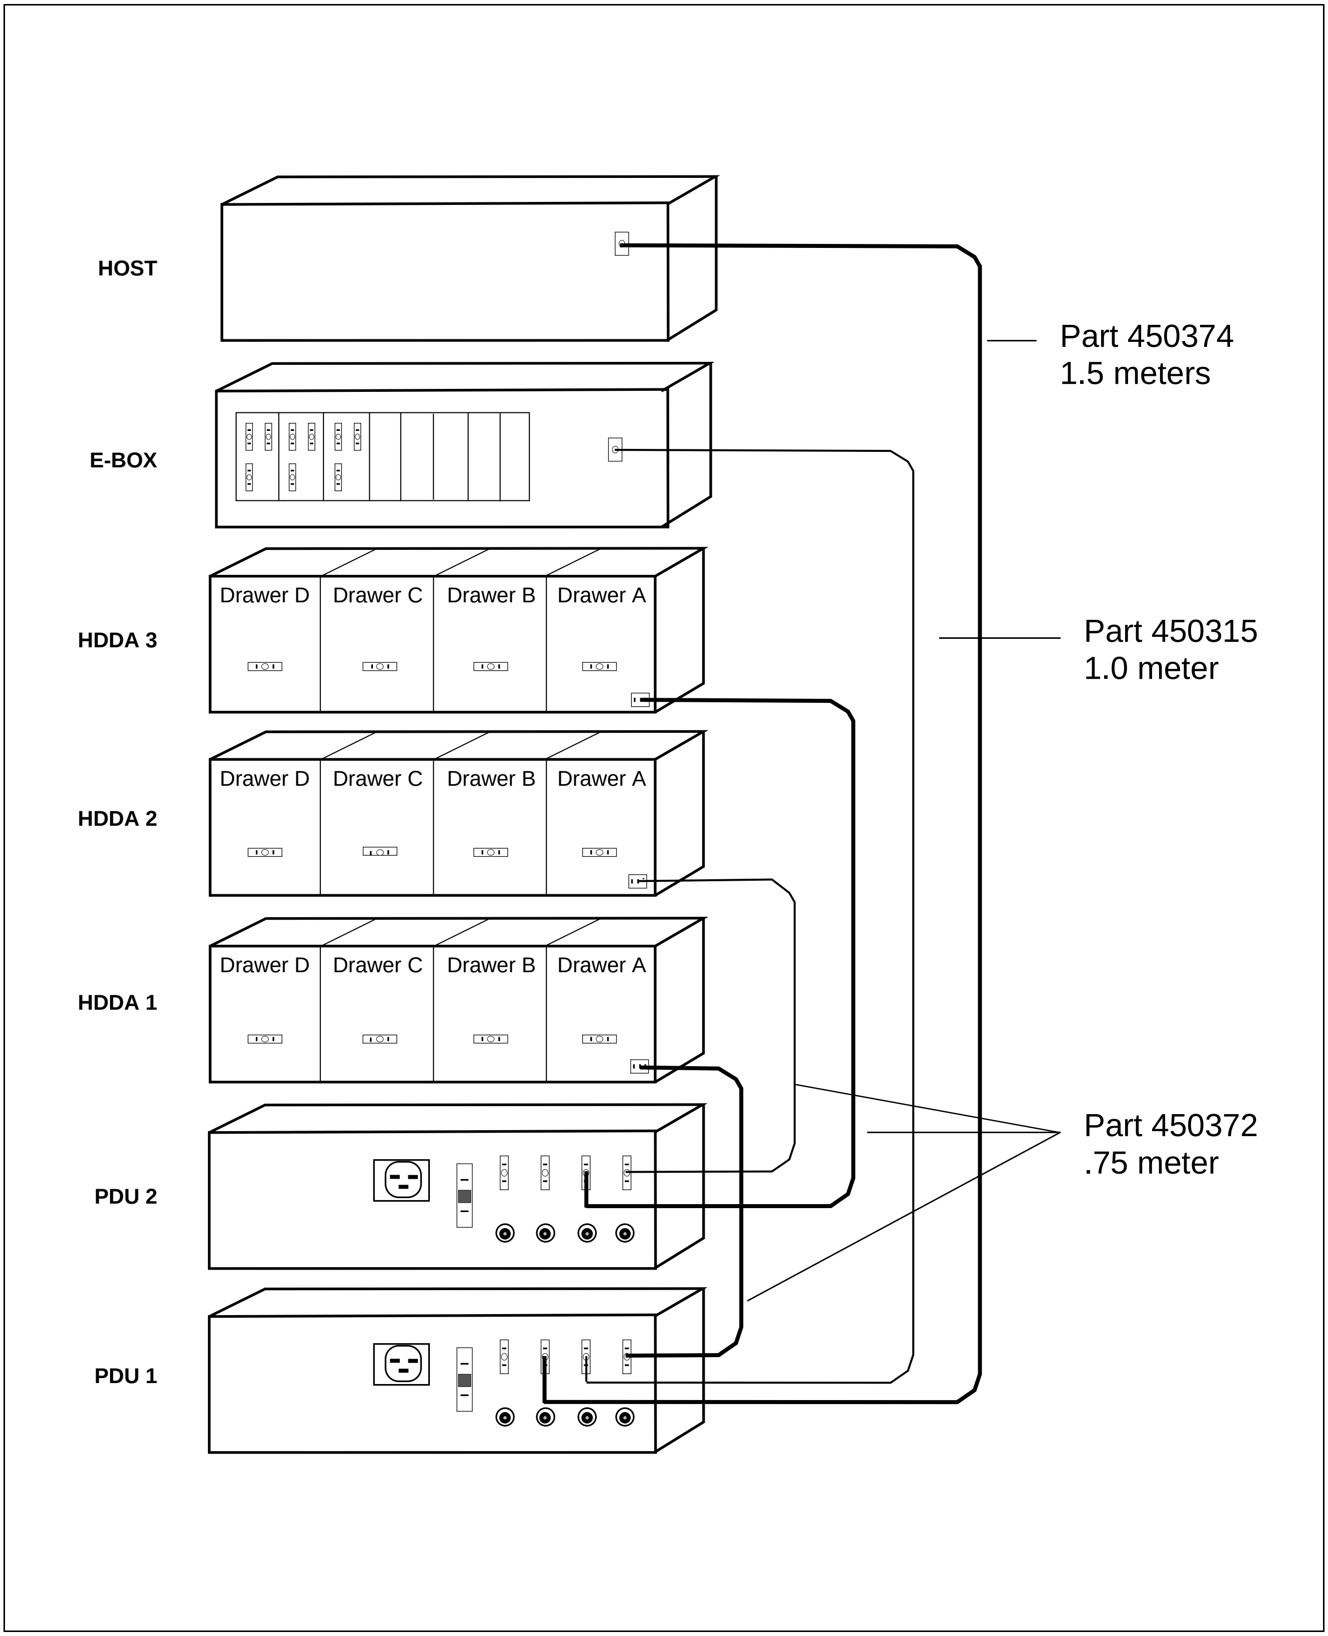

Base node power connections

Power Subsystem Cabling

Power cables (48VDC) route from the Power Distribution Unit(s) (PDU) to the various modules in the node.

Note: In the M2000, power cables do not cross nodes.

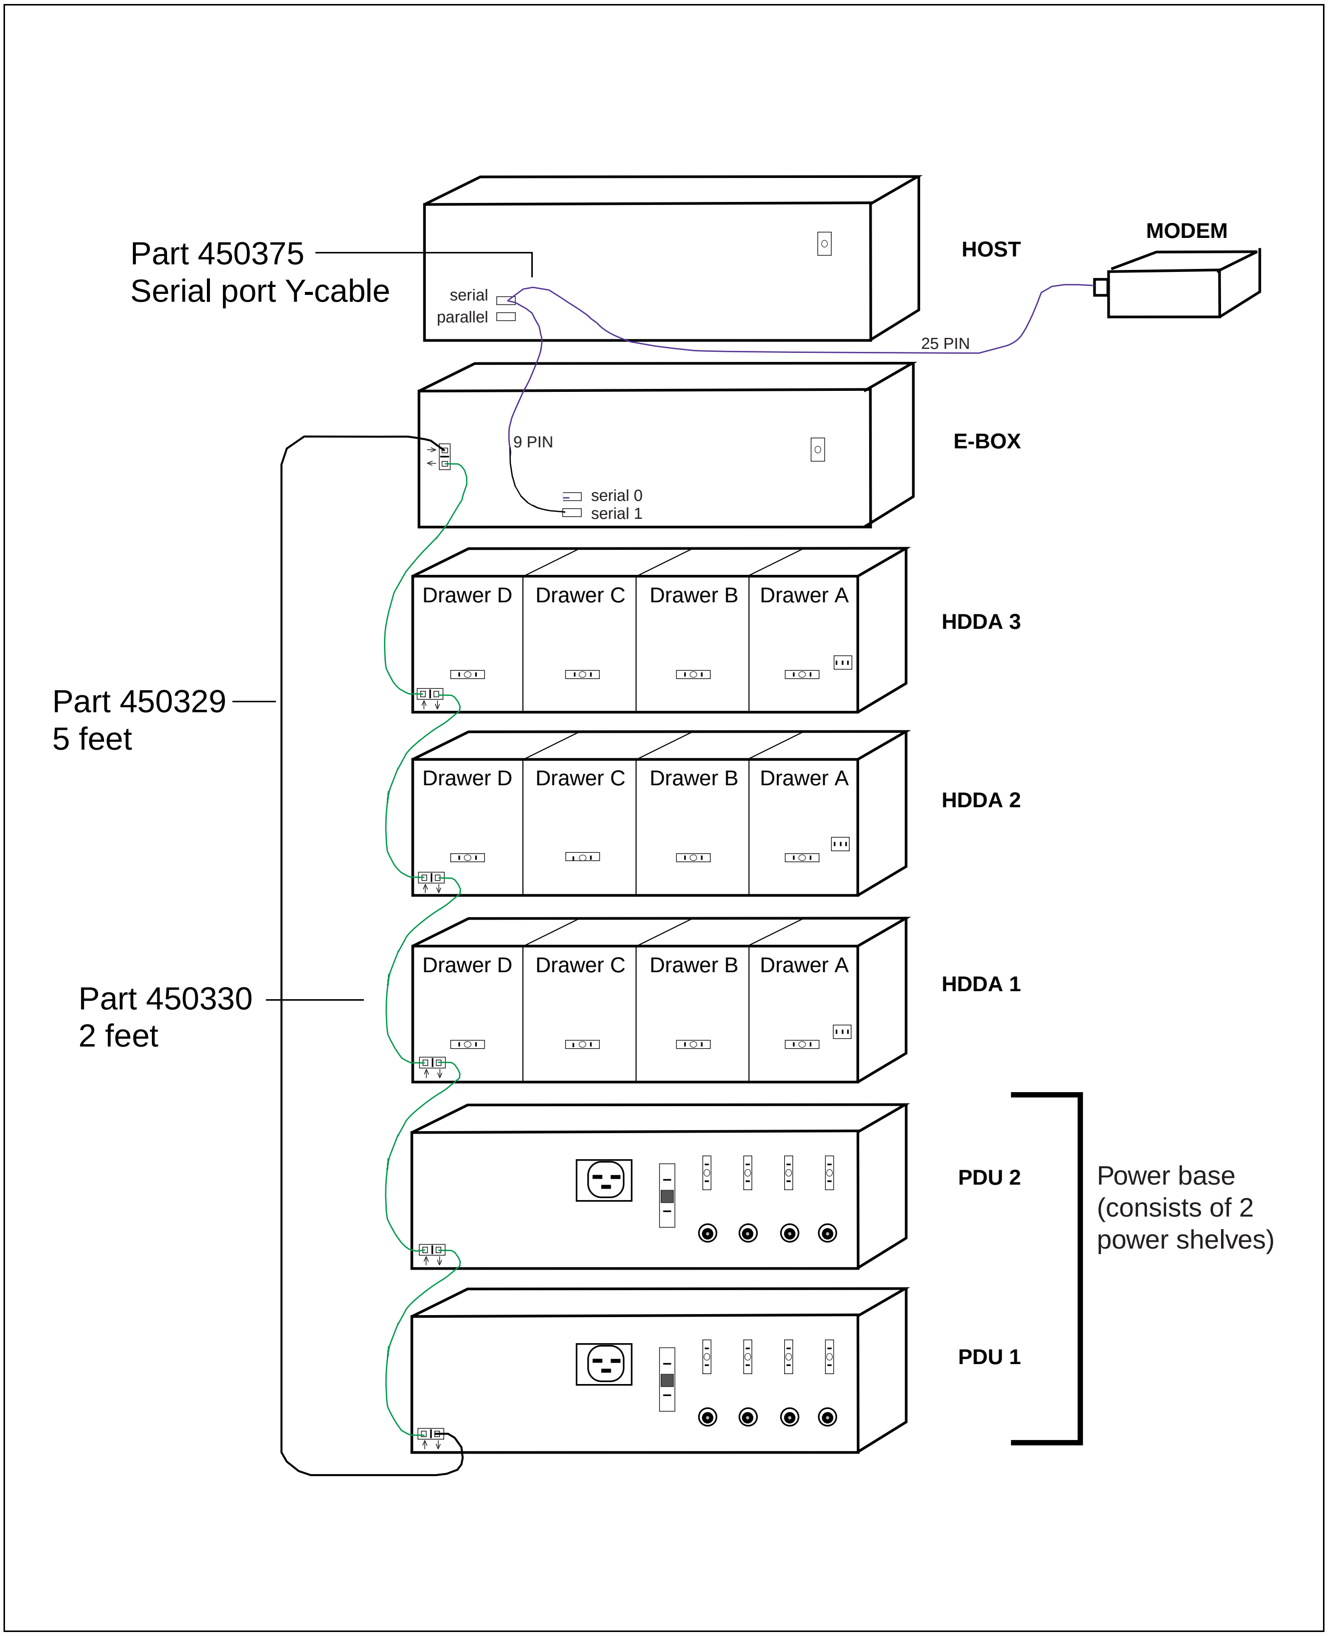

Single node EM Net cabling

Environmental Monitor Subsystem Cabling

The Environmental Monitor Subsystem is highly distributed and, in a multi-node system, requires a lot of cabling. The EM Net directly or indirectly involves every module in the M2000.

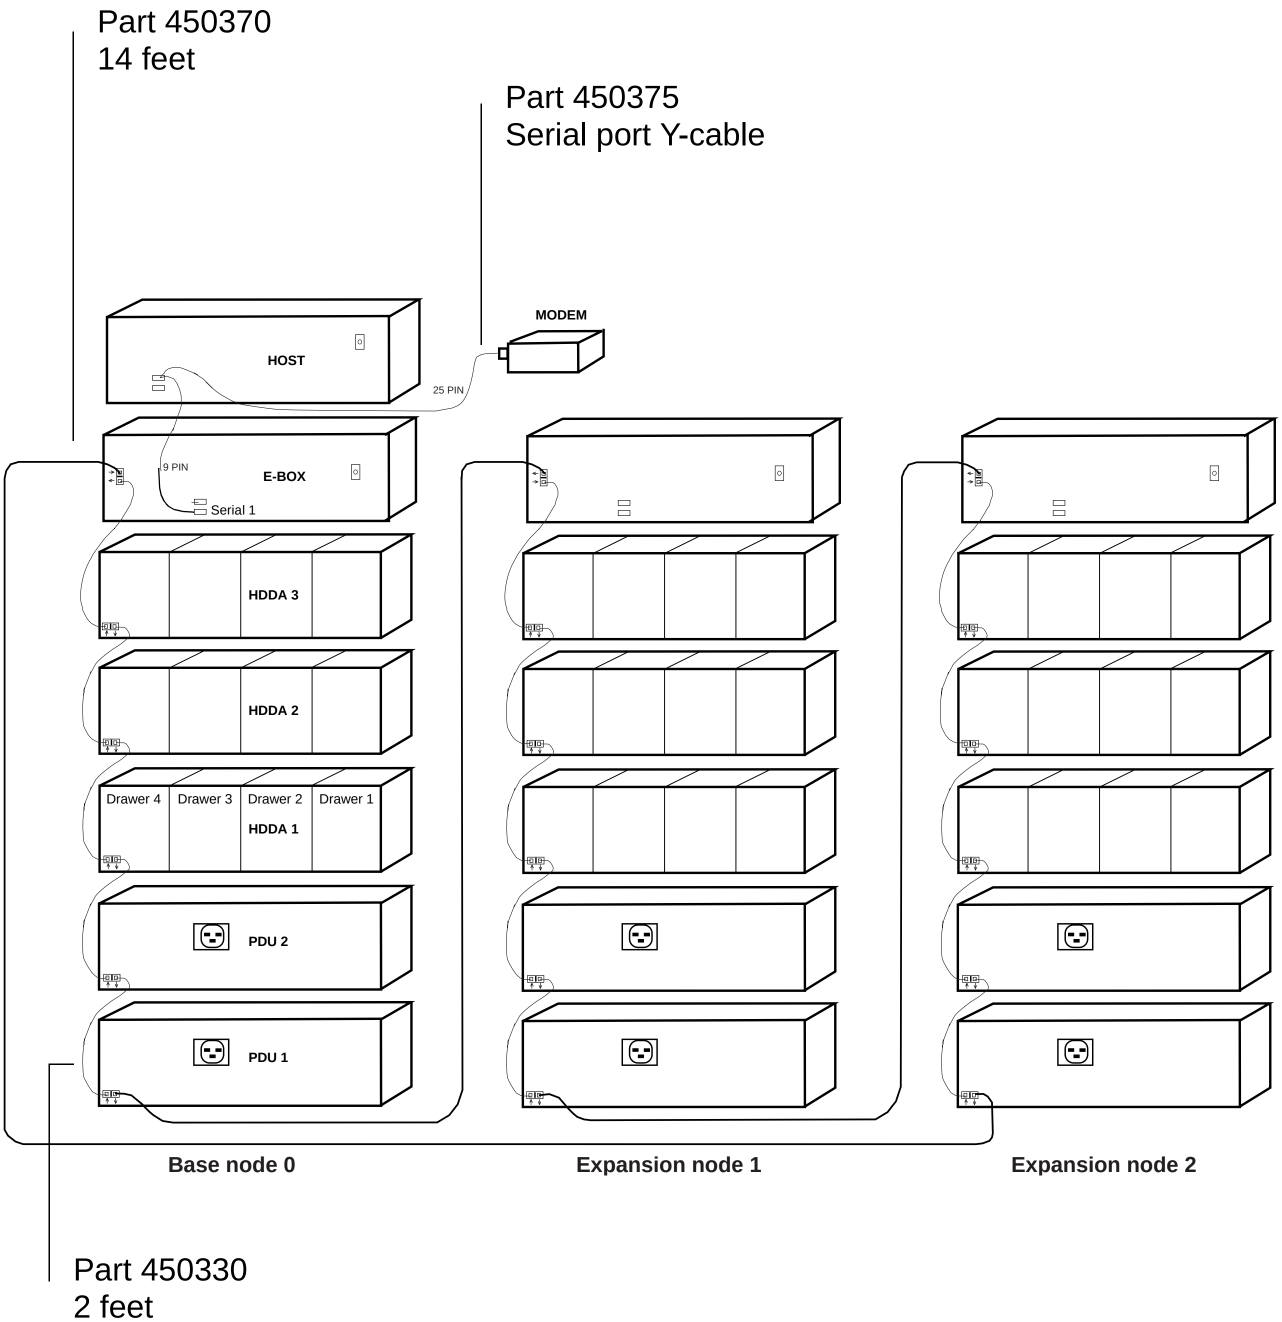

Three node EM Net cabling

6 Exercise: Introduction to Auspex and the M2000

Introduction

Thus far, you have received a lot of information about Auspex and an introduction to the M2000. This exercise is not a test of what you remember from the lecture; it is practice for the beneficial field work habit of referencing information and subsequently applying it.

Instructions

Answer the following questions independently and using the appropriate manual(s). Your instructor will discuss the answers in detail and ensure that you are comfortable with the information before proceeding to new material. Good luck and have fun with the exercise.

1. State the Auspex definition of quality.

2. The typical customer you service as an Auspex FE will have which support program for their servers?

3. Why is call closure important? What information is typically required when you are closing an Auspex service call?

4. What type of information is contained in the Auspex *Field Service Guide*? Describe *Field Service Guide Supplements*.

5. What information is contained in the *Storage Peripherals Manager's Guide*?

6. What is the primary difference between the *FRU Information* documents and the *Hardware Release Notes*? Why does Auspex not supply its customers with the *FRU Information* documents?

7. Name the M2000 modular building blocks.

8. Which of the M2000 modules does the SCI interconnect? Which of the modules does the Environmental Monitor interconnect?

9. In a multi-node M2000 system, where is the Host module located? How many Host modules does the M2000 support?

10. In a multi-node M2000 system, which subsystem's cabling connects across nodes?

11. State the pros and cons of using daisy-chain SCSI cables in the storage subsystem.

12. What factor limits NIC and SCSI controller scalability?

7 Processing Subsystem

Introduction

This chapter describes the Host, the E-Box, and the add-on boards that comprise the M2000 processing subsystem.

Objectives

After completing this chapter, you should be able to do the following:

- Describe the Host module from physical and functional perspectives.
- Name and describe the Host FRUs.
- Generally describe the Euclid Common Hardware Platform (ECHaP).
- Name and describe the E-Box FRUs.

Host Module

Provides basic host functions

Platform for applications (backup, DataGuard, etc.)

Uses Sun AXi motherboard; runs Solaris 2.6

Motherboard supports 128 to 512 MB main memory, dual FAST-20 Wide SCSI channels, six 32-bit PCI slots, standard I/O devices

Communication with E-Box(es) is via PCI-SCI card; system console support is via PCI-SVGA card, PS/2 mouse and keyboard connections

Internal environmental monitoring: temperature monitoring, voltage monitoring, fan control

In basic configuration, Host supports boot disk drive, CD-ROM drive, two hot-pluggable power supplies, SCI and SVGA cards, 128 MB main memory

Options include memory upgrades to 256 or 512 MB, and backup boot drive

Host Module

Functions and Features

The M2000 Host module provides basic host functions such as **boot**, **dump**, **fsck**, and **mount**. It also provides a platform for applications such as backup, GUI administration tools, Auspex DataGuard, command line utilities, internal (Host only) environmental monitoring, and diagnostics. Communication with the E-Box(es) is via an PCI-SCI card and SCI cabling.

The Host module uses a SUN Ultra AXi motherboard powered by a single Sun UltraSPARC-IIi 270 MHz processor. The processor runs the Solaris 2.6 operating system. The Host motherboard supports the following features:

- The main memory is configurable from 128 MB to 512 MB using 64 MB DRAM DIMMs
- 256 KB of processor cache
- On-board, dual FAST-20 Wide SCSI channels (one channel for internal devices, one for externally attached devices)
- Six 32-bit PCI slots. Only five PCI slots can be used because slot 6 is partially obstructed by the P/S 2 mouse and keyboard connectors. Supported PCI cards include SCI, SVGA, and, optionally, Differential SCSI.
- Standard I/O devices: two serial, one parallel, keyboard, and mouse ports

The Host is packaged in a custom-designed, highly serviceable enclosure and is stacked on top of the master E-Box in the base node.

Environmental Monitoring

Host environmental monitoring is performed by the Open Boot Process and is unrelated to any node Environmental Monitor functions. The Host

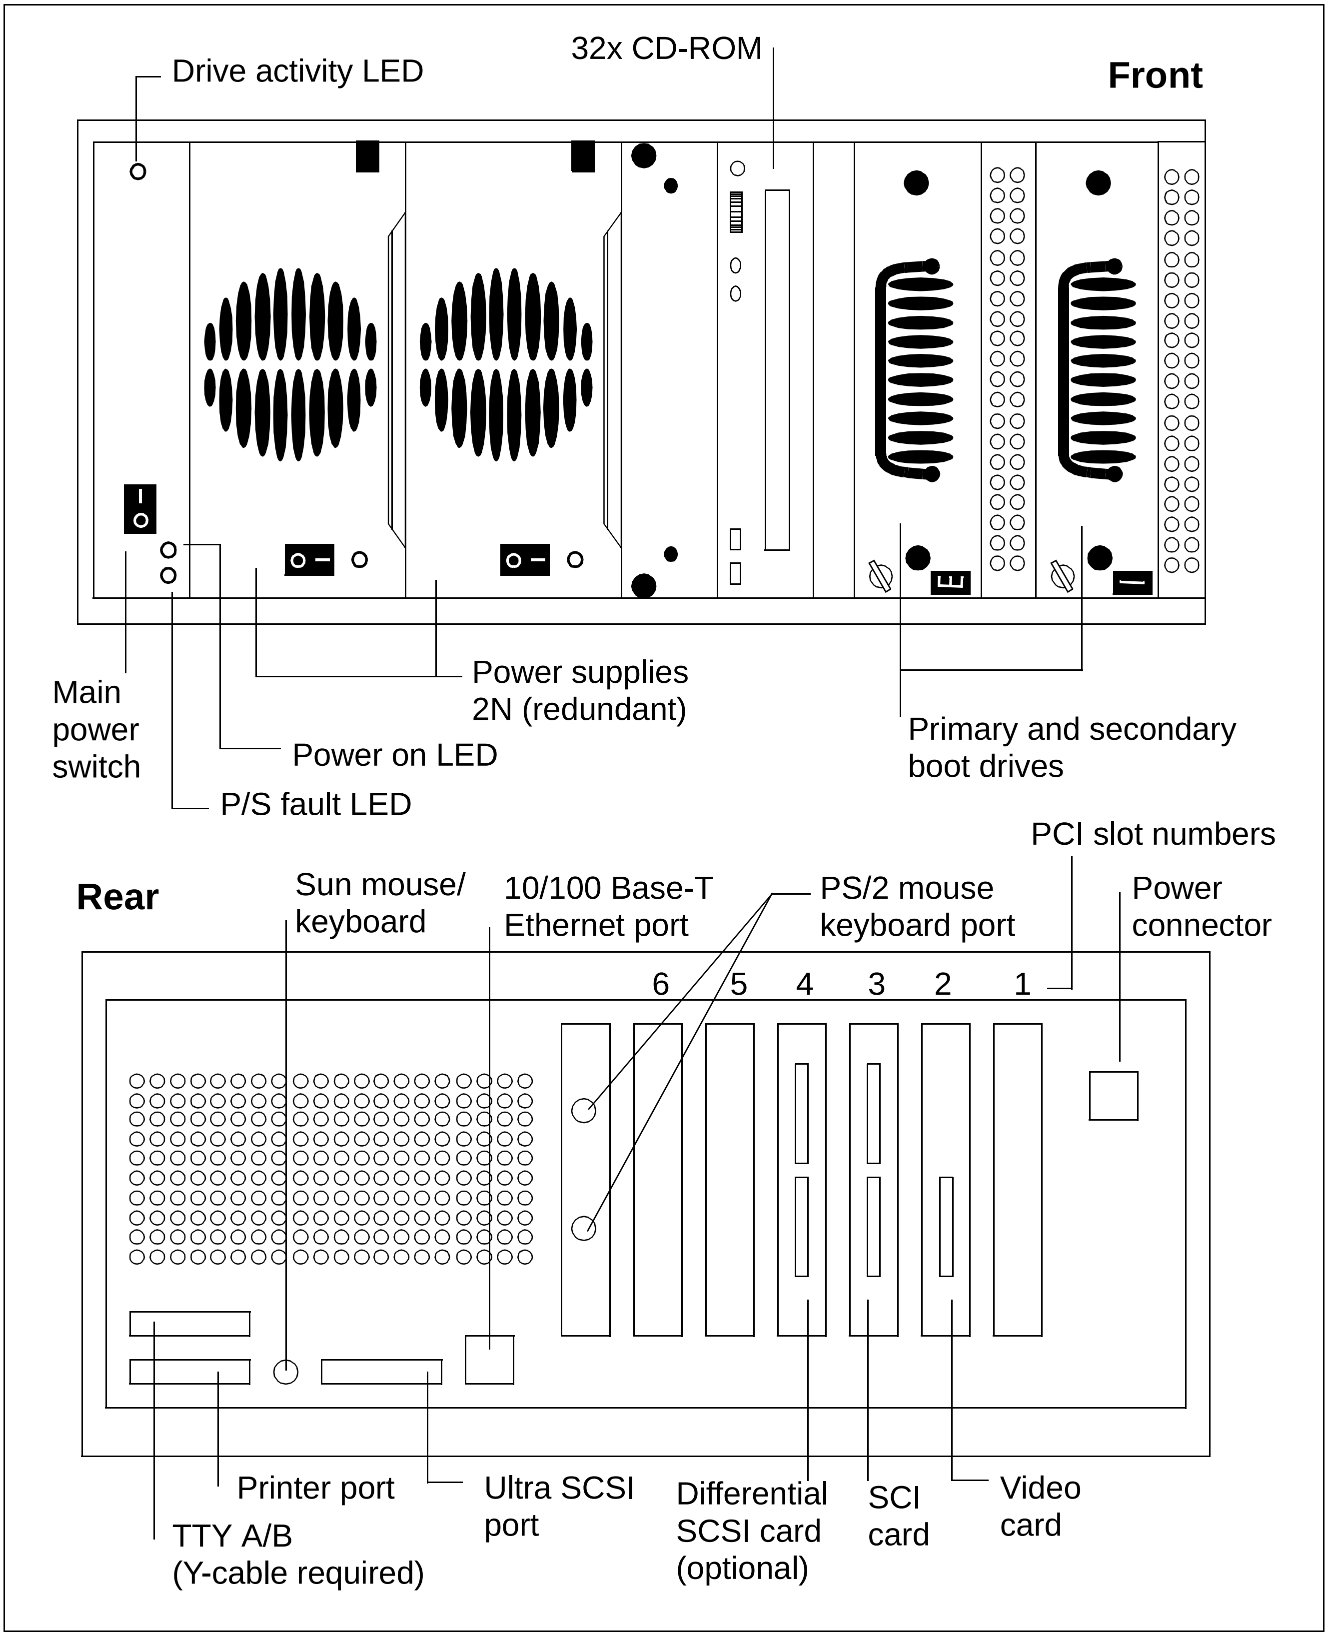

Host module (front, rear; bezels removed)

environmental monitor functions are viewed via a window on the Host system console. The following functions are supported:

- Temperature monitoring of CPU. The system provides a warning if the CPU overheats (exceeds 55^{o} C), and initiates a Host shutdown if temperature exceeds out-of-limits (58^{o} C).
- Voltage monitoring. The system provides a warning when low voltages are detected.
- The fan control function drives the CPU fan at one of four defined speeds. Speed is selected based on temperatures of four internal thermistors on the AXi motherboard. The system software controls the fan speed based on the CPU thermistor temperature readings. The fan speed increases when the CPU temperature rises and reduces when the CPU temperature drops. This results in lower overall operating noise and power consumption. The fan monitor function can also detect a non-operational CPU fan.
- The three chassis fans, located behind the drives and power supplies, operate at a fixed speed. Fan failure monitoring in the AXi motherboard flags if the fan behind the drives fails. The other two chassis fans are not directly monitored.

Hardware Configuration

The basic Host hardware configuration consists of:

- Host enclosure with motherboard. Includes 128 MB of main memory, and 256 KB of cache.
- Two 48VDC, 300-watt, hot-pluggable power supply modules in 2N (redundant) format
- One hot-pluggable, 9 GB SCSI disk drive mounted in a carrier (primary boot device)
- One 32X, SCSI CD-ROM drive
- One SVGA video card (PCI)
- One SCI card (PCI)

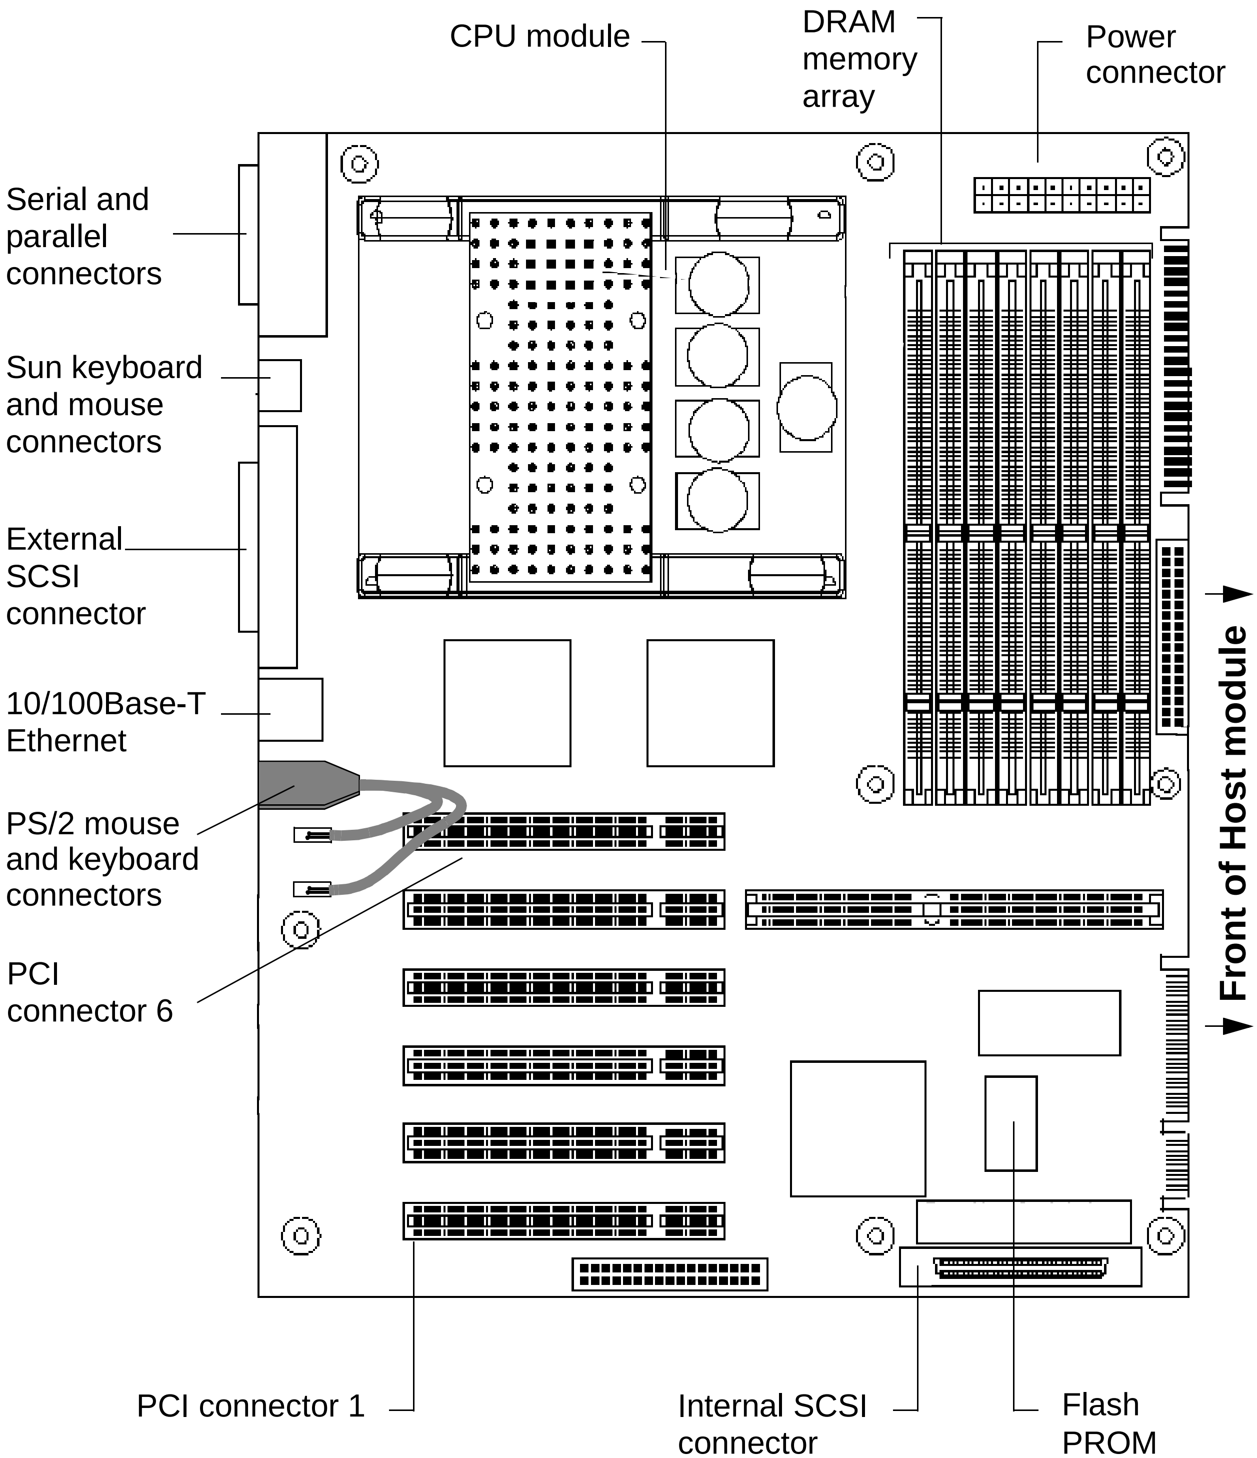

Host motherboard map (top)

- Modem for remote access

Available options include:

- An additional 9 GB SCSI disk drive assembly (backup boot device)
- Additional 64 MB DIMMs for main memory. Host memory options are 128 MB (standard), 256 MB, and 512 MB.
- Differential SCSI card to support tape backup

PCI Card Slot Assignments

The PCI card slot assignments are as follows (refer to the figures *Host module*, and *Host motherboard map*):

Slot 1: available

Slot 2: SVGA card

Slot 3: SCI card

Slot 4: Differential SCSI card (optional)

Slot 5: available

Slot 6: not used

Host Motherboard Map

The figure *Host motherboard map* shows the location of the significant motherboard components.

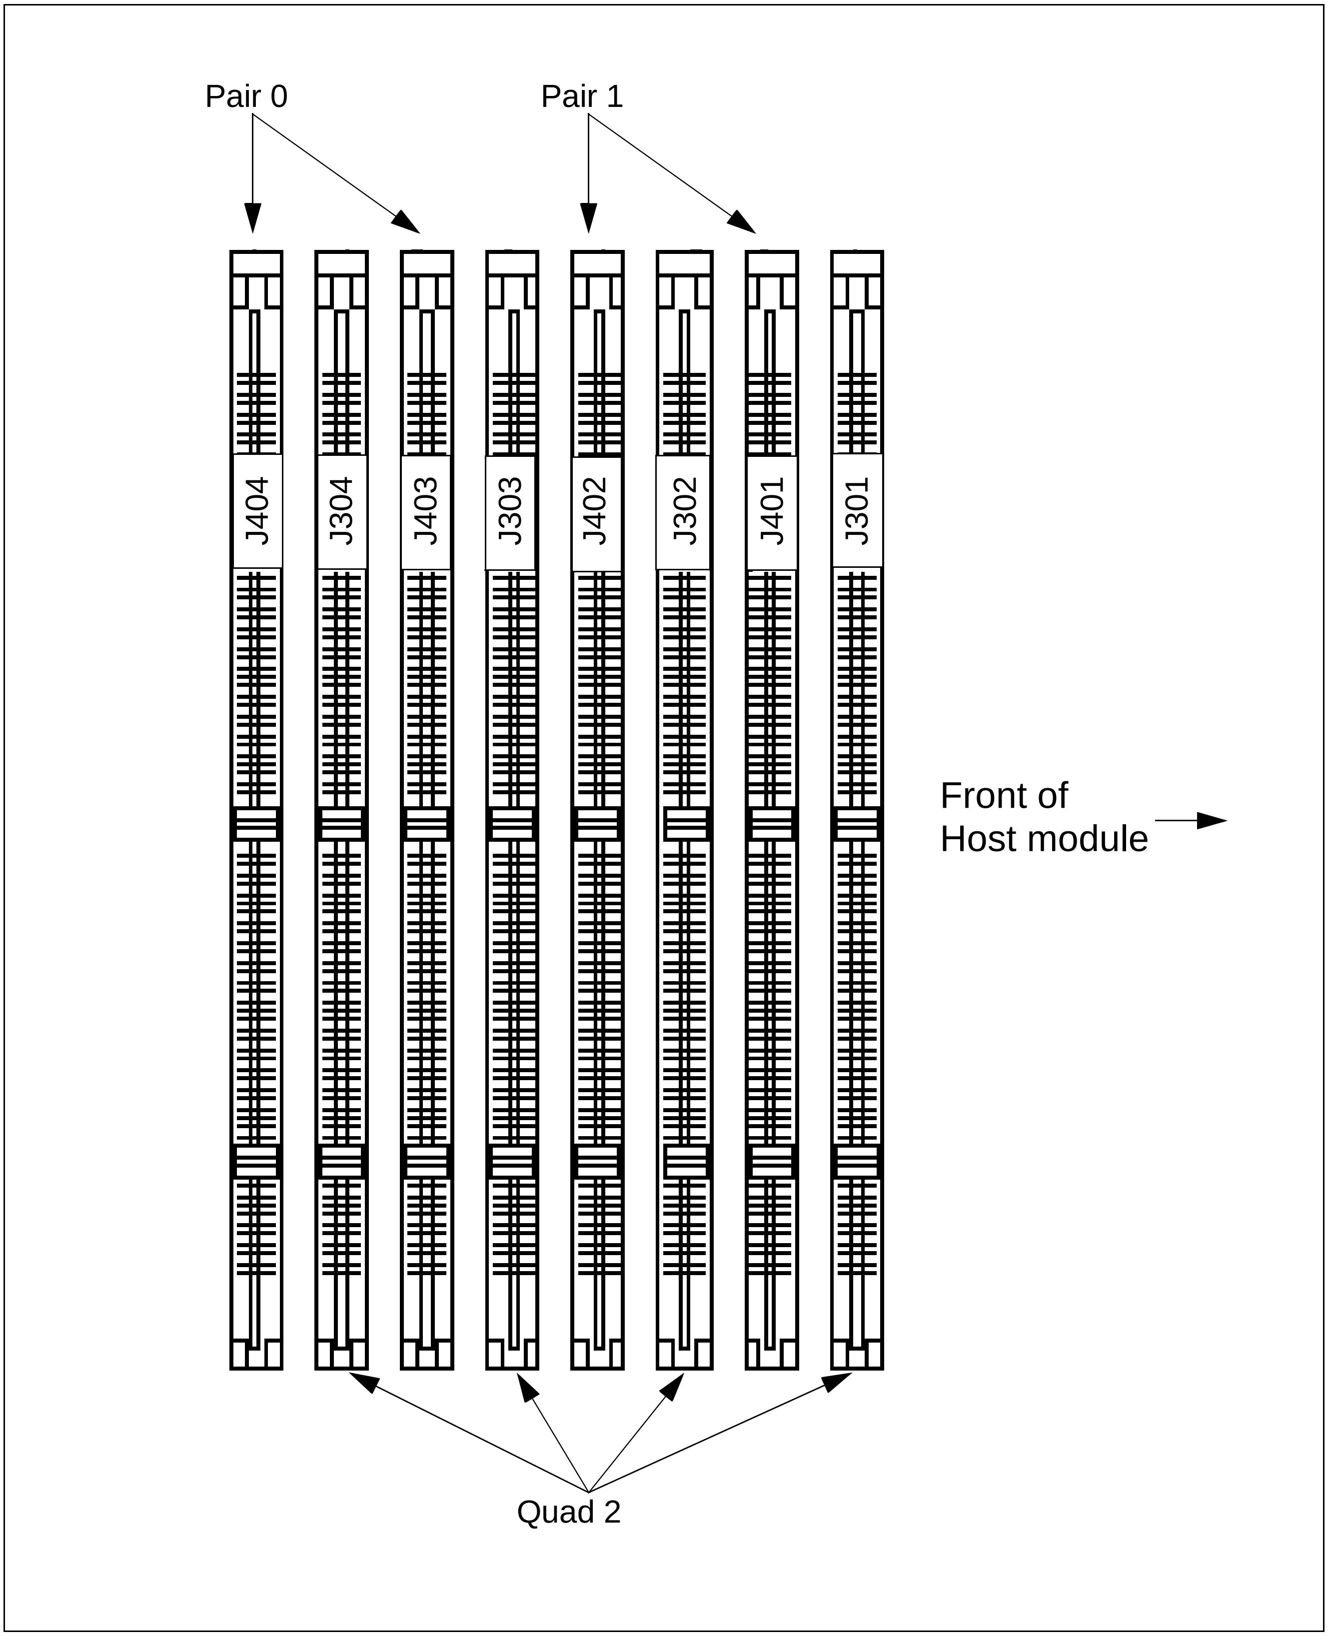

Host DIMM sockets pair/quad assignments

Host Main Memory Configuration

The Host uses pairs of 64 MB, 60 nanosecond, DRAM DIMMs to configure main memory. The standard (base) configuration is 128 MB, represented by *Pair 0* in the figure *Host DIMM sockets pair/quad assignments*. Additional supported configurations are:

- ▲ 256 MB; Pair 0 and Pair 1
- ▲ 512 MB; Pair 0, Pair 1, and Quad 2

Host FRUs		
Part number	**Part name**	**Description**
790131	Host subassembly (chassis)	Includes enclosure, motherboard, fans.
650029	Power supply module, 48VDC, 300W	Hot-pluggable power supply module installed in Host chassis.
600097	PCA, PCI-SVGA card, 4 MB	Graphics adapter that supports system console monitor.
600107	PCA, PCI-Differential SCSI card	Supports differential backup device from the Host.
600093	PCA, PCI-SCI cluster adapter	Supports SCI communication with E-Box(es).
600099	DIMM, 60nS, 8x72, EDO	Used to configure Host main memory: 128 MB, 256 MB, and 512 MB.
670018	CPU fan assembly	Hooded fan assembly that sits above the CPU module.
500933	Disk drive assembly, root	Hot-pluggable, carrier-mounted disk drive installed in Host chassis.
630052	Drive assembly, CD-ROM	CD-ROM drive for Host.
670019	Chassis fan assembly	Chassis cooling fans for Host.
640188	Modem, 56 KB, external	Modem allows M2000 remote access.
450375	Cable assembly, serial, “Y”	Y-cable (serial) interconnects Host, master E-Box, and modem.
640185	Monitor, SVGA, 17-inch	Part of system console
640160	Keyboard, 104-key, PS/2	Part of system console
640161	Mouse, 3-button, PS/2	Part of system console
450379	Cable assembly, PS/2 extension	Extension cable for PS/2 mouse and keyboard.
450380	Cable assembly, SVGA extension	Extension cable for system console monitor.

Host FRUs

The table *Host FRUs* lists the Host hardware items that are considered field-replaceable.

Note: Some of the FRUs listed in the Host FRUs table may appear in other FRU tables in this manual.

E-Box Module

Provides network, file, and storage processing for the M2000 system

Contains the following major hardware elements:

E-Box enclosure

ECHaP motherboard

DC-to-DC converter (also called Logic DCC)

EM/LCD assembly

PCI cards:

SCI card

Hardware RAID/SCSI controller

Differential SCSI converter (tape backup support)

NICs: Quad Ethernet card, Gigabit Ethernet card, FDDI card

AFX8000 card (write accelerator)

E-Box Module

Overview

The ECHaP Box (E-Box) provides network, file, and storage processing for the M2000 system. It is comprised of the following major hardware elements:

- E-Box enclosure
- Euclid Common Hardware Platform (ECHaP) motherboard
- DC-to-DC converter (also called Logic DCC)
- EM/LCD assembly (the master E-Box includes an E-LP daughter card on the EM/LCD assembly)
- PCI cards
  - Hardware RAID/SCSI controller
  - Differential SCSI converter
  - SCI card
  - Quad Ethernet card
  - Gigabit Ethernet card
  - FDDI cards
  - AFX8000 card (nonvolatile memory for the file system)

Each of the elements listed above is described in the following sections.

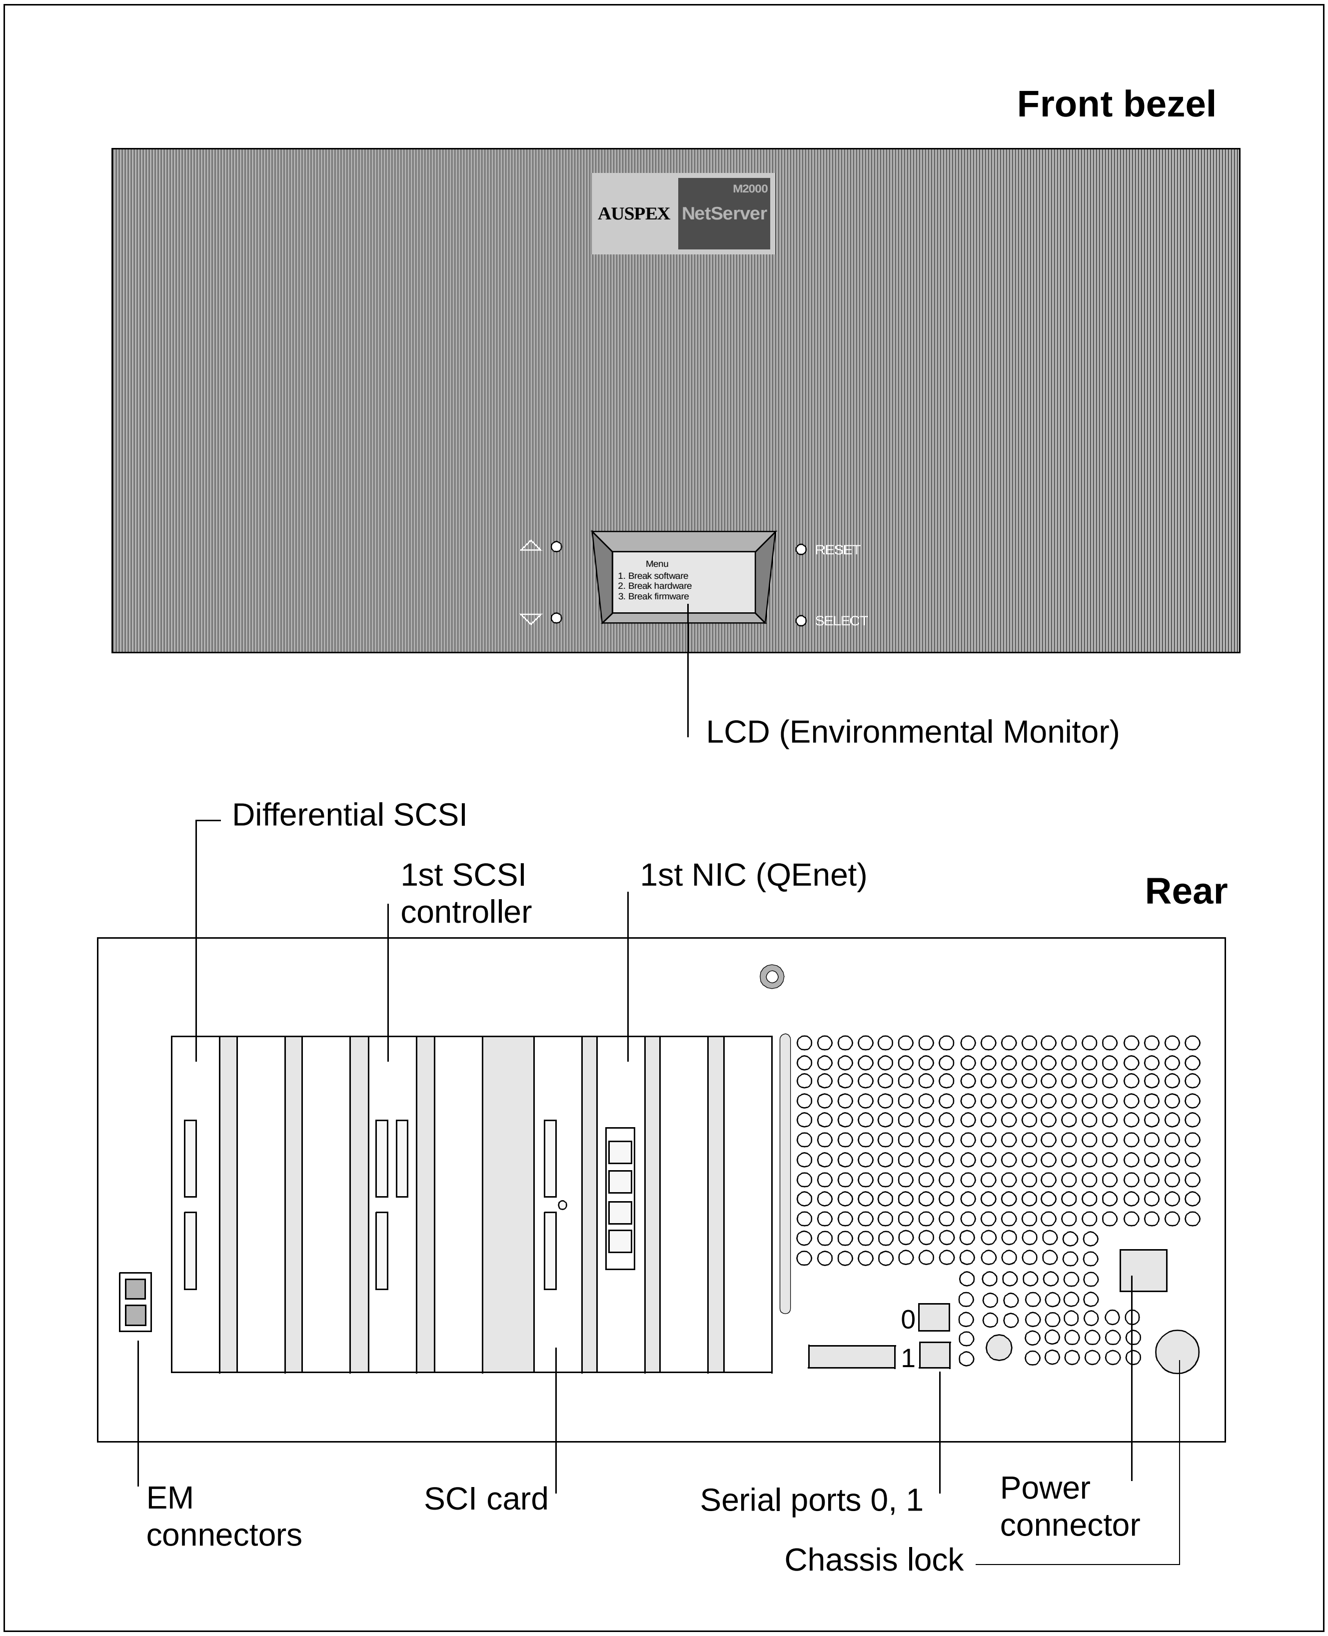

E-Box module (front, rear; rear bezel removed)

E-Box Physical Map

The figure *E-Box module* shows a front and rear view of the E-Box and highlights the key components.

ECHaP Features

Two CPUs running at 200 MHz

512 KB L2 CPU processor cache

Scalable, on-board main memory array

Bridged primary and secondary PCI buses

Flash memory for main boot PROM

Intelligent board status registers

Integrated diagnostic support

Performance monitoring for PCI buses and main memory utilization

JTAG implementation for major components

ECHaP

Summary of Features

The Euclid Common Hardware Platform (ECHaP) is the motherboard for the E-Box, also called the Network, File, and Storage Processor (NFSP). ECHaP features include:

- Two CPUs clocked at 200 MHz
  - One CPU provides NP services (network access via NICs), the other CPU provides FSP services (storage access via SCSI controller)
  - 512 KB L2 (level 2) CPU processor cache
- An on-board main memory array supporting up to 1024 MB (1 GB) of ECC-protected DRAM
- Bridged primary and secondary PCI buses:
  - Three slots on the primary bus, one of these having an auxiliary connection to the secondary bus to accept Auspex custom hardware (the AFX8000 card)
  - Four slots on the secondary bus
- Flash memory enabling field upgrades to the Main PROM firmware code
- Hardware-assisted, software-based FIFO
- Intelligent board status registers
- Integrated diagnostic support
- Performance monitoring capabilities for PCI buses and main memory utilization
- JTAG implementation for major components

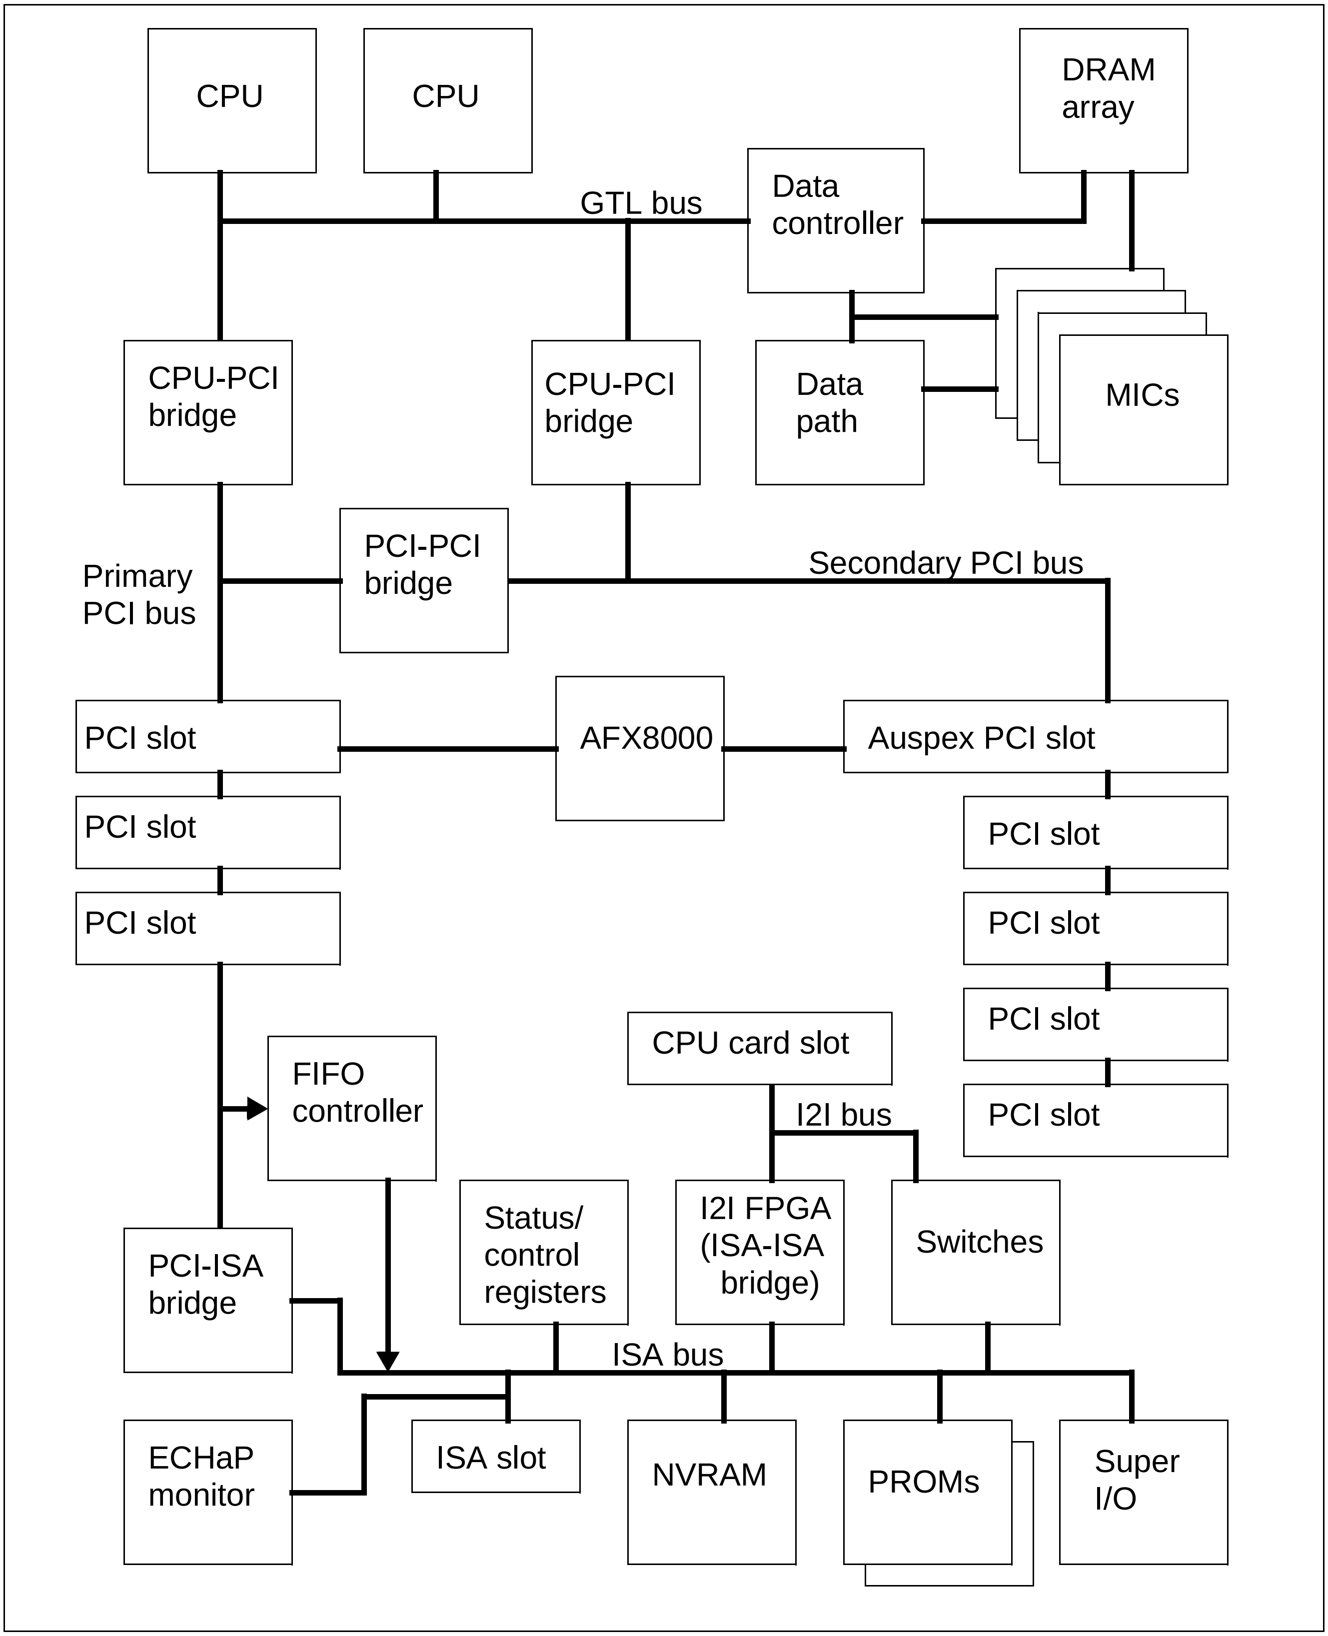

ECHaP architecture

Architectural Overview

Refer to the figure *ECHaP architecture* for a pictorial correlation of the material in this section. Data for the major functional areas of the ECHaP are provided in the section *ECHaP Functional Blocks*, later in this chapter.

The CPUs, memory controller chipset, and the main memory array are interconnected by a specialized, high performance bus using Gunning Transceiver Logic (GTL+). CPU communication with the Peripheral Component Interconnect (PCI) buses is via the CPU-PCI bridge chips. The primary and secondary PCI buses are connected through a PCI-PCI bridge. This connection allows data transfers between the buses without interfering with CPU bus traffic.

Circuit cards plug into the PCI bus expansion slots. There are three slots on the primary bus, and four slots on the secondary bus. One of the slots on the primary PCI bus has an auxiliary connection to the secondary bus called the *Auspex PCI slot*. The AFX8000, a write buffer card, plugs into this slot and the adjacent primary PCI slot, thus bridging the PCI buses.

Note: The SCI card connects only to the primary PCI bus. Inserting it into a secondary PCI slot would create 64-bit address contention with the PCI-PCI bridge.

The PCI-ISA bridge connects the following functions from the ISA bus:

- Status and control registers
- NVRAM
- Main and Auxiliary PROMs
- ID PROM
- Super I/O
- I2I FPGA[1]

1. The *I2I* (ISA-to-ISA) *FPGA* is a device that bridges two ISA buses and has a small FIFO for transferring commands and data between the single-board computer (SBC) host and the ECHaP. It is unlikely this slot will ever host an SBC, but the functionality is available. For more information, see the section *Peripheral I/O Devices and Registers*, later in this chapter.

ECHaP Bus and Interconnection Technology

GTL+ bus

Interconnects the CPUs, memory controller chipset, and main memory array

Synchronous, split-transaction, 36-bit address, 64-bit data bus

Clocked at 66 MHz; peak throughput of 532 MB/s

PCI buses

Synchronous, 32-bit address and data buses clocked at 33 Mhz; peak throughput of 132 MB/s

Primary PCI bus communicates with CPU-PCI bridge, command FIFO, and ISA bus

Primary bus has three expansion slots, supports the SCI card, a NIC card, and AFX8000 (AFX8000 bridges primary and secondary buses)

Secondary bus has four expansion slots and supports SCSI controller(s), NIC(s), and AFX8000

ECHaP Bus and Interconnection Technology

Gunning Transceiver Logic Bus (GTL+)

A GTL+ bus interconnects the CPUs, memory controller chip set, and the main memory array.

Gunning Transceiver Logic (GTL) technology is a reduced-voltage switching standard that provides high-speed backplane and point-to-point communications with low power dissipation.

GTL is a form of differential logic similar to Emitter Coupled Logic (ECL), except that it is implemented with CMOS transistors. Small amplitude signals are achieved by terminating device pins at a lower potential than the power supply.

The GTL+ bus is a synchronous, split-transaction[1], 36-bit address, 64-bit data bus. The bus runs at 66 MHz with a peak data transfer bandwidth of 532 MB per second.

PCI Buses

The ECHaP has primary and secondary PCI buses. The primary bus communicates with the CPU-PCI bridge, the command FIFO, and the ISA bus. The primary bus has three expansion slots. One of the slots supports an interface to the secondary PCI bus and is used for Auspex specific hardware (AFX8000). The primary bus also supports an SCI card and a Network Interface Card (NIC). The secondary PCI bus has four expansion slots, and supports SCSI controllers and NICs.

1. Split-transaction protocol is one method of increasing bus bandwidth. For a split-transaction bus, a requesting device transmits the request and address. After the memory acknowledges the request, the bus is released and free for other transactions. When the requested data is available, the memory signals the requesting device. After receiving the data, the requesting device signals that it has the data and the memory system releases the bus. Using this protocol, multiple bus transactions can be simultaneously processed.

ECHaP Bus and Interconnection Technology (continued)

ISA bus

Provides connectivity for PROMs, control and status registers, NVRAM, Super I/O, and other devices

Scalable Coherent Interface

SCI is an interconnection standard that uses unidirectional, point-to-point signaling instead of traditional bus architecture

M2000 SCI implementation uses PCI-SCI bridge cards plugged into the Host PCI bus and the primary PCI bus of the ECHaP. The cards are interconnected via cables thus establishing intermodule communication.

The PCI buses are synchronous, multiplexed 32-bit address and data buses running at a maximum of 33 MHz (half the speed of the GTL+ bus). The peak data transfer bandwidth is 132 MB per second at the maximum clock rate.

ISA Bus

The Industry Standard Architecture (ISA) bus provides connectivity for PROMs, control and status registers, NVRAM, Super I/O, an other devices. The key components connected to the ISA bus are shown in the figure *ECHaP architecture*, and are described later in this chapter.

Scalable Coherent Interface (SCI)

ANSI/IEEE Std. 1596-1992 defines Scalable Coherent Interface.

SCI is an distributed interconnection standard that provides the expected computer bus services but avoids traditional bus limitations by using point-to-point links and a packet protocol. SCI is designed for high-performance, highly parallel multiprocessors, but scales down for use with uniprocessors.

Prior to SCI, computer buses were essentially made by wiring together corresponding connector pins on a backplane or motherboard. Systems bused in this fashion cannot carry signals at speeds high enough to accommodate the new generation of computers.

SCI uses unidirectional point-to-point signaling, from one transmitter to one receiver, to simulate a bus without employing traditional bus architecture. This simplifies electrical problems so transmission speed can be greatly increased, but requires changes to basic bus protocol.

SCI Implementation

For the M2000, SCI is implemented with PCI-SCI cards plugged into the Host PCI bus and the E-Box primary PCI bus. The cards are interconnected via cables thus establishing intermodule communication.

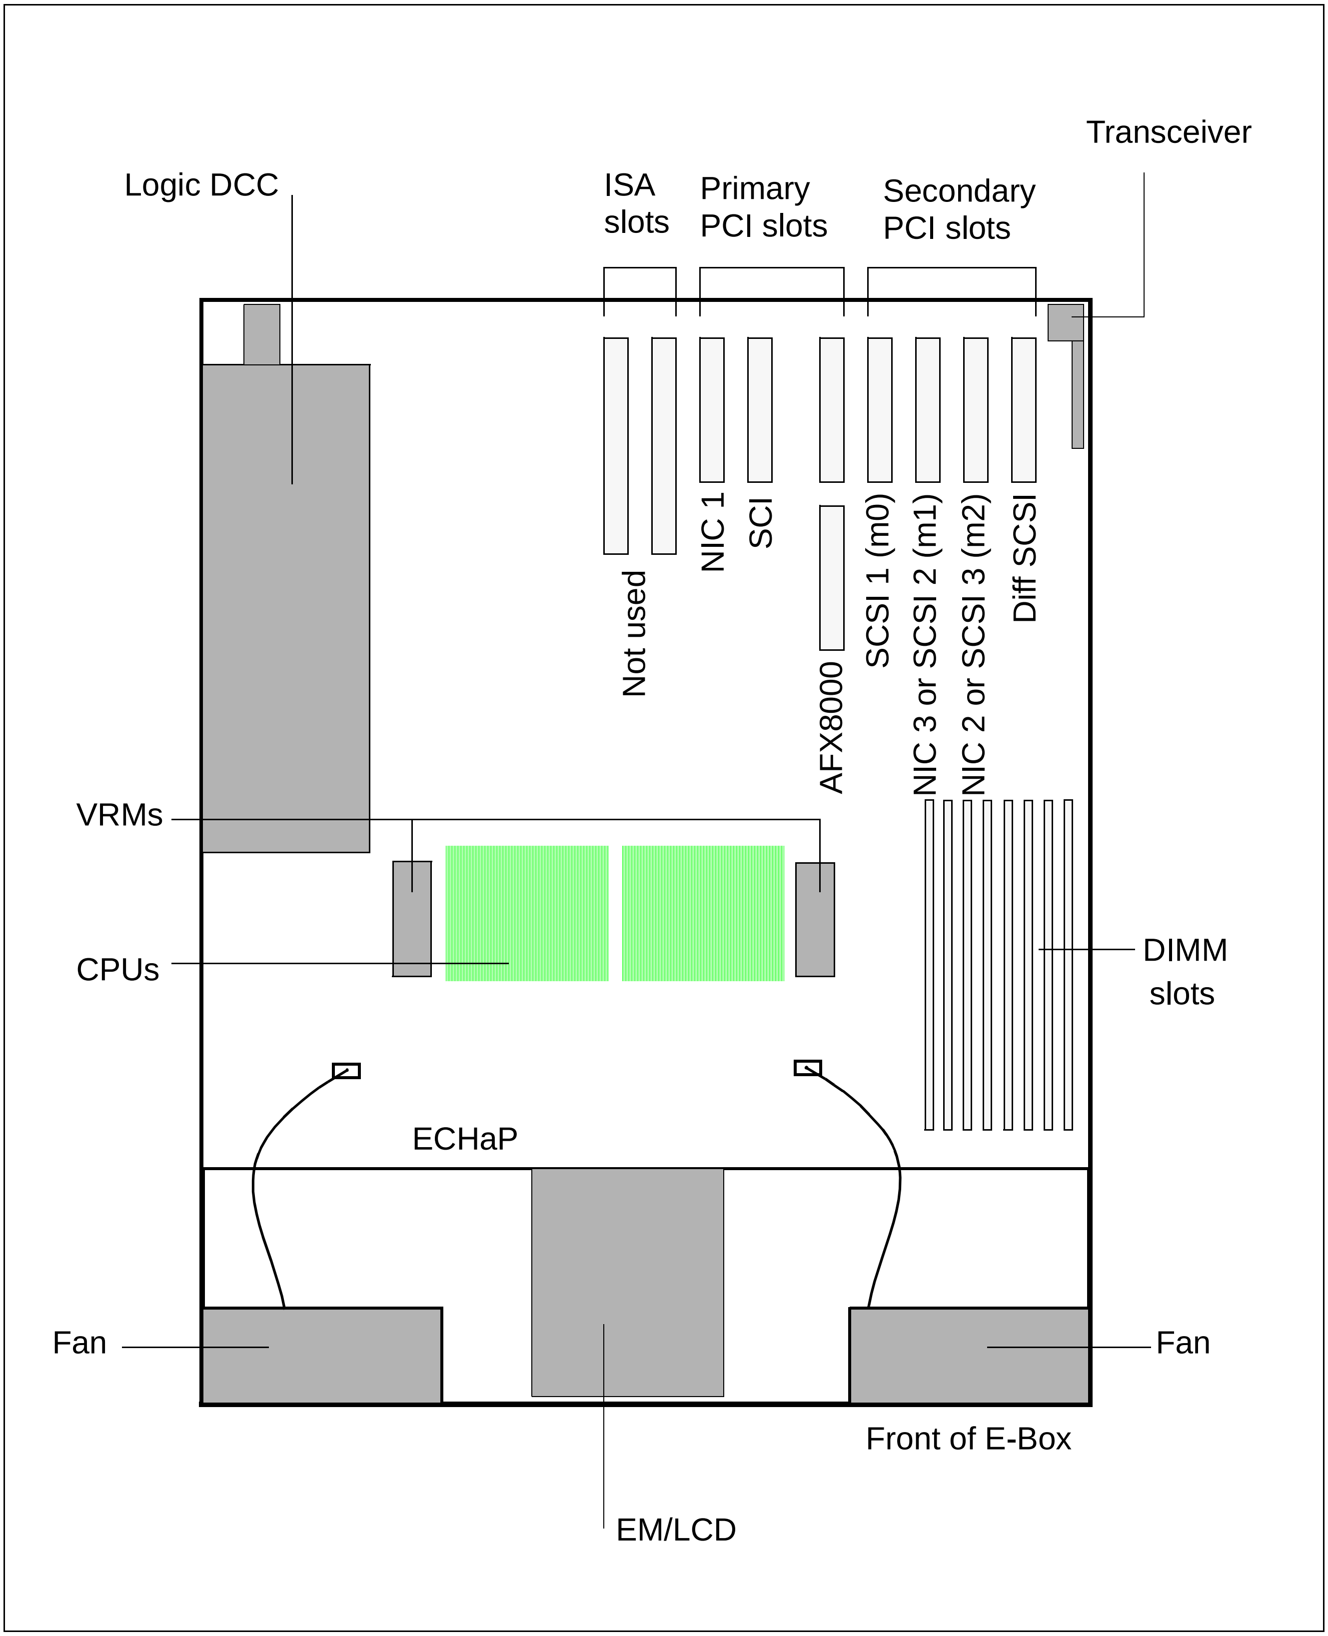

ECHaP motherboard map

Upon system initialization, SCI detects all operating nodes and all the paths between adjacent nodes. SCI's adaptive routing allows it to automatically route around failed nodes without reinitializing. It may be helpful to think of the SCI-interconnected server as a memory-mapped ring of transmitting and receiving nodes.

ECHaP Motherboard Map

The figure *ECHaP motherboard map* shows the location of the key ECHaP components.

ECHaP Functional Blocks

CPUs and memory controller chipset

Main memory array

Primary and secondary PCI buses

System command FIFO

Peripheral I/O devices and registers (ISA bus)

ECHaP Functional Blocks

The sections that follow provide data on functional sections of the ECHaP. To correlate the information, refer to the figure *ECHaP architecture*, earlier in this chapter.

The functional blocks of the ECHaP are:

- CPU and memory controller chipset
- Main memory array
- Primary and secondary PCI buses
- Command FIFO controller
- Peripheral I/O devices and registers

CPU and Memory Controller Chipset

The CPU and memory controller chipset consists of the following components:

- Two CPUs
- Two CPU-PCI bridges
- A Data Controller (DC)
- A Data Path (DP)
- Four Memory Interface Components (MICs)

The CPUs are clocked at 200 MHz, the interconnecting GTL+ bus is clocked at 66 MHz. The CPU-PCI bridge divides the GTL+ bus clock by two for the PCI buses, resulting in their 33 MHz rate.

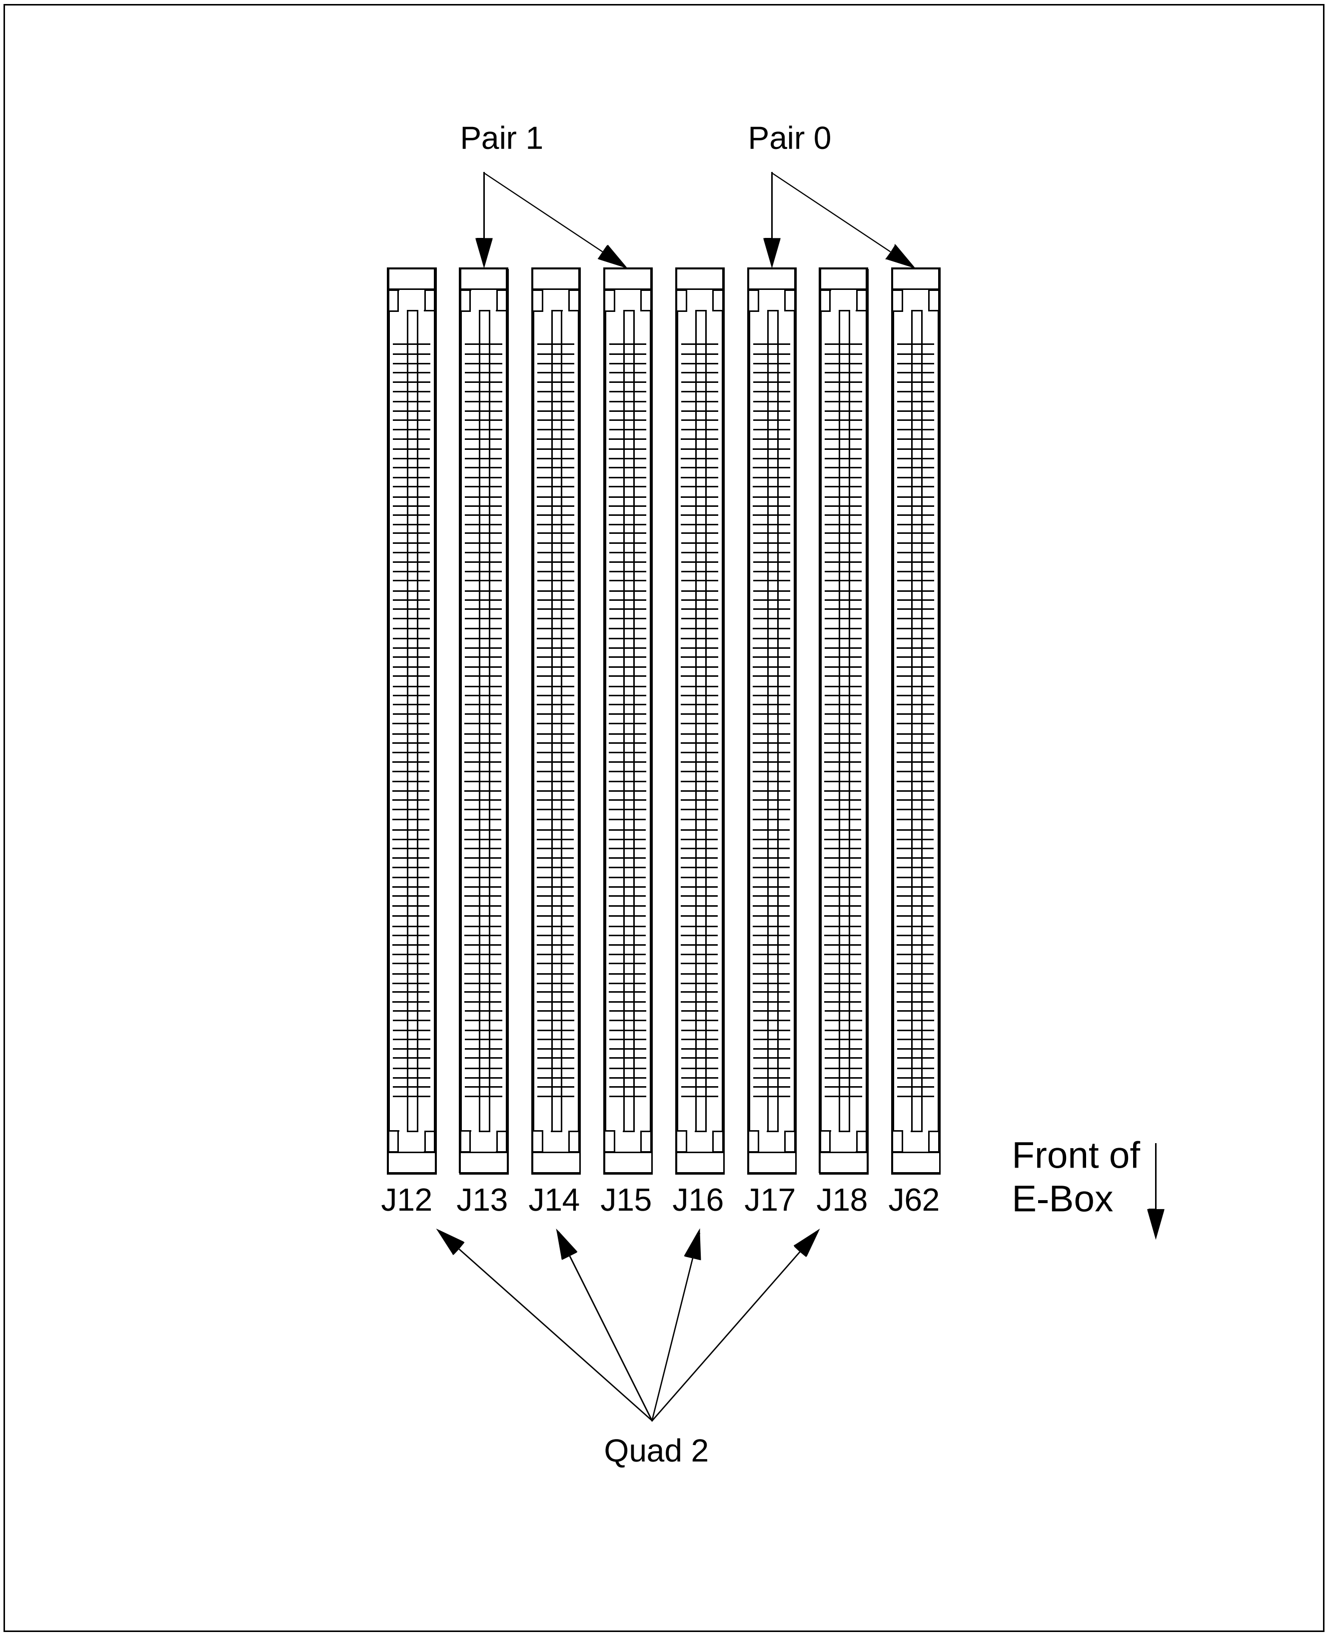

ECHaP DIMM sockets pair/quad assignments

Main Memory Configuration

The main memory configuration is illustrated in the figure *ECHaP DIMM sockets pair/quad assignments*. The main memory has the following features:

- ▲ Supports up to eight Dual In-Line Memory Modules (DIMMs) organized as four sets of two.
- ▲ The array uses 3.3 VDC, 60 nanosecond, double-sided DIMMs with a capacity of 128 MB.
- ▲ The standard main memory configuration for ECHaP motherboards is 256 MB
- ▲ Optional, expanded memory configurations are:
  - – 512 MB (pair 0, pair 1)
  - – 1024 MB (pair 0, pair 1, quad 2)

PCI Buses

The PCI buses are described earlier in this chapter in the section *ECHaP Bus and Interconnection Technology: PCI buses*.

Command FIFO Controller

The System Command FIFO on the ECHaP board is memory based. It is designed to accept writes from the primary PCI bus, then write data out to a circular buffer in the main memory (DRAM). The FIFO provides positive acknowledgment of message reception via the SCI transaction protocol.

Peripheral I/O Devices and Registers

The ISA bus implements connectivity for a number of functions; refer to the figure *ECHaP architecture* for pictorial detail. The key ISA functions are as follows:

Student notes

- The **NVRAM** maintains a set of default parameters for the ECHaP between power cycles. The NVRAM has a memory capacity of 128 MB.
- The **Main PROM** is a flash PROM that contains the most current firmware code. The **Auxiliary PROM** is programmed with the firmware that is current at the time of system shipment. If a problem occurs with the Main PROM, the system can boot from the Auxiliary PROM. The Auxiliary PROM is not field-upgradeable.
- The **Super I/O** provides the most commonly used I/O peripherals in PC design:
  - Keyboard controller
  - Real-time clock
  - Two Universal Asynchronous Receiver-Transmitters (UARTs)
  - Two serial ports
  - Parallel port
- The **ID PROM** is a 32-byte device that contains the ECHaP board serial number.
- The **I2I Field Programmable Gate Array (FPGA)** is used as a bridge between two ISA buses (*I2I* stands for ISA-to-ISA). This provides communication between the ECHaP and a special ISA slot. This switchable slot serves as a standard ISA slot or host to an ISA-based single board computer (not used in the M2000).
- The **Status and Control Registers** hold mostly Auspex-specific data. The registers are as follows:
  - PCI interrupt pending status - low
  - PCI interrupt pending status - high
  - SCI status register
  - Last FIFO entry register
  - FIFO control and status register
  - FIFO available space register

Student notes

- PCI-to-PCI bridge flush register
- FIFO size register
- Interrupt pending mask register
- LM78 registers (the LM78 provides ECHaP environmental monitoring)
- Clear board error register
- Environmental monitor registers
- Performance monitor registers
- DIMM presence detect registers
- Board control register
- Board and assembly revision register
- ISA FPGA revision register
- Soft reset and soft NMI register

Configuration Components

Dual In-Line Memory Module (DIMM)

SCI card

SCSI/RAID controller card

Differential SCSI card

Quad 10/100 Base-T card

Gigabit Ethernet card

FDDI card

AFX8000

Configuration Components

This section contains a brief description of the components used to configure the HFSP and the NP.

Dual In-Line Memory Module (DIMM)

The DIMMs are described earlier in this chapter in the section *Main Memory Configuration.*

SCI Card

PCI-SCI cluster adapter cards plug into the Host module PCI bus and into the primary PCI bus in the E-Box. These cards support the SCI-PCI interface and provide the backbone for processor module communication. The SCI card has the following features and specifications:

- Addressing support for 64,000 nodes
- Dual links supporting copper, Shielded Twisted Pair (STP) cable connections
- Peak transfer bandwidth of 1.6 gigabits per second per link (3.2 gigabits per second, duplex)
- 32-bit PCI interface

Hardware RAID/SCSI Controller Card

The M2000 uses a RAID/SCSI controller with the following features:

- Three single-ended, Wide UltraSCSI channels that support RAID levels 0, 1, 5, 6 (Virtual Partition) and 7 (JBOD: just a bunch of disks)
- High performance RISC processor running at 232 MHz

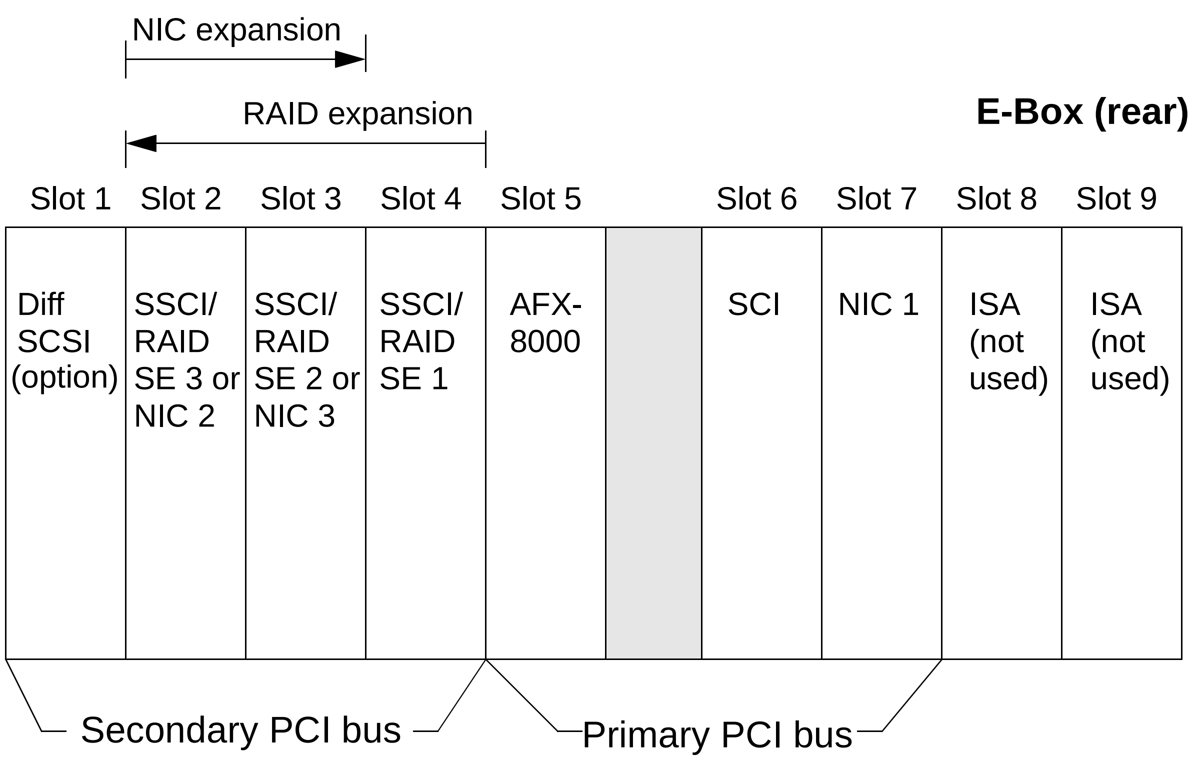

PCI card slot assignments

- 64 MB, battery-backed SDRAM cache memory module
  - Backup battery is a 3.6V nickel-cadmium
- The external SCSI connectors are female, Very High Density Cable Interconnect (VHDCI) connectors.

Differential SCSI Card

The Differential Wide UltraSCSI controller card has two independent SCSI channels bridged to the PCI bus. The maximum data transfer rate is 80 megabytes per second (40 megabytes per second, per channel). On the front panel are two 68-pin, high density connectors. This card supports differential tape libraries such as the ACL 4/52 or the ACL 7/100.

Quad 10/100 Base-T Card

The Quad 10/100 Base-T card has four Ethernet ports that autosense between 10 and 100 megabits per second hub connections. This card supports RJ-45 (UTP) connectors. It has a PCI-to-PCI bridge that extends the ECHaP PCI bus and reduces its load.

Gigabit Ethernet Card

The M2000 E-Box supports a Gigabit Ethernet card with these features:

- Separate read and write DMA paths support full duplex Gigabit Ethernet connectivity. Each DMA channel can sustain over 1 gigabit per second throughput across the PCI bus.
- Completely interoperable with existing equipment using standard Ethernet minimum to maximum frame size (64 to 1518 bytes), and standard frame format
- Optional *Jumbo Frame* support improves bulk data transfer performance and minimizes packet processing overhead
- Interrupt coalescing. The Gigabit Ethernet card adapts host interrupt frequency based on traffic load.

Student notes

FDDI Card

The PCI-FDDI Adapter allows the M2000 access to a FDDI network and has the following features:

- ▲ Single (SAS) or Dual (DAS) Attachment
- ▲ Initial support for fiber media connectors and eventual support for UTP connectors
- ▲ Simultaneous PCI bus and FDDI data transfers
- ▲ High performance Motorola FDDI chipset
- ▲ 32-bit DMA supports 128-byte bursts

Auspex AFX8000 Card

The AFX8000 serves as a non-volatile staging buffer for writes to any single disk or array. The card plugs into the special Auspex hardware slot on the ECHaP and is one of the cards that configures the ECHaP as a File Storage Processor (FSP).

The AFX8000 features are as follows:

- ▲ Write buffer comprised of 128 MB of battery-backed, non-volatile DRAM
  - – The battery is a 1.5 volt, rechargeable alkaline AA cell. Minimum data retention is 72 hours, typical is seven days. Ideal battery replacement interval is two years.
- ▲ Dual PCI ports; write buffer accessible from both ports

E-Box FRUs		
Part number	**Part Name**	**Description/Comments**
500817	ECHaP with E-Box enclosure	Enclosure and ECHaP motherboard used to build an E-Box.
650025	Power supply, Logic DCC	Logic DCC that converts 48VDC into operating voltages for the E-Box.
600077	128 MB DRAM DIMM, 16x72, 60ns, 3.3V	Plugs into the main memory array on the ECHaP motherboard.
600063	PCA, VRM8, DC-to-DC, 5VDC	Voltage regulator module that converts 5VDC to 3.3 VDC for CPUs.
100176	PCA, AFX8000	Write buffer card that bridges the PCI buses in the E-Box.
380082	Battery, 1.5V, alkaline, AA	Battery for AFX8000 memory.
600080	PCI Quad 10/100BT adapter	Network Interface Card (NIC) with four 10/100 BT ports.
600086	PCI dual differential Ultra SCSI adapter	Provides differential SCSI output for a backup tape device.
600093	PCI-SCI adapter	Plugs into E-Box primary PCI bus; supports SCI with Host.
600098	PCI RAID/SCSI controller	Three-channel hardware RAID/SCSI controller card.
600101	PCI Gigabit Ethernet adapter	Provides Gigabit Ethernet connectivity for M2000.
600104	PCI-FDDI adapter, SAS, fiber	Single Attachment FDDI for fiber media.
600105	PCI-FDDI adapter, DAS, fiber	Dual Attachment FDDI for fiber media.
670013	Fan assembly	Two of these fans are in each E-Box.
450326	Cable assy., SCI, 1 M	1 meter SCI cable (intranodal)
450369	Cable assy., SCI, 5 M	5 meter SCI cable (internodal)

8 Storage Subsystem

Introduction

Auspex's current-generation server products use the High Density Disk Array (HDDA) as the foundation for their storage subsystems. The use of Solaris 2.6 and a unique SCSI/RAID controller make the M2000 HDDA different with regards to SCSI cabling, drawer identification, and device numbering.

Objectives

After completing this chapter you should be able to do the following:

- ▲ State the major components of the HDDA.
- ▲ Describe the mechanical features of the HDDA.
- ▲ Explain the M2000 disk numbering convention.
- ▲ Describe M2000 SCSI cabling.
- ▲ State the precautions associated with handling the HDDA drives.

HDDA Overview

Each HDDA supports up to four drive drawers, each drawer houses up to seven 18 GB disk drives

Each M2000 node supports up to three HDDAs

Fully configured system storage capacity of 4.5 TB (three nodes, all with three HDDAs)

Each drive drawer uses a Disk DCC to convert the 48-volt input to 5 VDC and 12 VDC

PIC chip on each drive drawer motherboard gathers environmental data and communicates with the HDDA LON board

Disk drives are SCA-2 SCSI

Each HDDA has three 48-volt cooling fans

Each drive drawers has a jumper that sets SCSI termination and high- and low-order SCSI IDs

HDDA Overview

The High Density Disk Array (HDDA) was developed to provide as much storage capacity as possible in the available space. The HDDA chassis supports up to four drive drawer assemblies, each housing up to seven disk drives. Each M2000 node supports up to three HDDAs. A fully configured system (three nodes) stores 4.5 terabytes of data (252 disk drives, each storing 18 gigabytes).

Additional HDDA features are as follows:

- Each HDDA drive drawer supports a DC-to-DC converter (Disk DCC) that converts the 48-volt input to the 5 VDC and 12 VDC needed by the local disks.
- Each drive drawer motherboard has a Programmable Integrated Controller (PIC) that gathers environmental data, generates control signals, and communicates with the HDDA LON board plugged into the rear of the HDDA chassis. The *Environmental Monitor Subsystem* chapter, later in this manual, describes the HDDA LON board.
- The disk drives are SCA-2[1] SCSI, have a 18 GB storage capacity, and are equipped with rails and an ejector for simple replacement.
- Each HDDA has three cooling fans that operate on 48 VDC.
- Each drive drawer has a jumper that sets SCSI termination and high- and low-order SCSI IDs.

1. SCA is an abbreviation for Single Connector Attachment, a type of disk drive connector that includes connection pins for power and data wires. The current version of SCA, called SCA-2, uses 80 pins and is frequently used for high-end SCSI devices. SCA supports hot- and warm-plug for disk drives.

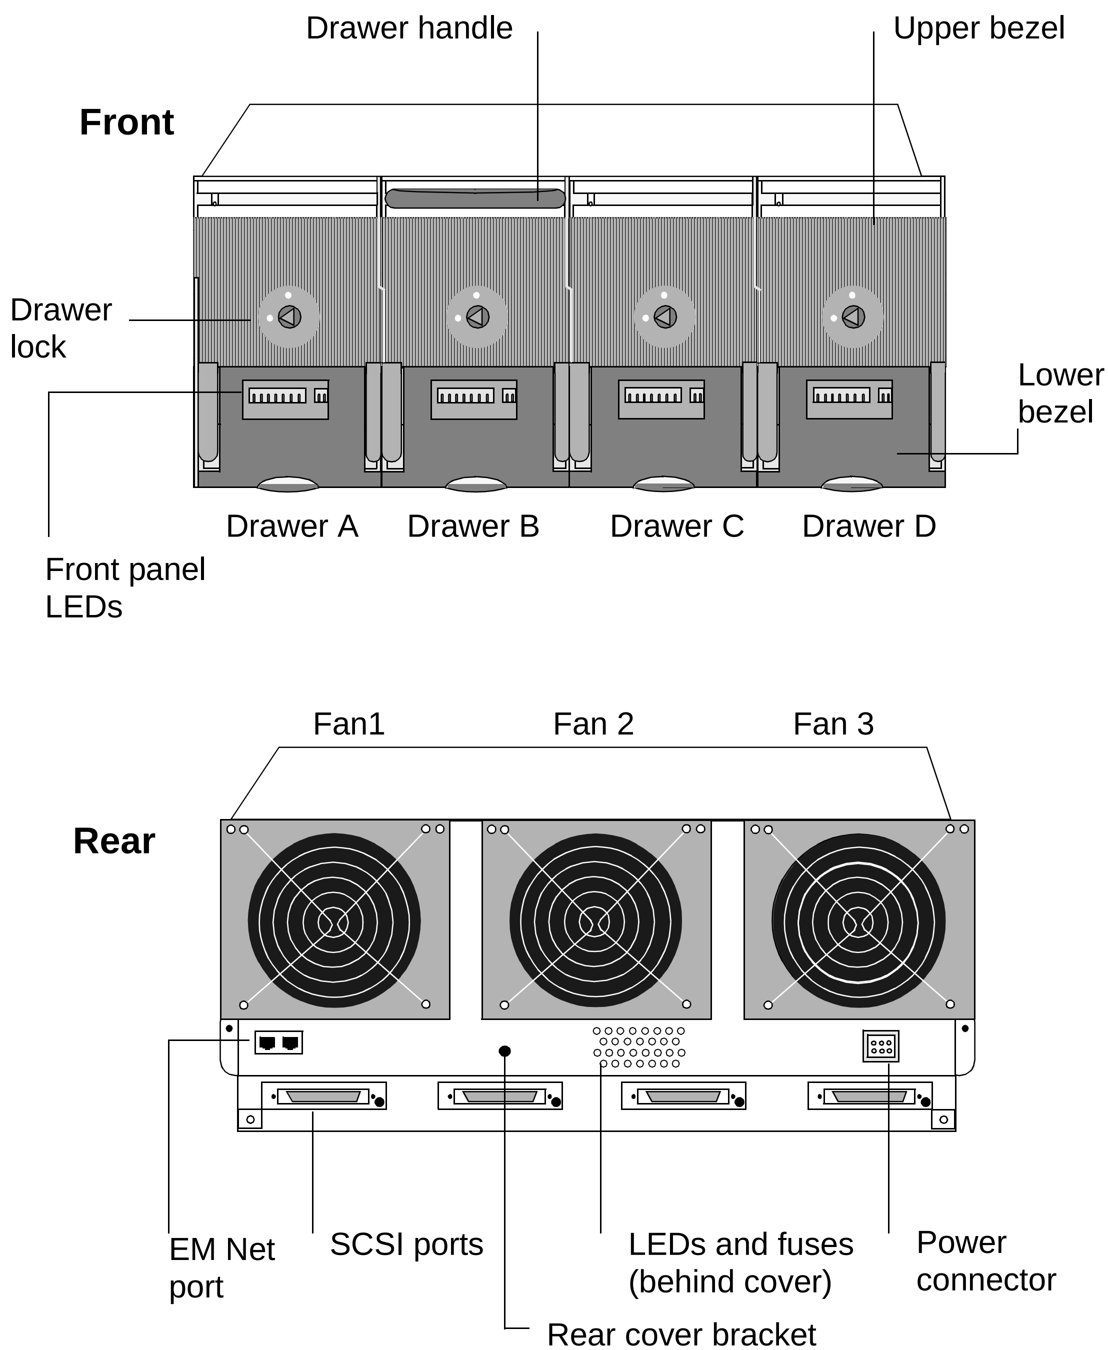

Component designation and numbering

Component Designation and Numbering

The figure *Component designation and numbering* shows the location of commonly referenced HDDA components.

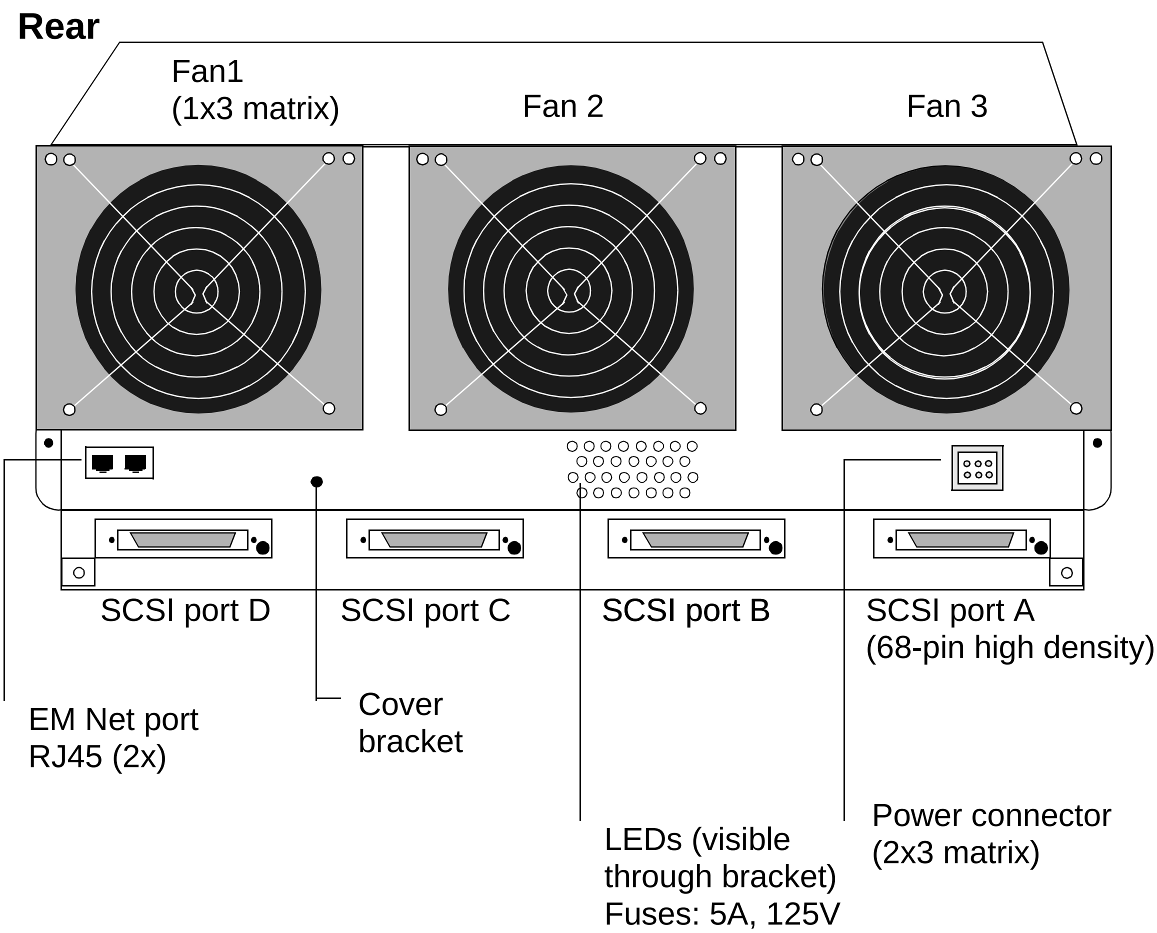

HDDA connections

HDDA Connections

The HDDA SCSI, power, and communication ports are on the rear of the chassis. Refer to the figure *HDDA connections* to correlate the following descriptions:

- ▲ Each HDDA supports an *HDDA LON Board* that plugs into a port at the rear of the chassis. This board is part of the environmental monitoring subsystem and provides two RJ45 connectors for the EM Net cables.
- ▲ Each HDDA has four 68-pin (Wide) SCSI connectors, one for each drive drawer.
- ▲ The HDDAs receive their 48-volt input via a cable connected to a Power Shelf.
- ▲ Three 48-volt, hot-pluggable cooling fans connect to the HDDA backplane, and are mounted to the back of the HDDA chassis. The fans draw cooling air through the front panel of the drive drawers then exhaust the heated air.
- ▲ Protecting the drive drawers are 5 amp, 125 volt fuses. There is one fuse per drawer and each has an adjacent LED. Any lit LED indicates that its corresponding fuse is blown. To replace one of these fuses, follow the procedure in the *Auspex NetServer M2000 Field Service Guide*. The LEDs are visible through a perforated section of the rear cover bracket.

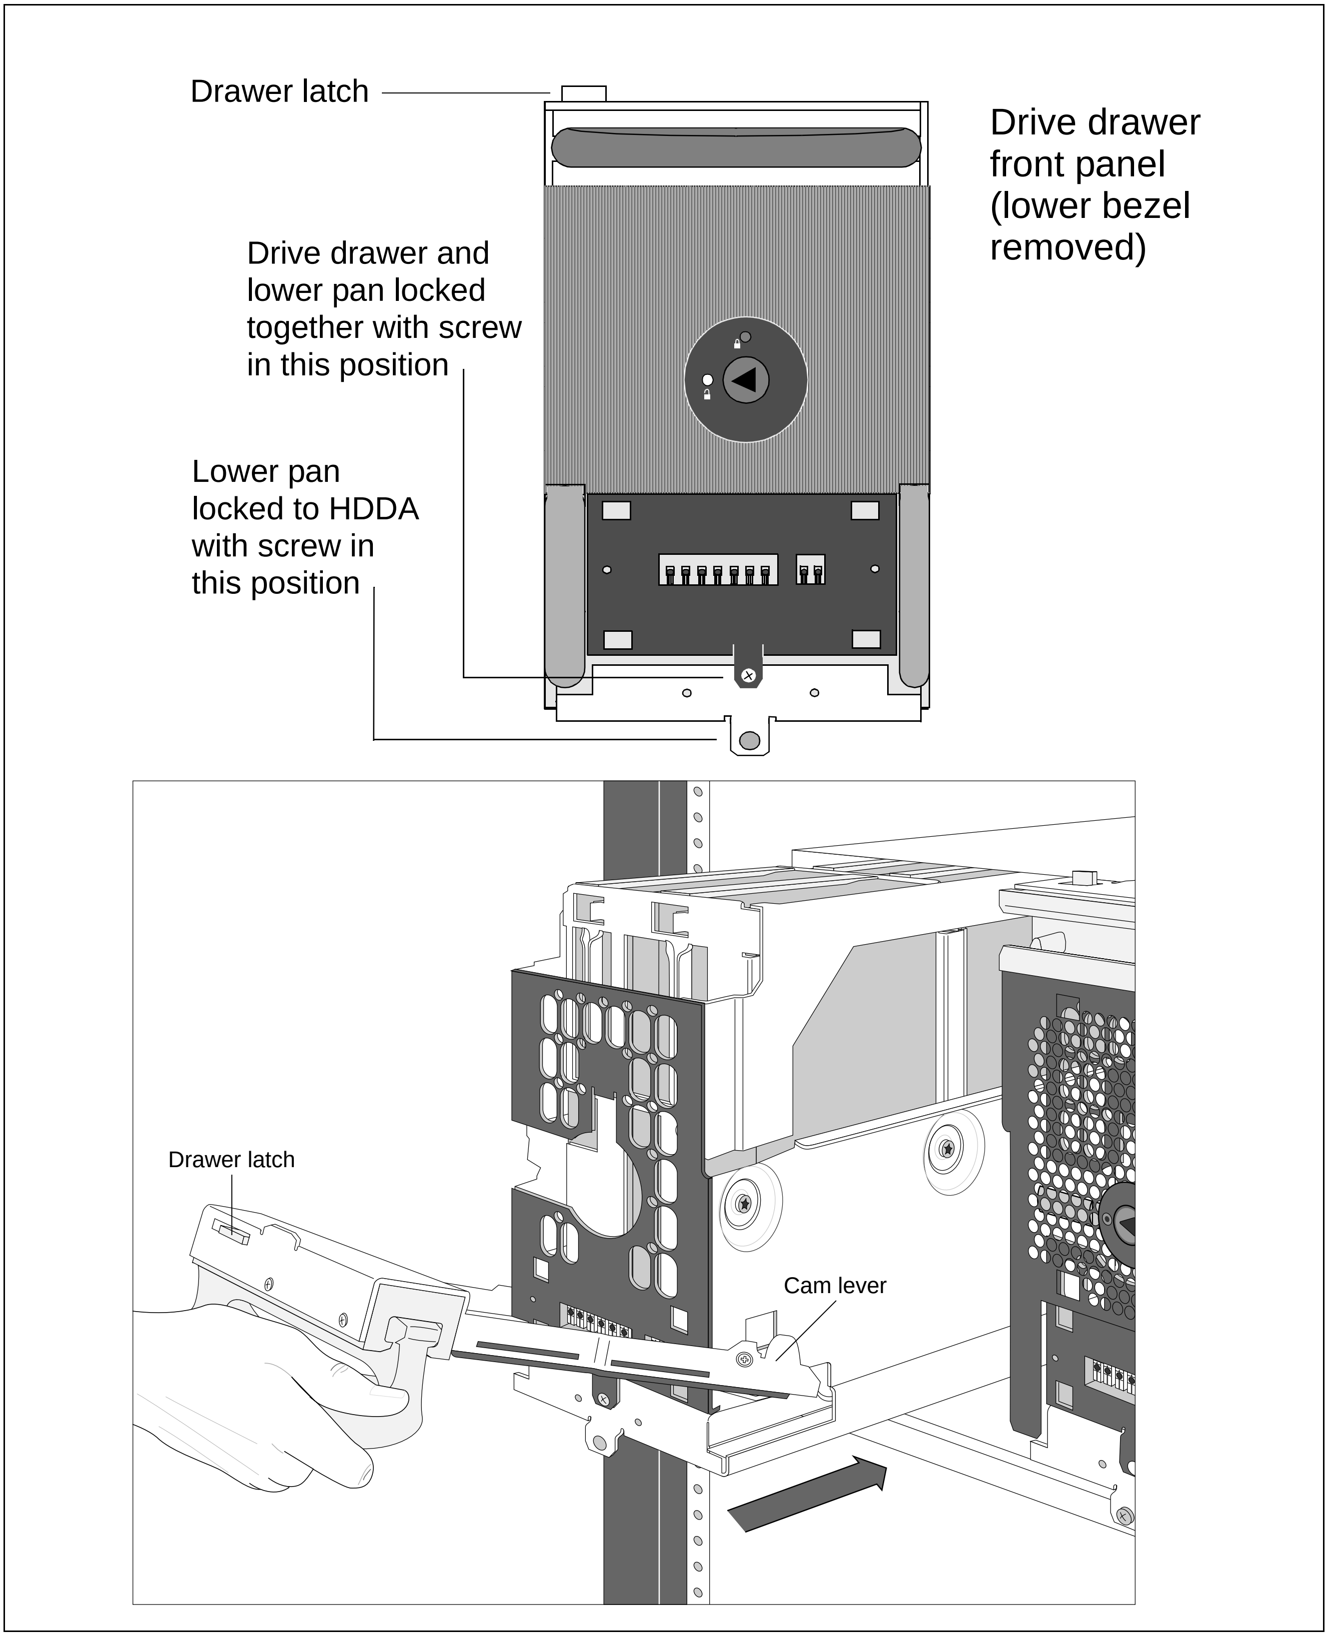

HDDA drive drawer mechanics

HDDA Drive Drawer Mechanics

The HDDA drive drawer assembly is comprised of two sections:

- ▲ The *drive drawer* that houses the drives and DCC power supply.
- ▲ The *lower pan* that houses the umbilical that routes 48 volts to the drive drawer motherboard.

When the drawer assembly is installed, the connector on the rear of the drive drawer engages a SCSI connector on the HDDA backplane. The connector on the rear of the lower pan engages a power connector on the HDDA backplane.

Refer to the figure *HDDA drive drawer mechanics*. When shipped, and prior to installation, the drive drawer and lower pan are locked together with a screw (upper position). After installation, the screw is removed to unlock the drawer assembly sections and moved to the lower position to secure the lower pan to the HDDA chassis.

With the lower pan secured to the HDDA chassis:

- ▲ The drive drawer can slide into and out of the chassis without interrupting power. This allows access to the drives for warm-swap replacement.
- ▲ Disengaging a drive drawer from the HDDA breaks the SCSI connection for that drawer. A drive drawer is disengaged by unlocking it, selecting it with the handle, and pulling downward on the handle.

During normal server operation, the lower pan will be secured to the HDDA chassis. If you are removing a drive drawer assembly from the HDDA, you must move the locking screw to the upper position to lock the two drawer assembly sections.

Caution: Check the position of the locking screw prior to opening a drive drawer. When the drive drawer and the lower pan are locked together and the drawer assembly is disengaged from the HDDA chassis, SCSI **and** power connections are broken.

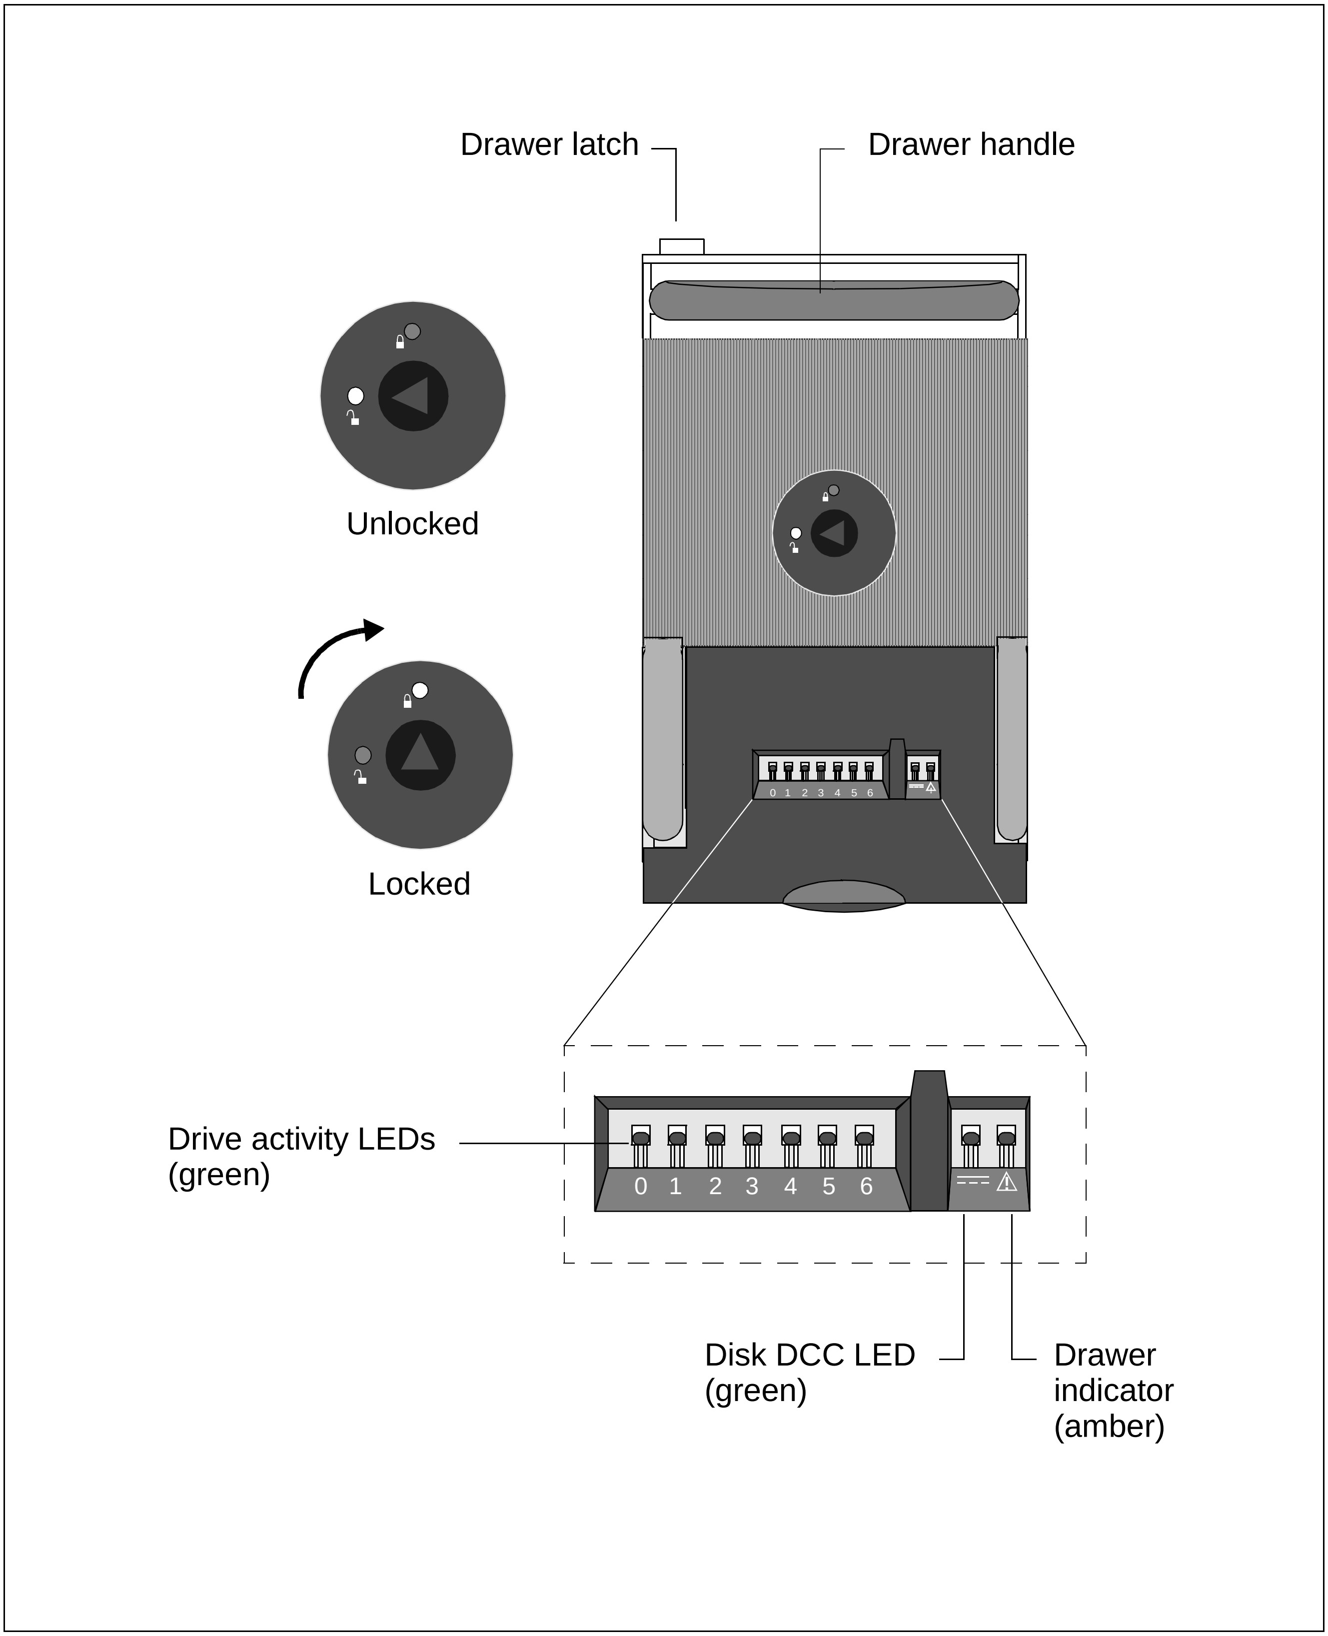

HDDA drive drawer lock and lower bezel detail

Drive Drawer Component Access

The following are important points concerning drive drawer component access and component replacement:

▲ Each drawer is key-lockable. The drawer lock (viewed from the front) is locked when the dot on the lock ring is in the 12 o'clock position, and unlocked when it is in the 9 o'clock position. Unless you plan to remove one of the drawers, it is prudent to keep them locked. Refer to the figure *HDDA drive drawer lock and lower bezel detail*.

▲ Each HDDA has one drawer handle that is shared among the drive drawers. A slotted rail mounted to each drive drawer front panel allows the handle to slide across the front of the drawers. When the handle is centered on a front panel, it causes the latch on the selected drawer to retract. An unlocked drawer can then be opened by pulling down on the handle. This action operates a cam lever that disengages the drive drawer from the HDDA backplane. The SCSI connection is interrupted when the drive drawer is disengaged from the backplane. It is therefore important to keep drive drawers locked to prevent someone from inadvertently pulling on or leaning on the handle. Auspex also recommends storing the drive handle off-centered with respect to the front panels so that drawer latches are not retracted.

Caution: Keep HDDA drive drawers locked whenever possible. Store the drive handle off-centered, with respect to the front panels, so that drawer latches are not retracted.

▲ The disk drives intended for use in the M2000 can not be installed in carriers. They are equipped with guide rails and a bail-locking ejector. The rails on the drive engage matching slots in the drive drawer. The plastic ejector serves as a cam lever to seat and unseat the drive. The lack of a carrier leaves the drive especially vulnerable to physical damage; handle the drives with extreme care. for more information, refer to the section *Drive Handling Precautions*, later in this chapter.

▲ The DCC power supply has the same mechanical features as the disk drives (ejector and guide rails).

Drive Handling Precautions

Wear a grounded wrist strap at all times

Use one hand to hold the ejector, the other to hold the head/disk assembly

Do not touch the circuit card

Take care when removing a drive from a drawer under power; the other disks are spinning

Allow 30-second spin-down time before removing a drive

Set drives down gently and only on antistatic mats

Do not allow a drive to tip over

Auspex recommends that you wear an antistatic lab coat when handling HDDA disk drives

Drive Handling Precautions

Practice the following when handling drives used in the HDDA:

- ▲ Wear a properly grounded wrist strap at all times while handling disk drives.
- ▲ When possible, use two hands: one hand to hold the ejector, the other to hold the head/disk assembly. Do not touch the circuit card.

  Caution: The drives used in the HDDA have exposed circuit cards. Take great care not to touch the card when handling the drives.

- ▲ Use extra care when removing a drive from a drawer under power; the other disk drives are still spinning.
- ▲ Before removing a disk drive from an HDDA drive drawer, allow about 30 seconds for spin-down.
- ▲ Set drives down gently and only on antistatic mats. Do no allow a drive to tip over: a fall of 1/4-inch or greater is sufficient to damage a drive.
- ▲ Auspex recommends that you wear an antistatic lab coat when handling HDDA drives.

Spare disk drives intended for use with the HDDA will have an ejector and guide rails.

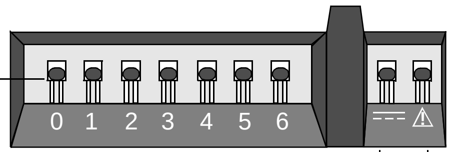

HDDA front panel LEDs

LED	Normal operating status	Failure indication	Action required upon failure
Drive activity	Off when corresponding drive is inactive; blinking when corresponding drive is active	All LEDs off could indicate a drawer fuse or SCSI failure	Check and replace drawer fuse per *Auspex NetServer M2000 Field Service Guide*
Disk DCC	Lit	Off	Check and replace Disk DCC per *Auspex NetServer M2000 Field Service Guide*
Drawer indicator	Off	N/A	Planned: LED flashes in response to commands to add or replace a drive in the affected drawer

Interpreting Front Panel LEDs

The table *HDDA front panel LEDs* provides information on interpreting front panel LED status and recommends a course of action.

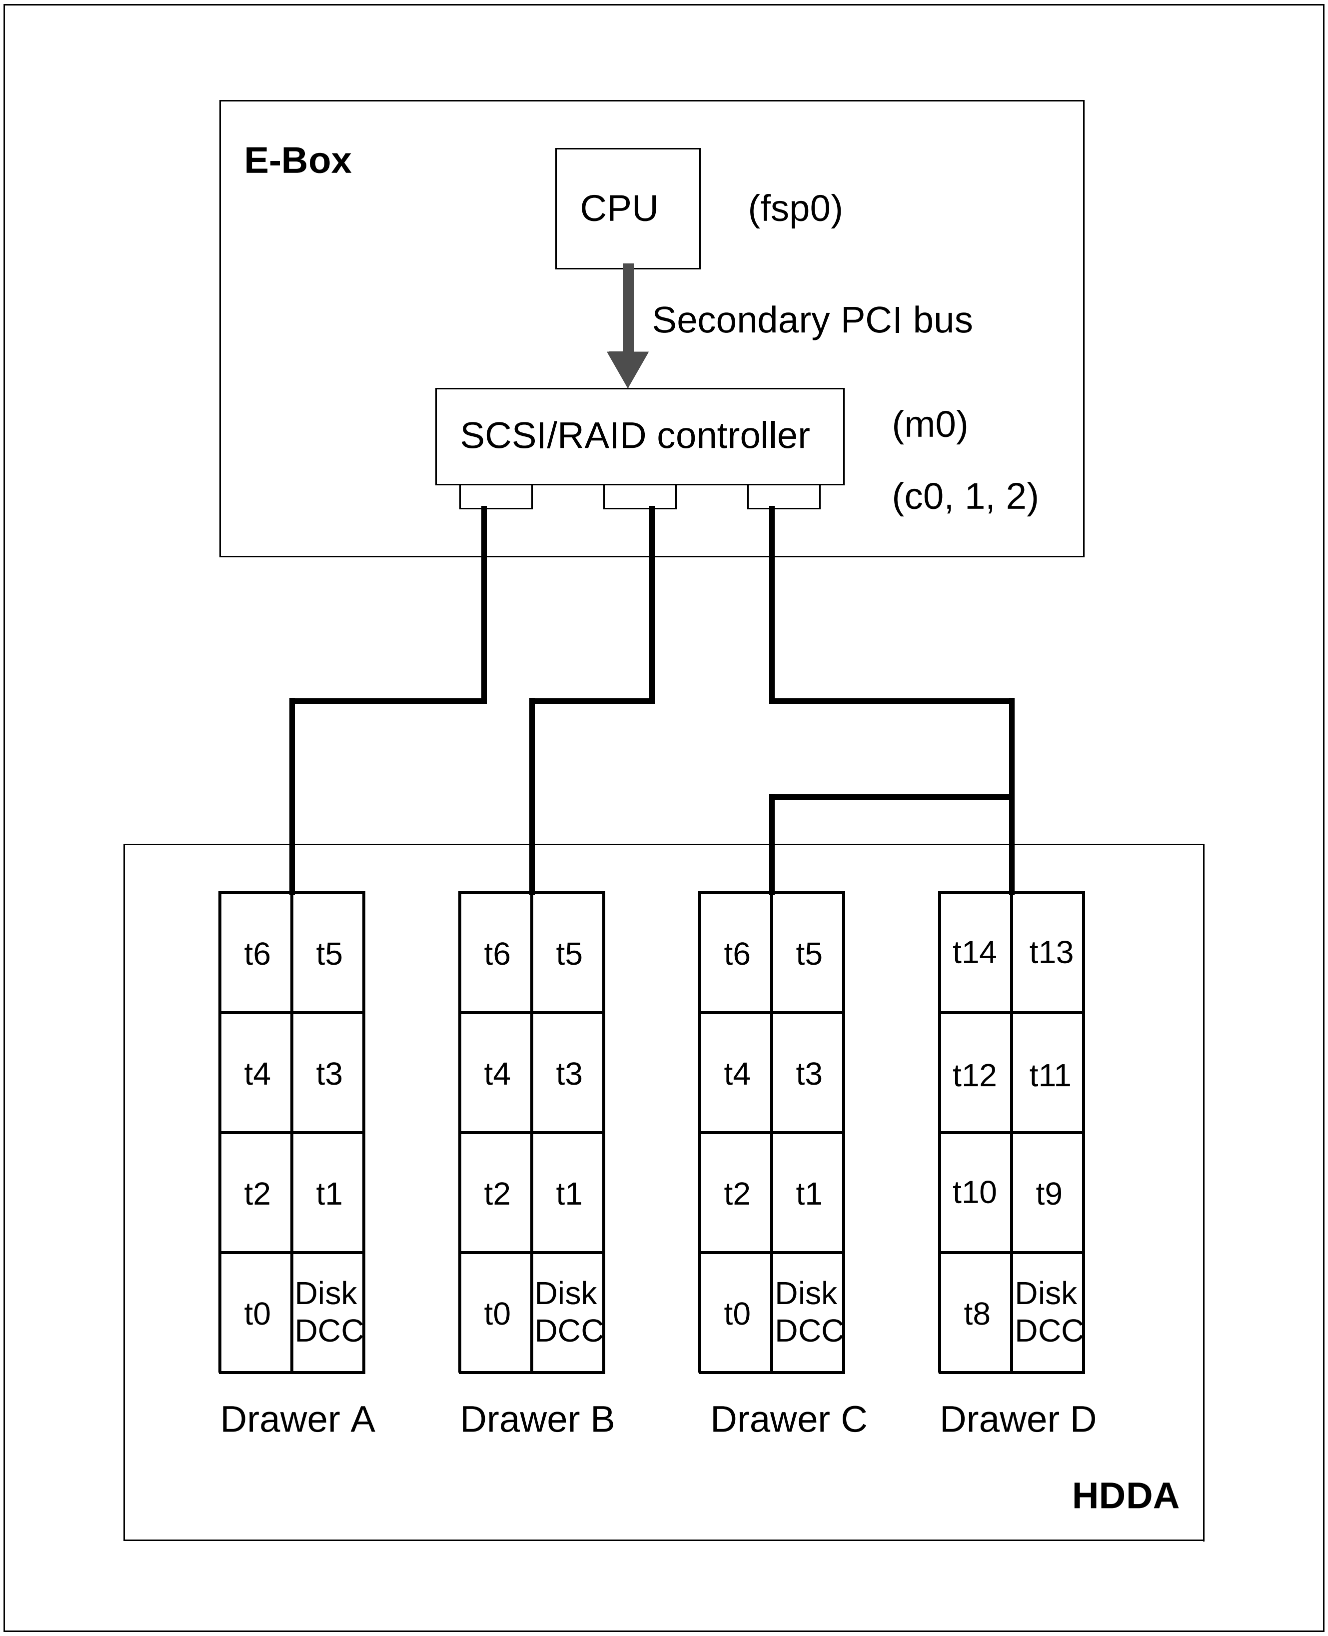

Base node, single-HDDA SCSI distribution

SCSI Cabling and Mapping

Overview

The M2000 storage subsystem is composed of these key elements:

- ▲ HDDAs that support up to four drawers, each with a SCSI port
- ▲ SCSI/RAID controller cards that have three single-ended, Wide UltraSCSI ports (channels)
- ▲ Standard, point-to-point SCSI cables that connect a SCSI port directly to an HDDA drawer
- ▲ Daisy-chain cables that connect two HDDA drawers (up to 14 disks) to a single port on the SCSI controller

Given the elements listed above, and that an M2000 supports up to nine HDDAs (three per node), the SCSI map permutations are numerous and complex. This section does not attempt to map all the possibilities, but present the student with foundational knowledge that will be useful in the field. To this end, this section of the manual presents several different pictorial and written descriptions of SCSI cabling, mapping, device numbering.

SCSI ID and Termination

The figure *Base node, single-HDDA SCSI distribution* shows SCSI control and distribution at its simplest level. The drive drawers at the bottom of the figure provide a physical map of disk location within the drawer, as viewed from the top of the drawer.

In the earlier chapter, *Processor Subsystem*, it was explained that one of the ECHaP CPUs provides FSP services, and the other CPU provides NP services. The FSP CPU in the base node is designated *fsp0* and provides storage access via the SCSI controller(s) plugged into the ECHaP's secondary PCI bus.

Each SCSI controller has three Wide UltraSCSI channels that connect to HDDA drawers with either point-to-point or daisy-chain cables. The

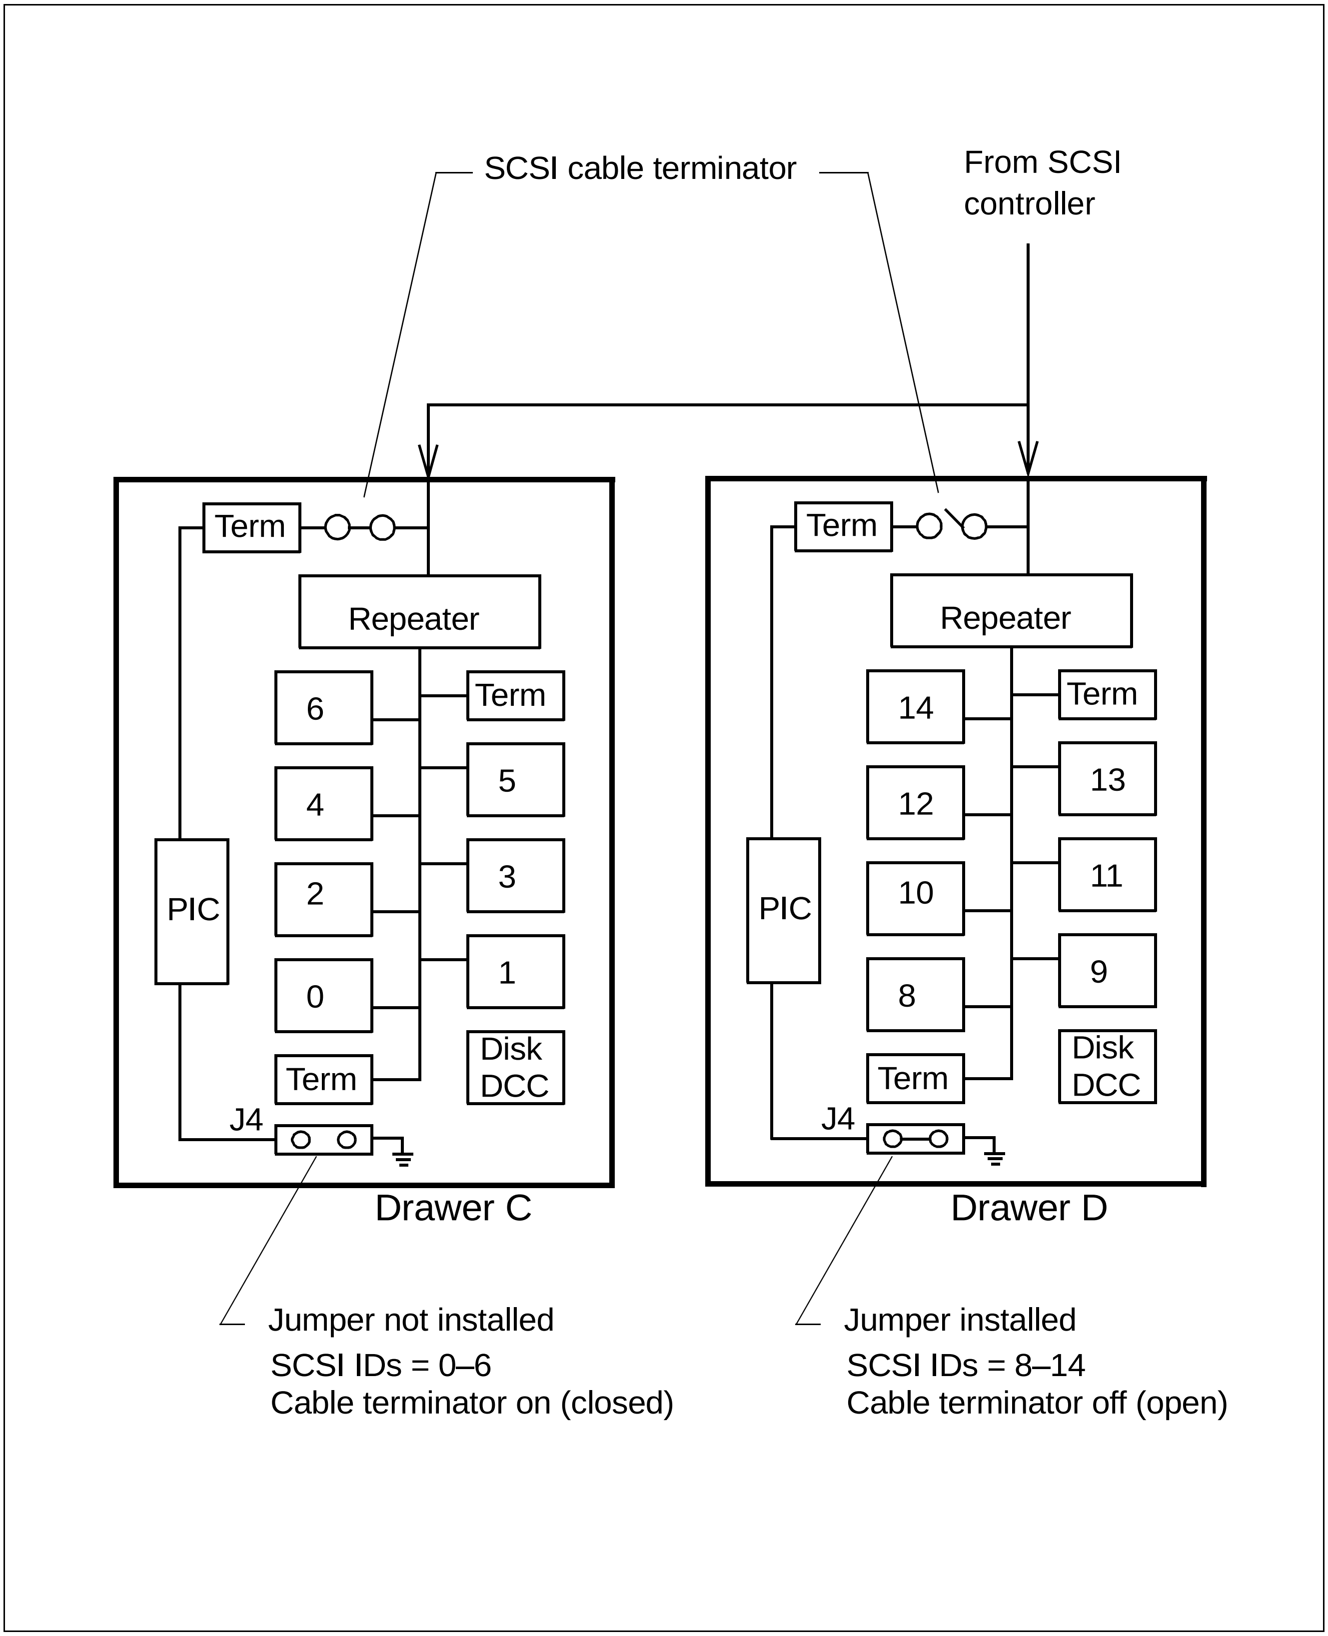

Daisy-chain SCSI ID and termination

drawers connected with point-to-point cables are configured and numbered as in *Drawer C* in the figure *Daisy-chain SCSI ID and termination*. This figure also shows that inserting the jumper on J4 opens the SCSI cable terminator and sets the high order IDs for the affected drawer. This jumper is inserted only when the SCSI bus is to be continued to the adjacent, daisy-chained drive drawer. Jumper pad J4 is located on the motherboard between drive 0 and the Disk DCC. Note that the cable from the controller initially connects to Drawer D then daisy-chains to Drawer C.

SCSI Repeater

The SCSI repeater chip installed in the HDDA drive drawer motherboards provides the bus extension necessary for the SCSI cable length and number of SCSI devices in a pair of daisy-chained drawers. Without a bus extender, the Wide UltraSCSI used in the M2000 supports only seven SCSI devices and 1.5 meters of cable length per channel.

Caution: Unless authorized by Auspex Technical Support, do not install in the M2000 an HDDA drive drawer from a different type of Auspex server.

fsp0m1c2t4

fsp refers to the File and Storage Processor (CPU) on the ECHaP. In this specific example, *fsp0* identifies node 0 or the *base node*. The FSPs for expansion nodes, nodes 1 and 2, would be identified as *fsp1* and *fsp2*, respectively.

m refers to Mylex, the manufacturer of the SCSI controller used in the M2000. An M2000 node supports up to three Mylex controllers, *m0*, *m1*, and *m2*.

c refers to a SCSI channel. The Mylex controller supports three SCSI channels; c0, c1, and c2.

t refers to the target drive. Each SCSI channel supports up to seven disk drives (one HDDA drawer) with a point-to-point cable, and up to 14 disks (two HDDA drawers) with a daisy-chained cable. Target drives for a given channel are numbered 0–6 (low-order SCSI IDs), and 8–14 (high-order IDs).

M2000 SCSI device numbering

SCSI Numbering

The figure *M2000 SCSI device numbering* shows how SCSI device numbers are developed under Solaris 2.6.

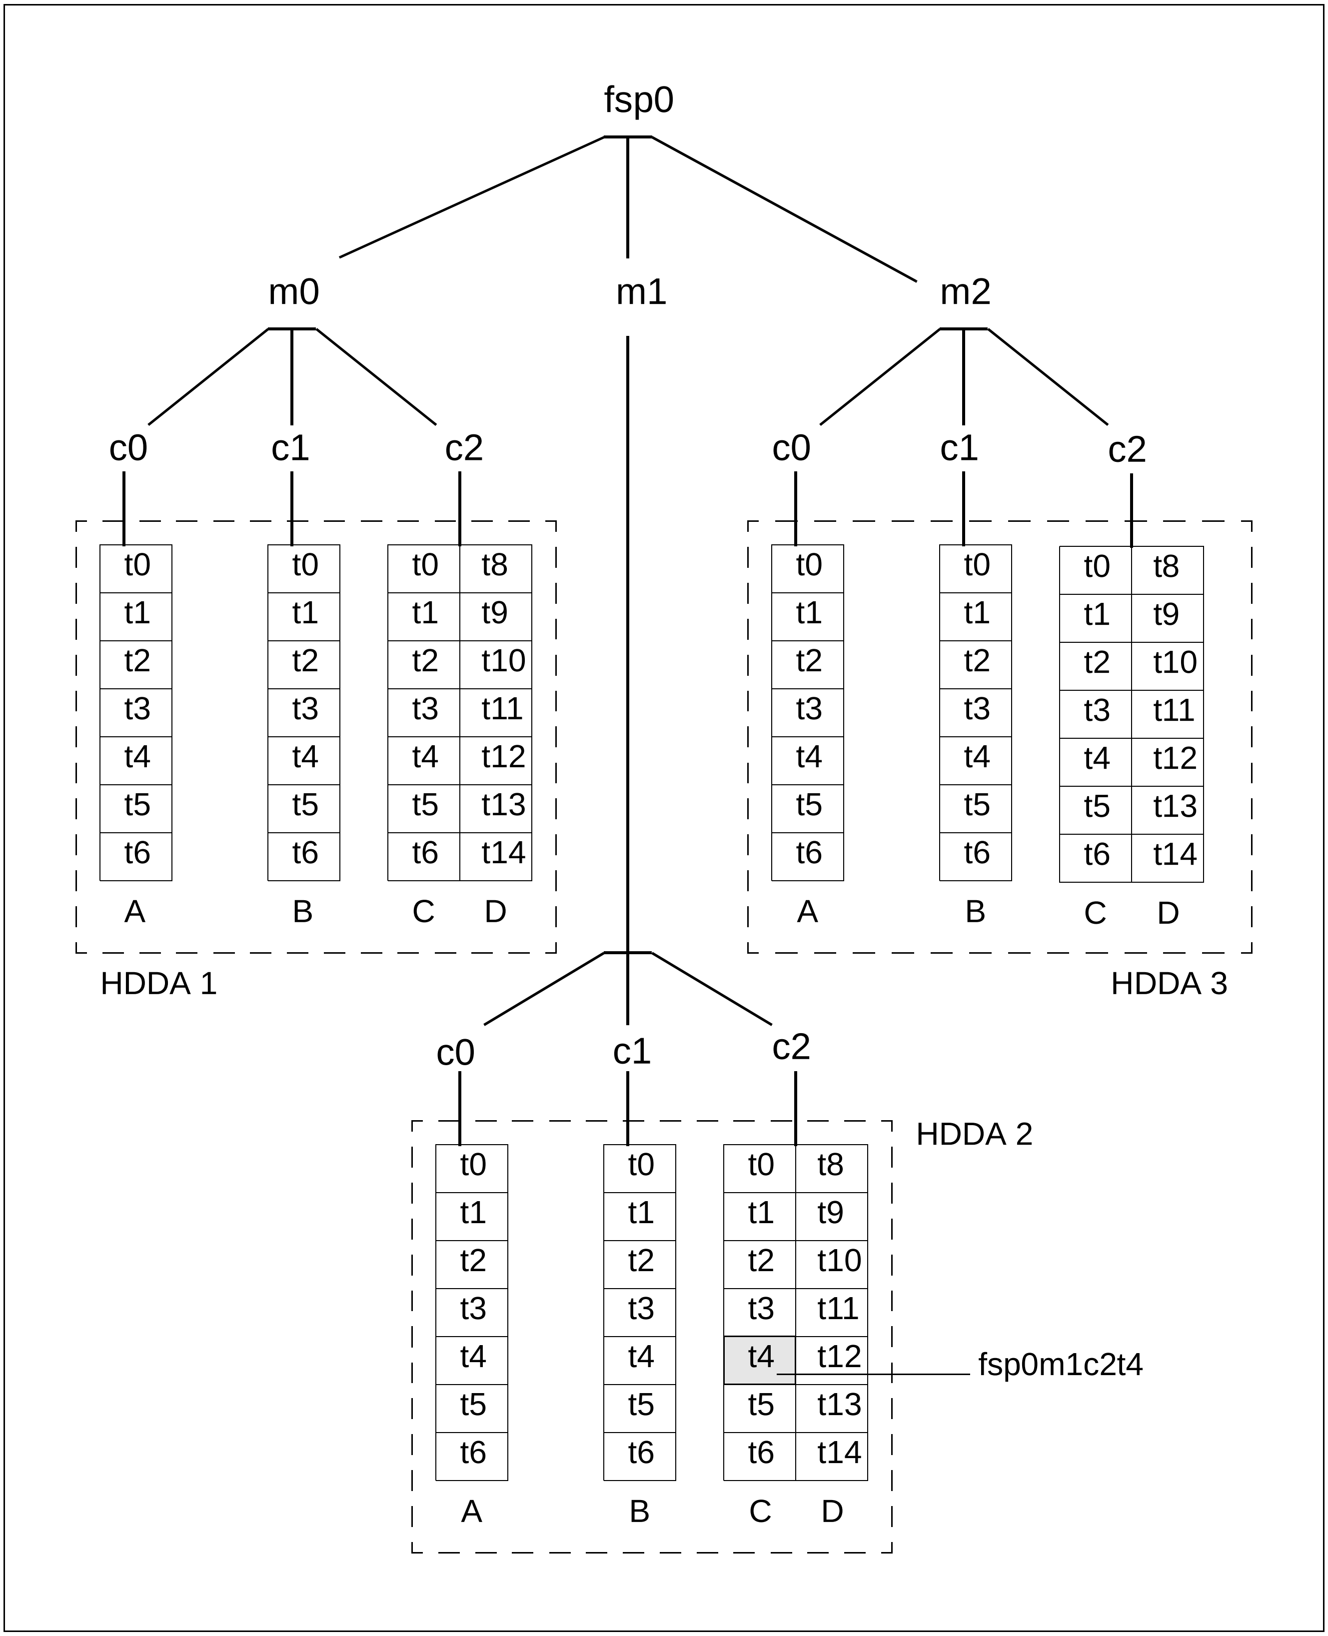

Base node SCSI map

Base Node SCSI Map

The figure Base node SCSI numbering shows the numbering scheme applied to a fully-configured base node. It also shows the map location of the numbering example from the previous figure, *M2000 SCSI device numbering.*

Storage subsystem FRUs		
Part number	**Part name**	**Description or function**
500705	Chassis assembly, HDDA	Houses up to four drive drawers.
500913	Drawer assembly, HDDA	Houses up to seven drives and one DCC power supply. Two sections: lower pan and drive drawer.
500782	Disk drive assembly, 9 GB, half-height,	Drive assembly for the HDDA only. Equipped with rails and an extractor.
500880	Disk drive assembly, 18 GB, half-height,	Drive assembly for the HDDA only. Equipped with rails and an extractor.
500783	DC-to-DC converter, 130W	Converts 48V to 5V and 12V for the HDDA drives; one DCC per drive drawer.
500706	Fan assembly, HDDA	Cooling fan for the HDDA chassis.
360013	Fuse, 5A, 125V, subminiature	Protects the drive drawers: one per drawer.
450361	Cable assembly, SCSI, 68 pin-VHDCI	Point-to point cable that connects the SCSI controller to the HDDA.
450373	Cable assembly, SCSI, 68 pin-VHDCI	Daisy-chain cable that connects the SCSI controller to two HDDA drawers.
100158	Drive array environmental monitor LON	EM network logical node for the HDDA.

9 Power Subsystem

Introduction

The M2000 has a distributed power subsystem with three main elements; each element performs a different function in the server.

Objectives

After completing this chapter you should be able to do the following:

- ▲ Generally describe the M2000 power subsystem components.
- ▲ State the voltage(s) provided by each power subsystem component and which M2000 device(s) use the voltage.

Power Subsystem Overview

Highly distributed subsystem with three main elements:

Bulk power supply modules that plug into the Power Shelves and provide 48 VDC to the Host, the E-Box(es), and the HDDA(s)

The *Disk DCC* that converts 48 VDC into 5 and 12 volts required by the HDDA disk drives

The *Logic DCC* that converts 48 VDC into the operating voltages required by the E-Boxes

Other power subsystem elements include:

The *VRMs* that regulate ECHaP CPU voltage

The DC-to-DC converter on the HDDA LON board

+5 VSB (standby voltage) for critical components of the Environmental Monitor

Overview

The highly distributed M2000 power subsystem uses three main elements to distribute and regulate server power:

- ▲ The bulk power supply modules that plug into the Power Shelves and provide 48 VDC to the Host, the E-Box(es), and the HDDA(s)
  - – The 48 VDC output of the bulk modules in a given Power Shelf is common; a specific module does not supply power to a specific system component. However, Power Shelves do not generate a common output; each shelf has a separate AC input and separate 48 VDC output.
- ▲ The DC-to-DC converter (Disk DCC) that converts 48 VDC into 5 and 12 volts required by the HDDA disk drives. Each HDDA drive drawer contains a Disk DCC.
- ▲ The DC-to-DC converter (Logic DCC) that converts 48 VDC into the operating voltages required by the E-Boxes. The Logic DCC mounts to the ECHaP motherboard.
- ▲ Other power subsystem elements include:
  - – The Voltage Regulator Modules (VRMs) that regulate ECHaP CPU voltage
  - – The DC-to-DC converter that is mounted on the HDDA LON board and converts 48 VDC to 5 VDC for use on this board
  - – +5 VSB (5 volts standby) from the Logic DCC for critical parts of the Environmental Monitor (EM/LCD board and transceivers). This standby voltage allows the EM/LCD to turn the E-Box on or off, and may allow segments of the EM/LCD to stay active after a +5 VDC failure.

System Power Requirements

The Power Base in an M2000 node supports two Power Shelves:

One shelf with three bulk power supply modules provides N+1 redundant power for Host, E-Box, and one HDDA

Second HDDA requires additional shelf with one power supply module to operate HDDA and one module for redundancy. Third HDDA requires one more power module.

Redundant operation:

N+1 redundancy allows a single power supply to fail and be replaced without disrupting server operation

The 48 VDC bulk supplies have N+1 redundant capacity; the Logic DCCs and Disk DCCs cannot be made redundant

System Power Requirements

The Power Base in an M2000 node supports two Power Shelves. One shelf with three bulk power supply modules provides N+1 redundant power for Host and E-Box modules, and one HDDA. Adding a second HDDA requires an additional shelf with one power module to operate the HDDA and one module for redundancy. Adding a third HDDA (maximum configuration) requires one more power module.

Power Subsystem Redundancy

N+1 redundancy means that for a specific server configuration, the minimum number of power supplies are installed plus one more is added for redundant operation. This allows a single power supply to fail and be replaced without disrupting server operation.

In the M2000, the 48 VDC bulk supplies have N+1 redundant capacity; the Logic DCCs and the Disk DCCs cannot be made redundant.

Note: Each M2000 Power Shelf has its own power cord and requires a separate AC source. Likewise, each shelf provides 48 VDC to different M2000 modules. Because their output is not common, each Power Shelf requires an extra power supply module to achieve N+1 redundancy.

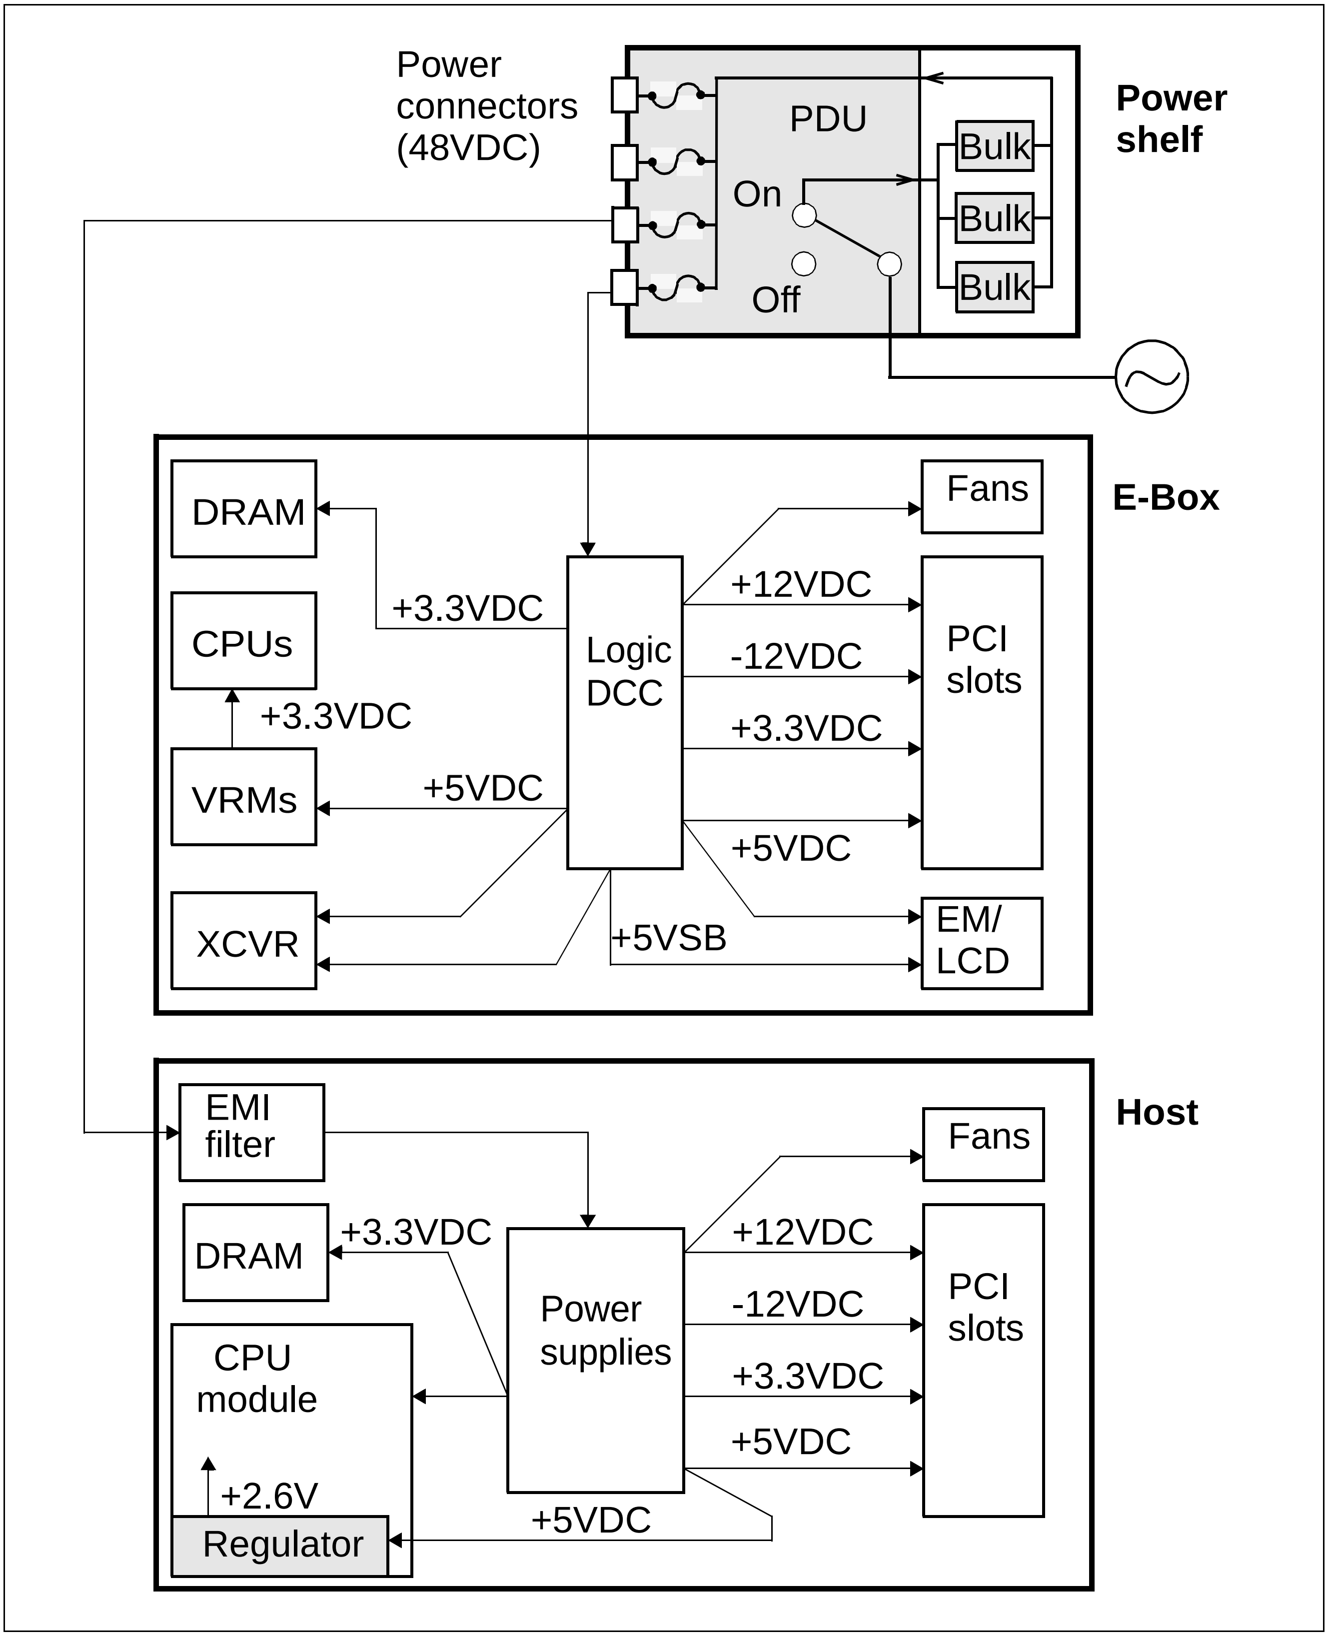

Processing subsystem power distribution

Power Distribution

Processor Subsystem

Refer to the figure *Processor subsystem power distribution* to correlate the following description.

The PDU and the bulk power supplies plug into the Power Shelf. The PDU receives the AC input from the wall outlet and routes it to the bulk power supplies. The bulk supplies convert the AC input to 48 VDC and route it to the power connectors on the rear of the PDU. From the power connectors, 48 VDC is cabled to the E-Box and Host module.

E-Box Power Distribution

In the E-Box, the Logic DCC receives 48 VDC and converts it to +3.3 VDC, +5 VDC, +5 VSB, +12 VDC, and -12 VDC. Power is routed to the major processor components as follows:

- +3.3 VDC, +5 VDC, +12 VDC, and -12 VDC to the PCI card slots
- +12 VDC to the fans
- +3.3 VDC to the DRAM DIMMs (main memory) and some data path devices (PCI chipset and GTL+ bus termination)
- +5 VDC to the VRMs
- +5 VDC and +5 VSB to the Environmental Monitor/Liquid Crystal Display (EM/LCD) board
  - The +5 VSB provides standby voltage to the critical areas of the EM/LCD which allows it to turn the E-Box on or off and may allow it to stay active after a 5 VDC failure.

Host Power Distribution

In the Host, the 48 volts from the Power Shelf routes through an EMI filter to a pair of hot-pluggable power supplies. The supplies convert the

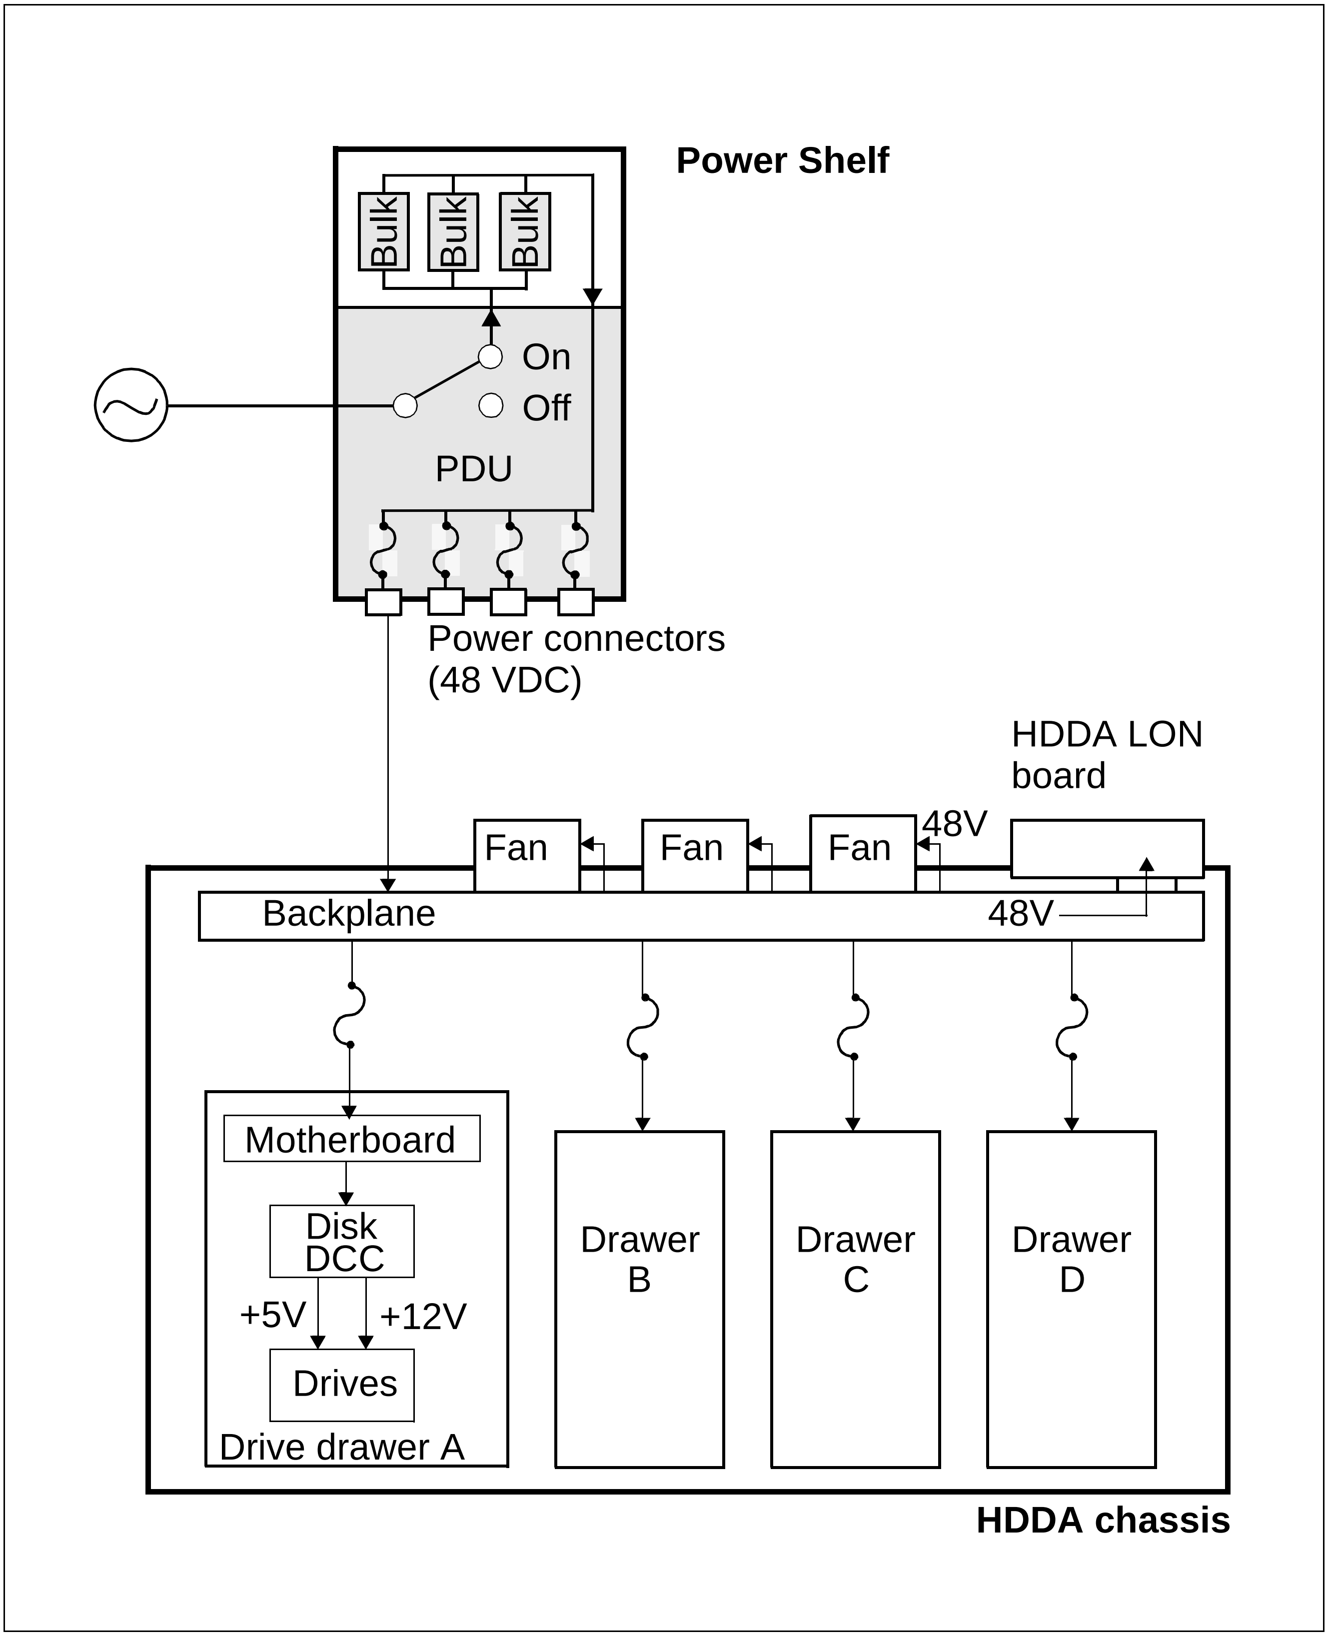

HDDA power distribution

48 volts to +3.3 VDC, +5 VDC, +12 VDC, and -12 VDC. Power is routed to the major Host components as follows:

- ▲ +3.3 VDC, +5 VDC, +12 VDC, and -12 VDC to the PCI card slots
- ▲ +12 VDC to the fans
- ▲ +3.3 VDC to the DRAM DIMMs (main memory)
- ▲ +3.3 VDC and +5 VDC to the CPU (the 5 volts is regulated to 2.6 VDC for the CPU)

Storage Subsystem

Refer to the figure *Storage subsystem power distribution* to correlate the following description.

The PDU and the bulk power supplies plug into the Power Shelf. The PDU receives the AC input from the wall outlet and routes it to the bulk power supplies. The bulk supplies convert the AC input to 48 VDC and route it to the power connectors on the rear of the PDU. From the power connectors, 48 VDC is cabled to the HDDA(s).

HDDA Power Distribution

The HDDA backplane receives 48 VDC, routes it to the fans that plug into the rear of the HDDA, and to the HDDA LON board. The HDDA LON board has a DC-to-DC converter that converts the 48 volts into 5 volts for use on the board.

In each HDDA drive drawer, the Disk DCC converts the 48 VDC input into 5 and 12 VDC for the HDDA disk drives: 5 VDC for the drive logic and 12 VDC for drive motors.

Fuses are electrically in the 48V path for each drive drawer umbilical cable. This protects the HDDA drive drawers from a shorted umbilical cable.

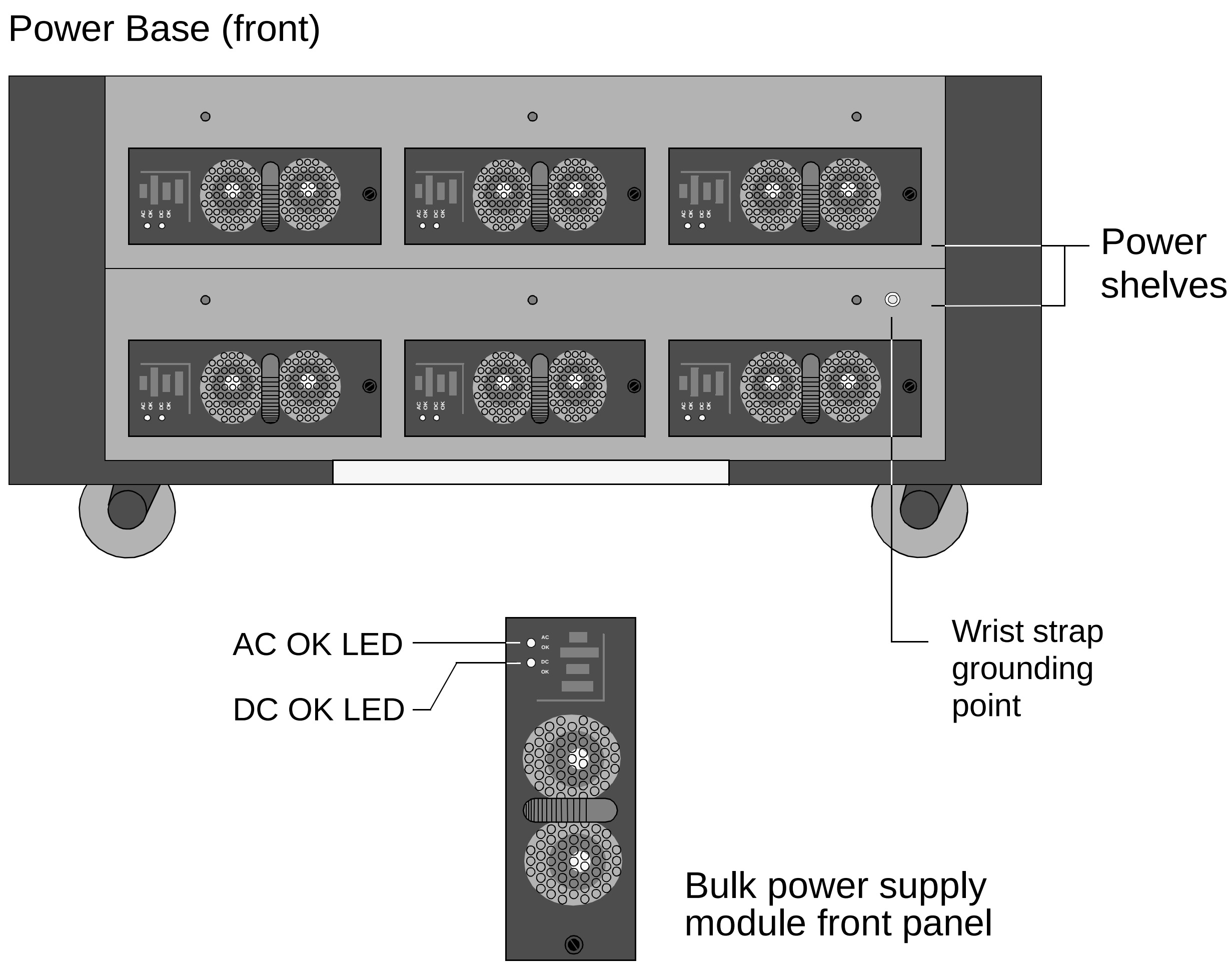

Bulk power supply status indicators

LEDs	Power supply status	Required action
AC OK green and DC OK green	Fully operational.	No action required.
AC OK off or DC OK off	Power supply failure or over-temperature condition.	Unplug module for 30 seconds, then reinstall it. If it doesn’t operate properly, replace it.
None	The supply is not receiving AC power.	Check the power cable connection at the wall outlet and system power switch on the PDU.

Bulk power supplies

Power Subsystem Components

Power Shelf

The Power Base, a crate structure at the bottom of the M2000 stack, supports one or two Power Shelves, depending on system requirements.

The Power Shelf is an enclosure that houses up to three 48 VDC bulk power supply modules and a pluggable Power Distribution Unit (PDU). For the M2000, one shelf is required; a second shelf is added to accommodate additional HDDAs (up to three). In this case, the second (optional) shelf is stacked above the standard shelf.

Bulk Power Supply Modules

The bulk power supplies convert the AC input into 48 volts DC for distribution to the Host, E-Box, and HDDA(s). The bulk power supply modules do not have power on/off switches; the operator must turn server power on or off at the main power switch on the PDU. The modules are hot-pluggable if the configuration is N+1 redundant.

Checking Bulk Power Supply Status

The figure *Bulk power supplies* shows the location of the Power Shelves and bulk power supply modules. Each bulk supply has two LEDs: *AC OK* and *DC OK*. During normal operation, the green AC OK LED lights to indicate the supply is receiving AC power. The green DC OK LED lights to indicate the supply is providing DC power. If either LED is off, the power supply may have failed. For more information, refer to the table *Bulk power supply status indicators*.

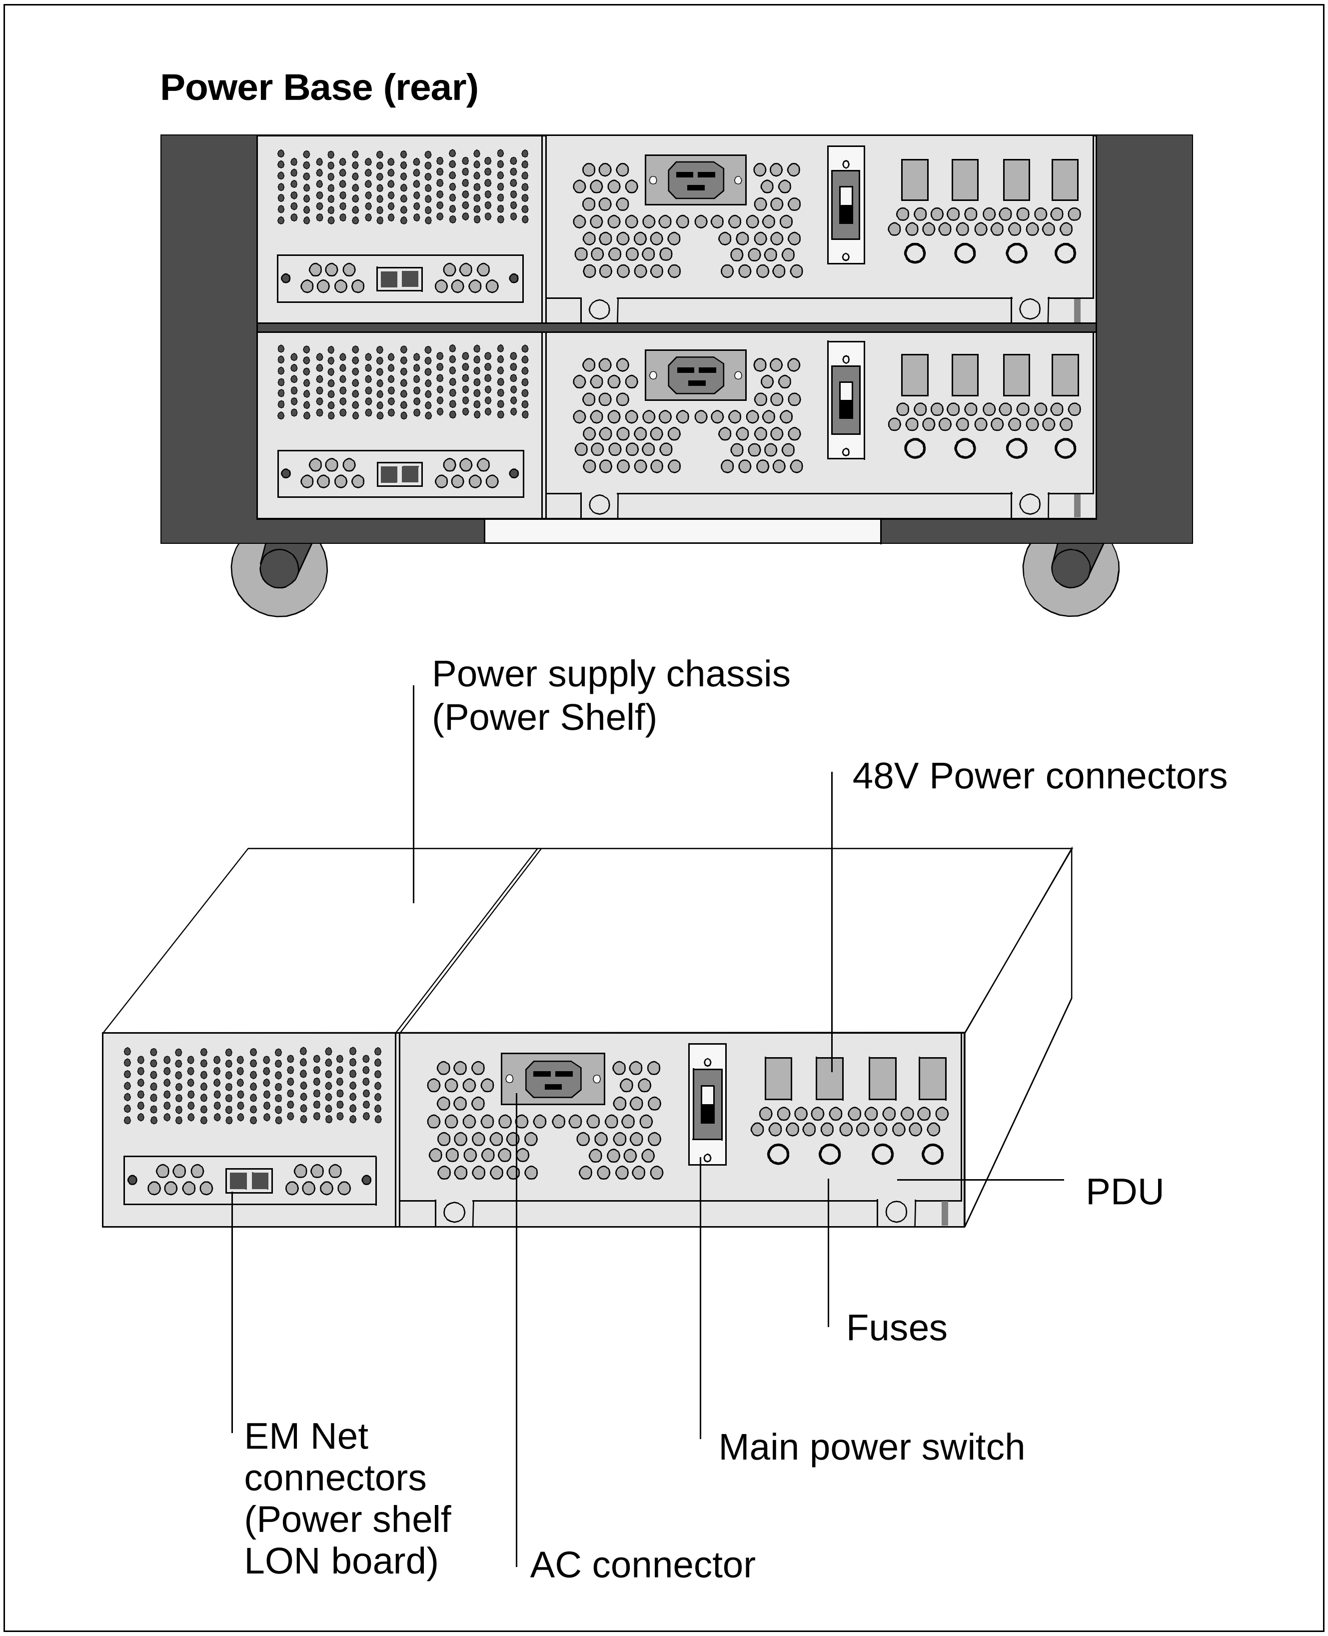

Power Distribution Unit (PDU)

Power Distribution Unit (PDU)

The figures *Processor subsystem power distribution* and *HDDA power distribution* show a simple block diagram of the Power Distribution Unit (PDU).

The PDU plugs into the Power Shelf. It directs the AC input voltage to the bulk power supplies. The output of the bulk supplies (48VDC) is routed back to the PDU and is available from fused power connectors on the rear of the PDU. From these connectors, power is cabled to the Host, E-Box, and the HDDA(s).

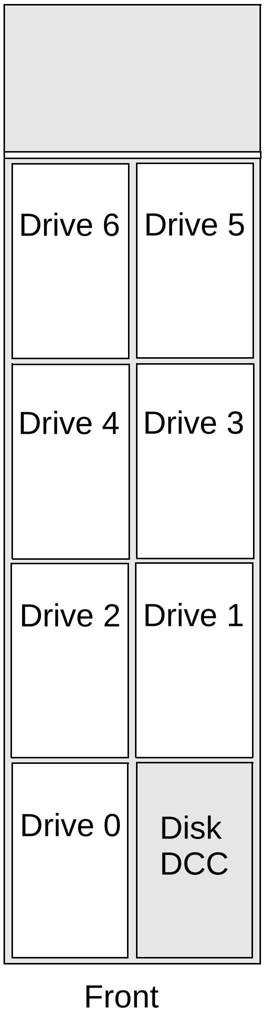

HDDA drive drawer (top)

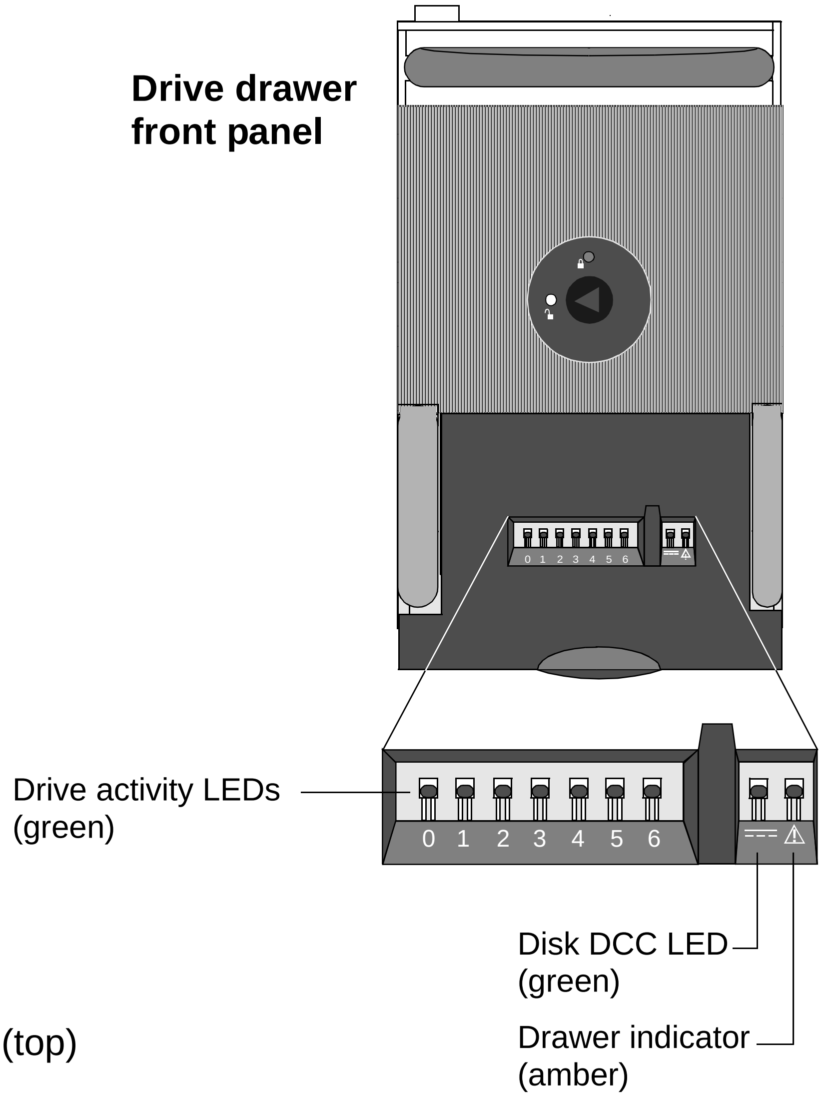

Disk DCC status indicator

LEDs	Power supply status	Required action
Green LED lit	DCC operating properly	None
Green LED extinguished	Power supply failure	Follow *Field Service Guide* instructions for checking and replacing Disk DCC
Amber LED	N/A	None. Flashes in response to command to add or replace drive in affected drawer.

Disk DCC and status indicator

Disk DC-to-DC Converter (Disk DCC)

In each HDDA drive drawer, a DC-to-DC module (Disk DCC) converts the 48-volt input from the PDU into 5 and 12 volts for the disk drives. Redundancy is not available for the Disk DCC subsystem. For more information and pictorial detail, refer to the figure *Disk DCC and status indicator.*

DCC Status Indicator

On the front panel of each HDDA drive drawer, a green LED indicates the status of the Disk DCC in that drawer. During normal operation, the LED lights to indicate the supply is receiving 48-volt power and providing 5- and 12-volt power. If the LED is off, the supply has failed. For more information, refer to the table *Disk DCC status indicator*.

Drawer Indicator

To the right of the DCC status indicator is the drawer indicator; an amber LED. The drawer indicator is extinguished during normal operation. The indicator flashes when a command is entered to replace or add a drive in the affected drawer (planned function; not yet operational).

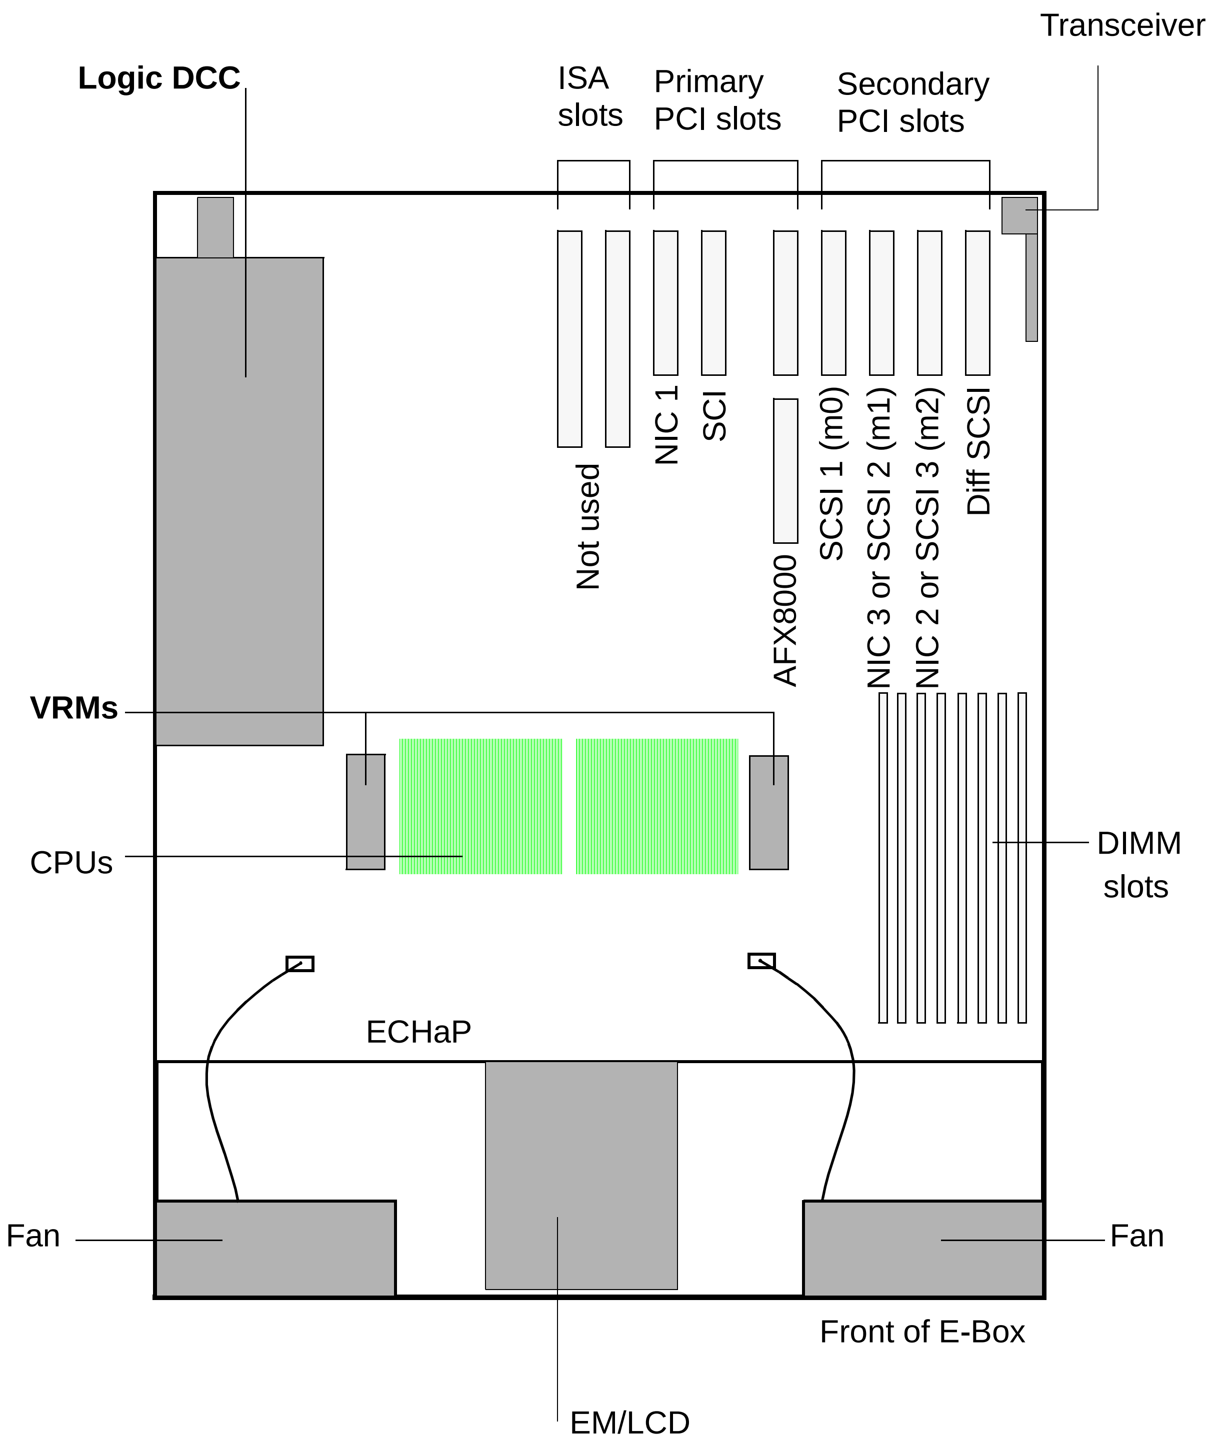

Voltage Regulator Module and Logic DCC

Logic DC-to-DC Converter (Logic DCC)

The Logic DCC mounts to the ECHaP motherboard and directly receives the 48-volt input from the Power Shelf. It converts the 48 volts into the operating voltages required by the E-Boxes: +3.3 VDC, +5 VDC, +5 VSB, +12 VDC, and -12 VDC.

Incorporated into the Logic DCC is a +5 VSB (5 volts standby) source to the EM/LCD module. This allows the EM/LCD to continue operating if there is a 5 VDC failure on the Logic DCC (assuming the + 5 VSB source is not coincidently affected). If the EM/LCD remains operational, the system operator can use the EM/LCD interface to troubleshoot the server and turn the E-Box on and off.

Voltage Regulator Module (VRM)

A Voltage Regulator Module (VRM) regulates the input voltage to each ECHaP CPU. It receives +5 VDC from the Logic DCC and has a +3.3 VDC, 12.4 amp output. The VRMs are field replaceable units (FRUs) and their location is shown in the figure *Voltage Regulator Module and Logic DCC*.

Power subsystem FRUs		
Part number	**Part Name**	**Description/Comments**
640154	Power supply chassis, 3-slot	Power shelf enclosure that houses the PDU and bulk power supplies.
650024	Power supply module, 48 VDC, 600 watts	Bulk supplies that power the Host, E-Boxes, and HDDA(s).
690006	PDU, 3-slot power supply chassis	Routes AC power to the bulk supplies; supports four fused, 48-volt power connectors.
500783	Power supply, DC/DC, 130 W	Disk DCC that converts 48-volt input to 5 and 12 volts for the HDDA disk drives.
650025	Power Supply, Logic DCC	Logic DCC that converts 48-volt input into operating voltages for the E-Box.
600063	Voltage Regulator Module, 5V	Regulates input voltage to each ECHaP CPU.
450374	Cable assembly, power, 1.5 meters	Power cable that connects the PDU to the Host module.
450315	Cable assembly, power, 1 meter	Power cable that connects the PDU to the E-Box.
450372	Cable assembly, power, 0.75 meter	Power cable that connects the PDU to the HDDAs.
450321	Power cord, IEC309	Power cord for Japan, China, South America, Australia, Italy, Switzerland
450322	Power cord, NEMA 5-20	Power cord for North America, Taiwan, and Saudi Arabia
450323	Power cord, CEE7/7	Power cord for Israel, and most of Europe
450324	Power cord, 2.5 M, 13A	Power cord for UK, Hong Kong, Malaysia, Singapore

10 Environmental Monitor Subsystem

Introduction

The Environmental Monitor (EM) subsystem monitors the M2000 internal environment, provides a user interface, and generates control signals. The EM subsystem network, like the SCI, provides a communication backbone for the M2000.

Objectives

After completing this chapter, you should be able to do the following:

- State the key functions of the EM subsystem.
- Describe the major components of the EM subsystem.
- Give an operational overview of the EM subsystem.

Key Features

Front panel: user interface on each E-Box

Environmental measurements: temperatures, voltages, fan speeds

EM Console: character-based interface organized in menus

Generates control signals: power supply control, ECHap reset

Robust, independent operation

Validates EM Network is contiguous

Provides access to NP and FSP consoles

Allows verification of firmware-level compatibilities

Key Features

The M2000 Environmental Monitor subsystem:

- provides a standalone front panel operation mode for debug of E-box startup problems
- measures the server’s internal environment
  - Temperatures
  - Voltages
  - Fan speeds
  - Other, specific measurements for the E-Box, HDDA, and Power Shelves
- provides the EM Console, a character-based user interface organized in menus
- generates platform control signals based on input from the EM Console or the Front Panel
  - Power supply control
  - ECHaP reset
  - Select Auxiliary PROM
- provides robust operation, independent of all network, SCI, and SCSI connections:
  - M2000 system continues normal operation under most EM subsystem failures.
  - EM Network signals are slew rate limited to reduce EMI radiation.
  - The EM Network can withstand and recover from many transient interruptions such as electrostatic discharge (ESD).
- validates EM Network is contiguous
- provides access to NP and FSP consoles

Student notes

- allows verification of firmware-level compatibilities

Note: Environmental monitoring for the M2000 Host is accomplished via the Host motherboard hardware and a third party software utility, and is independent of all other elements described in this chapter.

Major Functions

Front Panel

Normal Mode

E-Box Debug Mode

EM Console

Character-based user interface organized into menus

Displays environmental data from E-Boxes, HDDAs, and Power Shelves

Accepts keyboard input to control low-level hardware settings

Major Functions

Front Panel

The Front Panel is part of every E-Box and consists of a two row by 16-character LCD, and four push-buttons:

- *Reset* generates a hard reset signal to the ECHaP and the E-LP board but does not reset the EM/LCD assembly (the E-LP and EM/LCD are described later). Resetting the E-LP reinitializes the EM Network. The Reset function operates independently of any other Front Panel function or mode. The button is recessed and requires that you insert a foreign object to press it.
- *Scroll Up*
- *Scroll Down*
- *Select* is context-sensitive

Pressing the Scroll buttons simultaneously toggles between the Normal and Debug operating modes (described in the next section). Simultaneously pressing any combination of Scroll and Select keys generates a unique, context-sensitive input to the EM subsystem.

Operational Modes

The Front Panel operates in two distinct modes: *Normal Mode* and *E-Box Debug Mode*. Mode selection is implemented solely within the EM subsystem hardware and firmware.

Normal Mode

The Normal Mode:

- provides a local display for the BIOS, diagnostics, ECHaP Monitor (EMON), operating system drivers, and applications. While this mode is active, the LCD will flash the Auspex logo every few seconds to indicate the system is alive.

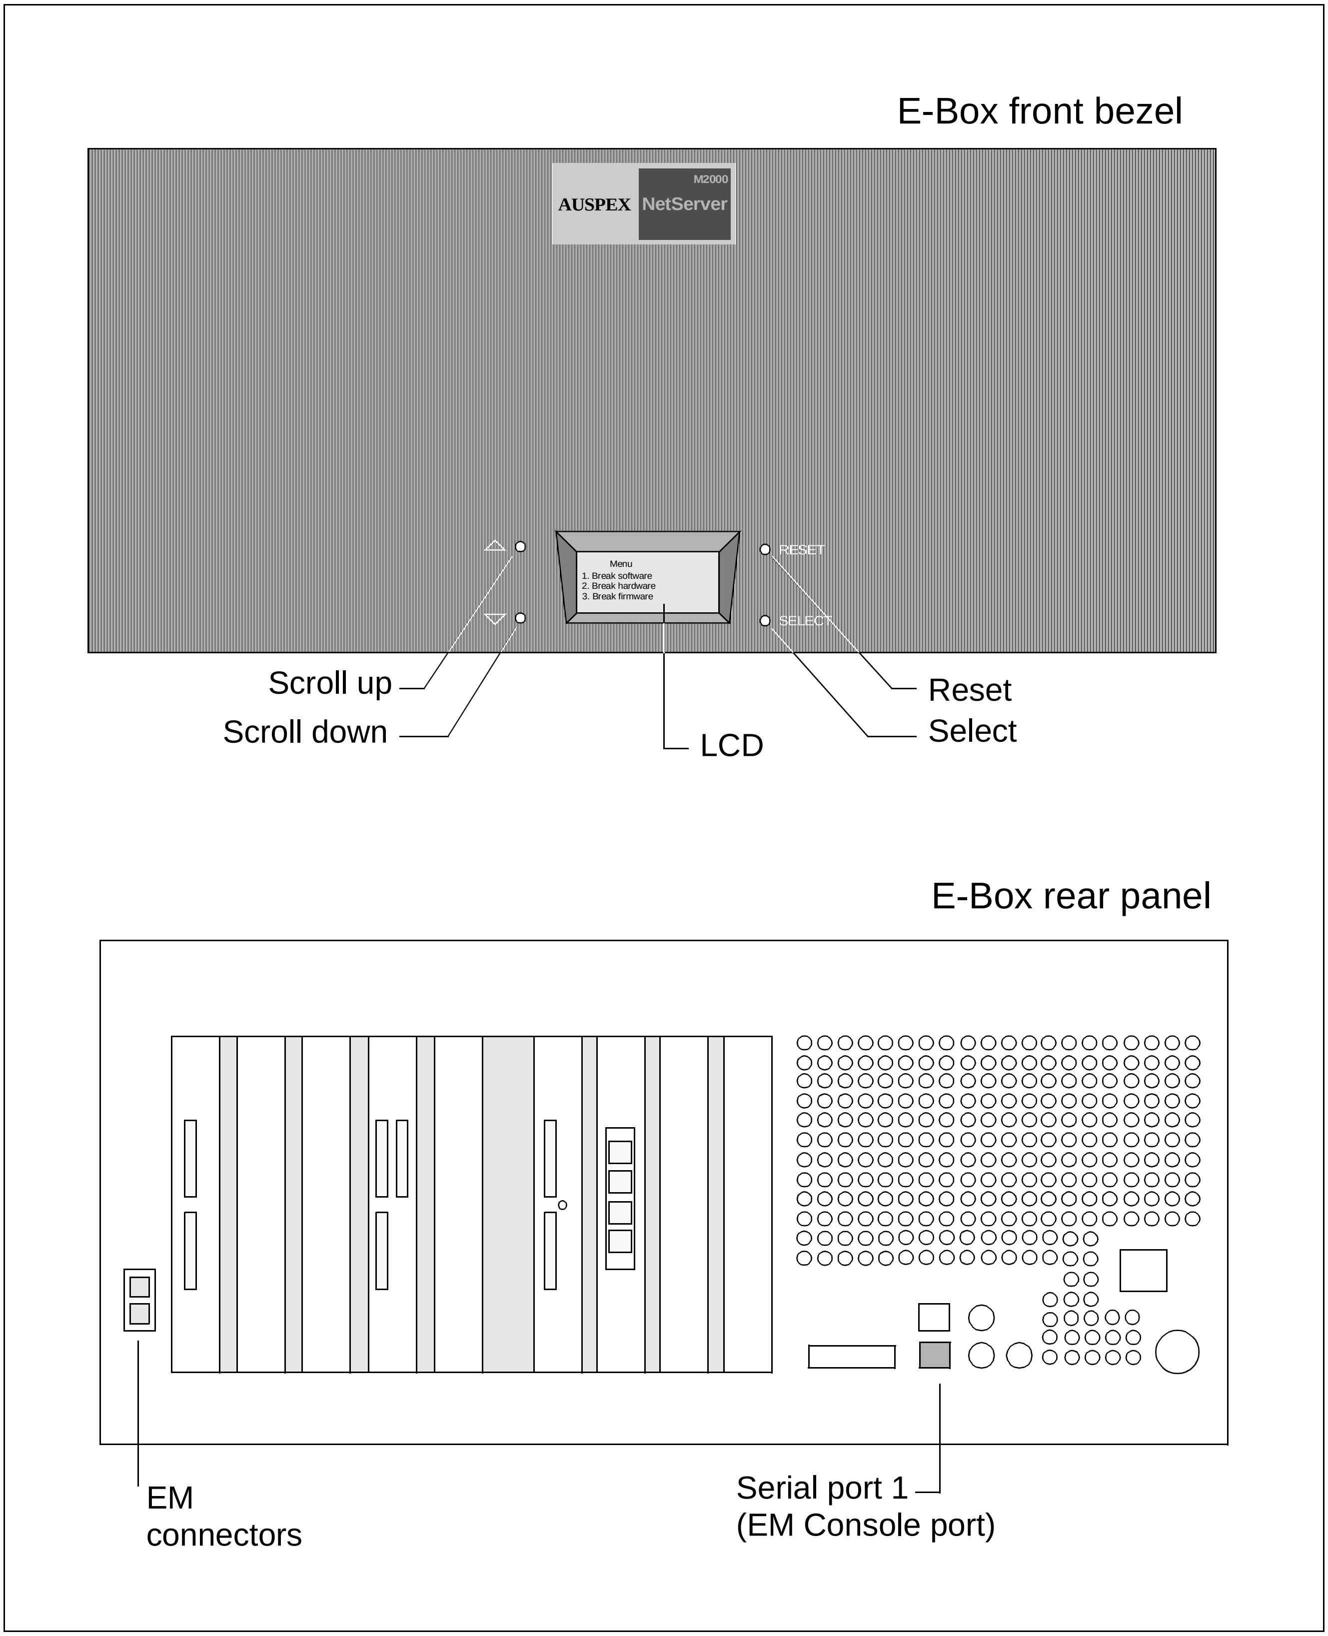

Front Panel and EM Console port

E-Box Debug Mode

The E-Box Debug Mode is for diagnosis of E-Box problems. The hardware and firmware that support this mode have limited memory and processing power but operate independent of ECHaP resources.

The E-Box Debug Mode:

- ▲ allows the operator to power the ECHaP CPUs on or off.
- ▲ can reset the ECHaP CPUs, but will **not** reset the E-LP board (as will the Reset button on the Front Panel).
- ▲ provides access to the most recent Hexadecimal code written to PORT 80 by the ECHaP CPUs.

EM Console

The EM Console is a character-based user interface organized into menus. It displays environmental data from the E-boxes, HDDAs, and Power Shelves. The EM Console accepts keyboard input to control low-level hardware settings such as voltage control. Refer to the tables that follow this section for lists of subsystem-specific functions displayed on the EM Console.

The EM Console displays:

- ▲ on a dumb terminal connected to serial port 1 on the master E-Box (the base node E-Box). Refer to the figure *Front Panel and EM Console port* for the location of serial port 1.
- ▲ on the Host console (system console). This requires a cable connection between the Host serial connector and serial port 1 on the master E-Box. A terminal emulation utility running under Solaris 2.6 displays the EM Console as a window on the Host console.

Control and display functions: E-Box

Function	Front Panel display	Front Panel control	EM Console display	EM Console control
E-Box internal temperature	Yes	N/A	Yes	N/A
CPU case temperature (CPUs 0 and 1)	Yes	N/A	Yes	N/A
Fan speed (fans 1 and 2)	Yes	N/A	Yes	N/A
Voltage: 5V DCC output	Yes	N/A	Yes	N/A
Voltage: 3.3V DCC output	Yes	N/A	Yes	N/A
Voltage of 12V DCC output	Yes	N/A	Yes	N/A
Voltage of -12V DCC output	Yes	N/A	Yes	N/A
Voltage of CPU controlled DCC (CPUs 0 and 1)	Yes	N/A	Yes	N/A
ECHaP DCC “good” signal	Yes	N/A	Yes	N/A
ECHaP DCC on/off	Back light on/off	Yes	Yes	Yes
CPU Status codes (CPUs 0 and 1)	Yes	N/A	Yes	N/A
Port 80 buffer (most recent byte)	Yes	N/A	Yes	N/A
ECHaP Reset (Both CPUs, 0 and 1)	Yes	Yes	Yes	Yes
Select Auxiliary PROM	No	Yes	Yes	Yes
EM Net: break in the ring	No	N/A	Yes	N/A
Check if EM LCD application is operational (ping)	No	N/A	Yes	N/A
View firmware revisions	No	N/A	Yes	N/A

Control and Display Functions

The tables on the facing and following pages list the E-Box, HDDA, and Power Shelf parameters that are measured or controlled via the EM Console and Front Panel.

Control and display functions: HDDA

Function	EM Console display	EM Console control
Temperature (drawers 1–4)	Yes	N/A
Voltage of 12V DCC output (drawers 1–4)	Yes	N/A
Voltage of 5V DCC output (drawers 1–4)	Yes	N/A
Voltage of dedicated EM 5V power (drawers 1–4)	Yes	N/A
Drawer type single-ended (SE) SCSI (drawers 1–4)	Yes	Yes
Drawer present (drawers 1–4)	Yes	N/A
Voltage ID (V_{ID}): drawer position within HDDA (drawers 1–4)	Yes	N/A
48V input present (drawers 1–4)	Yes	N/A
HDDA DCC “good” signal (drawers 1–4)	Yes	N/A
Drawer DCC power enabled (drawers 1–4)	Yes	N/A
Fan speed (fans 1–3)	Yes	N/A
HDDA chassis temperature (at HDDA LON board)	Yes	N/A
View firmware revision data	Yes	N/A
View unidentified drawer data	Yes	N/A

Student notes

Control and display functions: Power Shelf

Function	EM Console display	EM Console control
Current outputs; min and max (supplies 1–3)	Yes	N/A
Power supply present (supplies 1–3)	Yes	N/A
Power supply enabled (supplies 1–3)	Yes	Yes
AC input good (supplies 1–3)	Yes	N/A
DC output good (supplies 1–3)	Yes	N/A
Thermal warning (supplies 1–3)	Yes	N/A
Chassis temperature (measured at LON board)	Yes	N/A

Student notes

EM Components

Environmental Monitor Network (EM Net) that interconnects the major elements of the EM subsystem

Processor boards (LON boards) that control communication and processing of environmental data

Islands of circuitry on the ECHaP and HDDA drive drawer motherboards

EM Network

LON topology

Service topology

EM Net cables

EM Components

The Environmental Monitor is a distributed subsystem consisting of the following components:

- ▲ The Environmental Monitor Network (EM Net). This is the dedicated communications network that connects the major hardware elements of the EM subsystem.
- ▲ Processor boards (three unique board designs) that control the communication and processing of environmental data
- ▲ Islands of EM circuitry on the ECHaP and the HDDA drive drawer motherboards

Environmental Monitor Network (EM Net)

The Environmental Monitor Network (EM Net) is the communications network that interconnects the major hardware components of the EM subsystem.

A board at the rear of each shelf (E-Box, HDDA, Power Shelf) provides connectivity for a set of point-to-point cables. The EM Net cables form a ring topology at the rear of the M2000 stack by daisy-chaining together the supported shelves. A cable connecting the last node to the first node completes the ring.

There are four conductors in the EM Net cables: the two-wire Local Operating Network (LON) for status and control data, and the two-wire service network for topology management.

LON Topology

The LON is a shared bus topology. When any node transmits and detects collisions, it attempts to re-transmit at a later time. The LON is

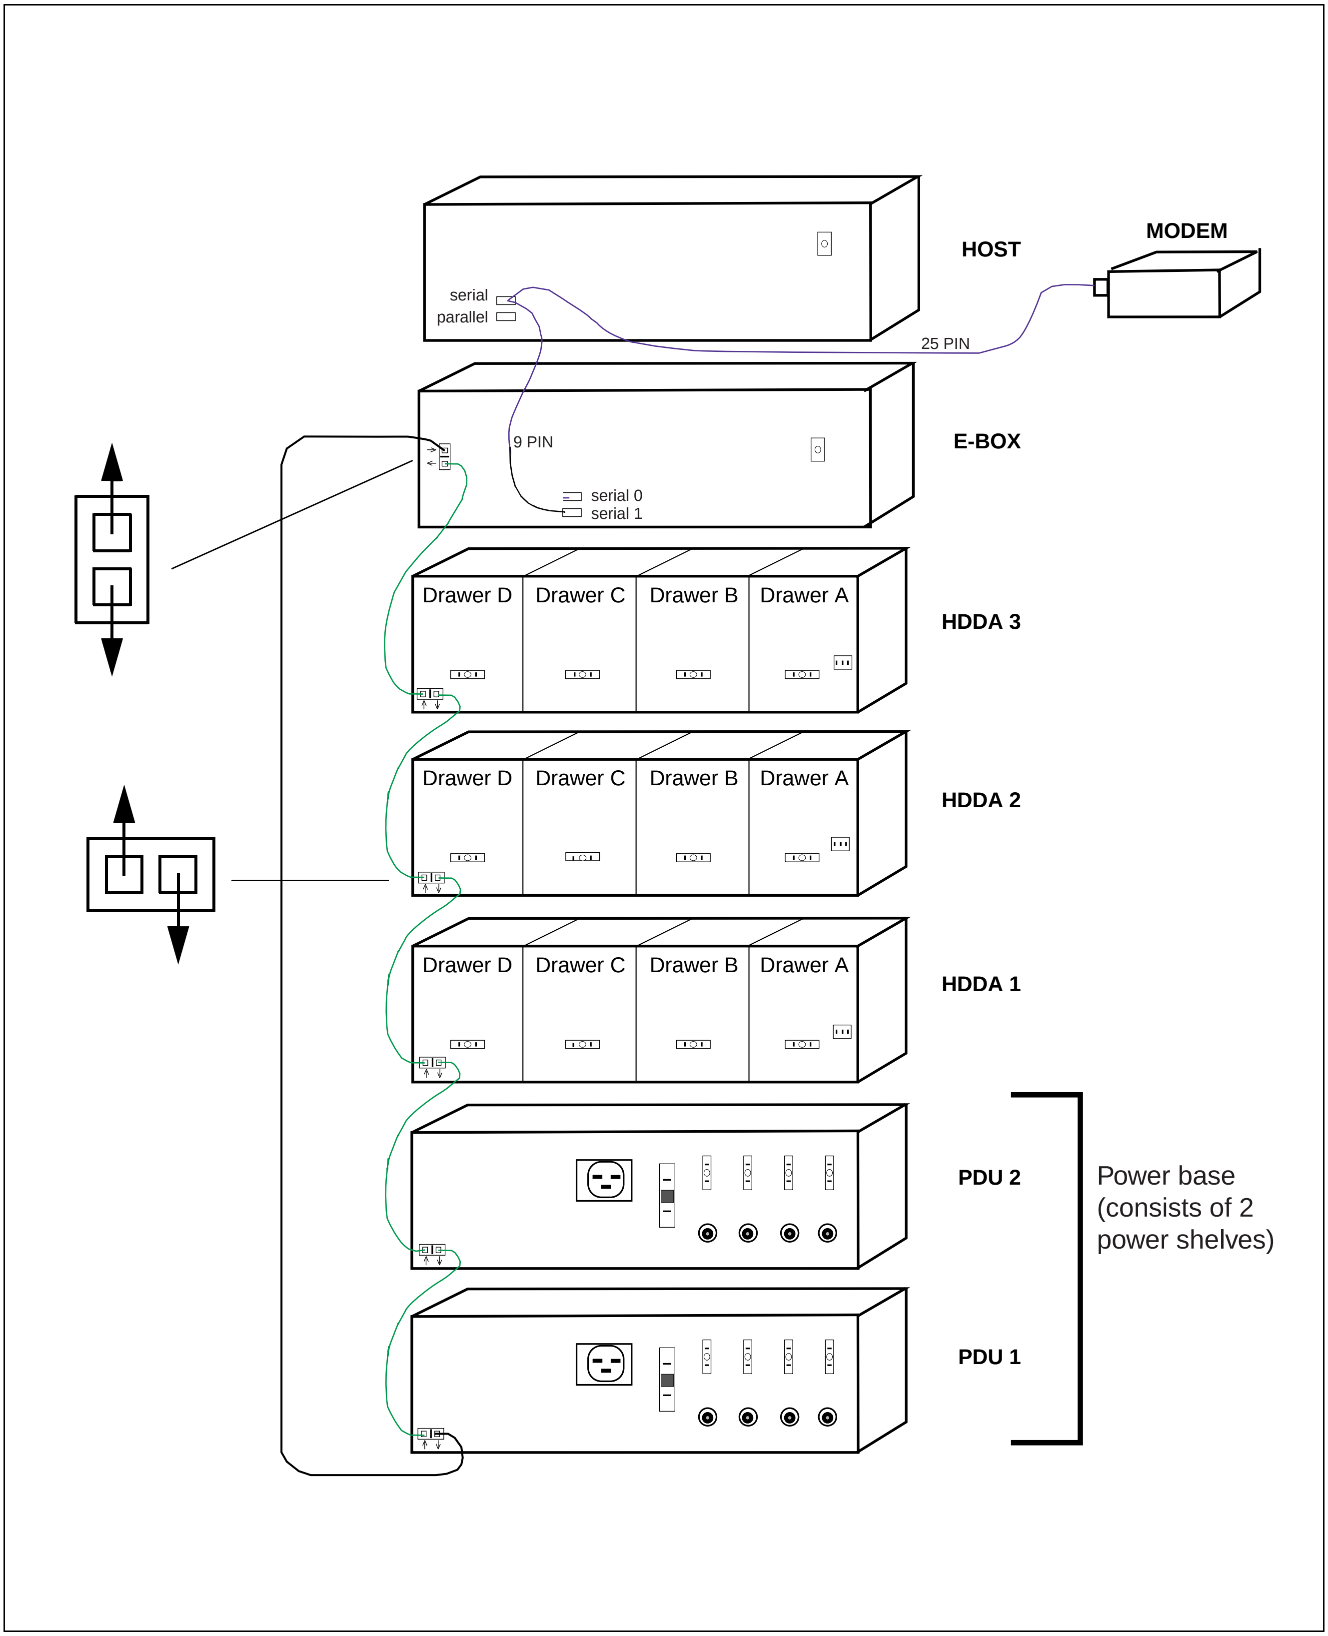

EM Net cable interconnection

terminated at the Master E-Box Transceiver board. All other transceivers operate as pass-through devices.

Caution: It is important to ensure that each M2000 system has exactly one Master E-Box Transceiver. If all transceivers are non-masters (pass through), the EM Net will function, but the unterminated EM Net cables collectively act as a large inductor. This leaves the M2000 vulnerable to electromagnetic interference (EMI) and radio frequency interference (RFI).

Service Network Topology

The service network is a set of point-to-point connections starting and ending at the E-LP board. These connections run through the same cables as the LON. The service network assists in mapping logical EM Net addresses to physical locations along the EM Net.

EM Net Connectivity and Communication

The figure *EM Net cable interconnection* shows the direction of connectivity. When the RJ45 connector pair is arranged horizontally (HDDA and Power Shelf), the left cable connects to the next higher module; the right cable connects to the next lower module. On the E-Box, the connector pair is vertically arranged: the bottom cable connects to the next lower module, the top cable connects to the bottom node in the stack to complete the ring.

Communication across the EM Net is relatively slow, but the EM subsystem does not require fast communication to be effective. The differential transceivers in the network are capable of transmitting at a maximum of 250 kilobits per second. The M2000 uses these transceivers at 156 kilobits per second.

EM Components (continued)

EM Hardware Elements

E-LP board

EM/LCD assembly

ECHaP monitor

Master E-Box transceiver

Slave E-Box transceiver

HDDA drive drawer PIC

HDDA LON board

Power Shelf LON board

EM Hardware Elements

E-LP Board

The Environmental LON Processor board (E-LP) is the master node of the EM Net, therefore each system has only one. The E-LP is assembled onto the Master E-Box's EM/LCD board as a daughter card, but it operates independently. It receives power and reset from the Logic DCC. The E-LP can be powered off or reset without affecting the EM/LCD assembly.

The E-LP board communicates with the EM Console through software on the ECHaP that transfers characters and commands between the E-LP (via ECHaP IDE bus) and ECHaP serial port 1.

EM/LCD Assembly

The Environmental Monitor Liquid Crystal Display Assembly (EM/LCD) consists of three components: the EM/LCD board, the LCD board, and the LCD. Because of the physical form of this assembly, it is sometimes referred to as the *sled*.

The EM/LCD board provides a LON slave node for status and control of an E-Box. It communicates with the ECHaP via the I^2C[1] bus and several dedicated signals. The EM/LCD assembly provides the Front Panel for each E-box.

ECHaP Monitor

The ECHaP contains a microprocessor system hardware monitor acting as a remote portion of the EM/LCD. It is a highly integrated IC that measures E-Box temperature, voltage, fan speed, and many key system signatures. This device also captures messages sent to the server's Port 80 address, and communicates error conditions via the I^2C bus.

---

1. The Inter-Integrated Circuit bus (I^2C bus) is a serial bus with associated protocol for communication among integrated circuits.

Student notes

Master E-Box Transceiver

The Master E-Box Transceiver, also called the E-LP transceiver, provides termination to both ends of the EM Net. It operates as an extension of the EM/LCD and E-LP boards. The board is physically located at the rear of the E-Box and provides the EM Net connection. Signals are carried between the transceiver and the EM/LCD and E-LP boards by traces on the ECHaP board.

Slave E-Box Transceiver

This board, also called the ECHaP transceiver, is similar to the Master E-box Transceiver board, but does not provide termination for the EM Net. It operates as a pass-through device and is used in all expansion E-Boxes in an M2000 system.

HDDA Drive Drawer PIC

A Programmable Integrated Controller (PIC) is installed on the motherboard in each HDDA drive drawer. This device communicates with another PIC on the HDDA LON board plugged into the rear of the HDDA chassis. The motherboard PIC provides status and control for the Disk DCC, disk drives, and the drawer/status indicator on the HDDA front panel (amber LED in the right window).

HDDA LON Board

The HDDA LON board plugs into the LON port on the rear of the HDDA chassis. It contains a PIC that performs the following functions:

- ▲ Communicates with the drive drawer PICs to gather drive drawer status
- ▲ Performs A-to-D conversion for analog measurements
- ▲ Generates control signals

EM Components (continued)

EM Firmware Elements

386EX code

Neuron

Programmable Integrated Controller (PIC)

Field Programmable Gate Array (FPGA)

The PIC-to-PIC communication is via a four-wire I^2C bus. The use of this bus reduces the number of signals routed across the HDDA backplane and the drive drawer motherboards. The Neuron chip on the HDDA LON interfaces the PIC to the EM Net.

The HDDA LON board receives 48VDC through the LON port on the HDDA backplane and has an on-board DCC converter (48V to 5V).

Power Shelf LON Board

The Power Shelf LON board plugs into the PDU (rear of the Power Shelf). It contains a PIC that performs the following functions:

- Gathers power supply status
- Performs A-to-D conversion for analog measurements
- Generates control signals

EM Firmware Elements

386EX Code (E-LP board)

The 386EX code implements the master node of the EM Net. After any E-LP reset, the 386 code discovers the physical address of each node (Neuron) on the EM Net, generates logical ID addresses, and sends them to each node for use in subsequent communication. It determines the shelf type along the EM Net ring, and captures periodic status messages sent from the node in each shelf. The hardware supports field firmware download updates of the 386 code.

The 386 code also implements the EM Console. The console is communicated to the ECHaP serial port through a hardware FIFO on the E-LP board and a software driver running on the ECHaP.

Student notes

Neuron

The *Neuron* is the interface on the LON portion of the EM Net (refer to the *LON Topology* section earlier in this chapter). A Neuron is used on the EM/LCD, HDDA LON, Power Shelf LON, and E-LP boards. The Neuron on the E-LP board is the EM Net master, while the Neurons on the other boards are slaves.

Status and control values for each shelf type are grouped into one or more *network variables*. Each slave Neuron receives updated network variable values from its local PIC. A slave Neuron propagates the current values up to the master upon command from the local PIC, or in response to a request from the master Neuron. Each shelf type (E-Box, HDDA, Power Shelf) requires a unique set of variables, and each shelf requires a different hardware interface between the Neuron and PIC. The Neuron code for each shelf type is therefore unique.

The hardware supports field firmware download updates of all the slave Neurons. The master Neuron on the E-LP has a larger memory and cannot be field-downloaded.

Each Neuron is programmed by the manufacturer with a unique 48-bit number that serves as a physical address.

Programmable Integrated Controller (PIC)

The PICs on the EM/LCD, HDDA LON, and Power Shelf LON boards gather status and generate control signals within each shelf. They each reside on the same board as the Neuron that provides their interface to the EM Net.

The PIC on each board type performs a unique set of functions over a unique set of interfaces, so the programmed PIC on each board type is unique. All the PIC parts use internal, one-time program memory for their code and cannot support field-download updates.

Student notes

Field-Programmable Gate Array (FPGA)

The FPGA on the EM/LCD board interfaces the ECHaP IDE bus to the Front Panel. It also interfaces the EM/LCD PIC to the Front Panel, and contains much of the EM/LCD control *glue* logic.

The FPGA on the E-LP board contains much of the control logic required to support the local Neuron and 386EX microprocessor. It controls a FIFO used by the E-LP to communicate with the EM Console over the ECHaP IDE bus.

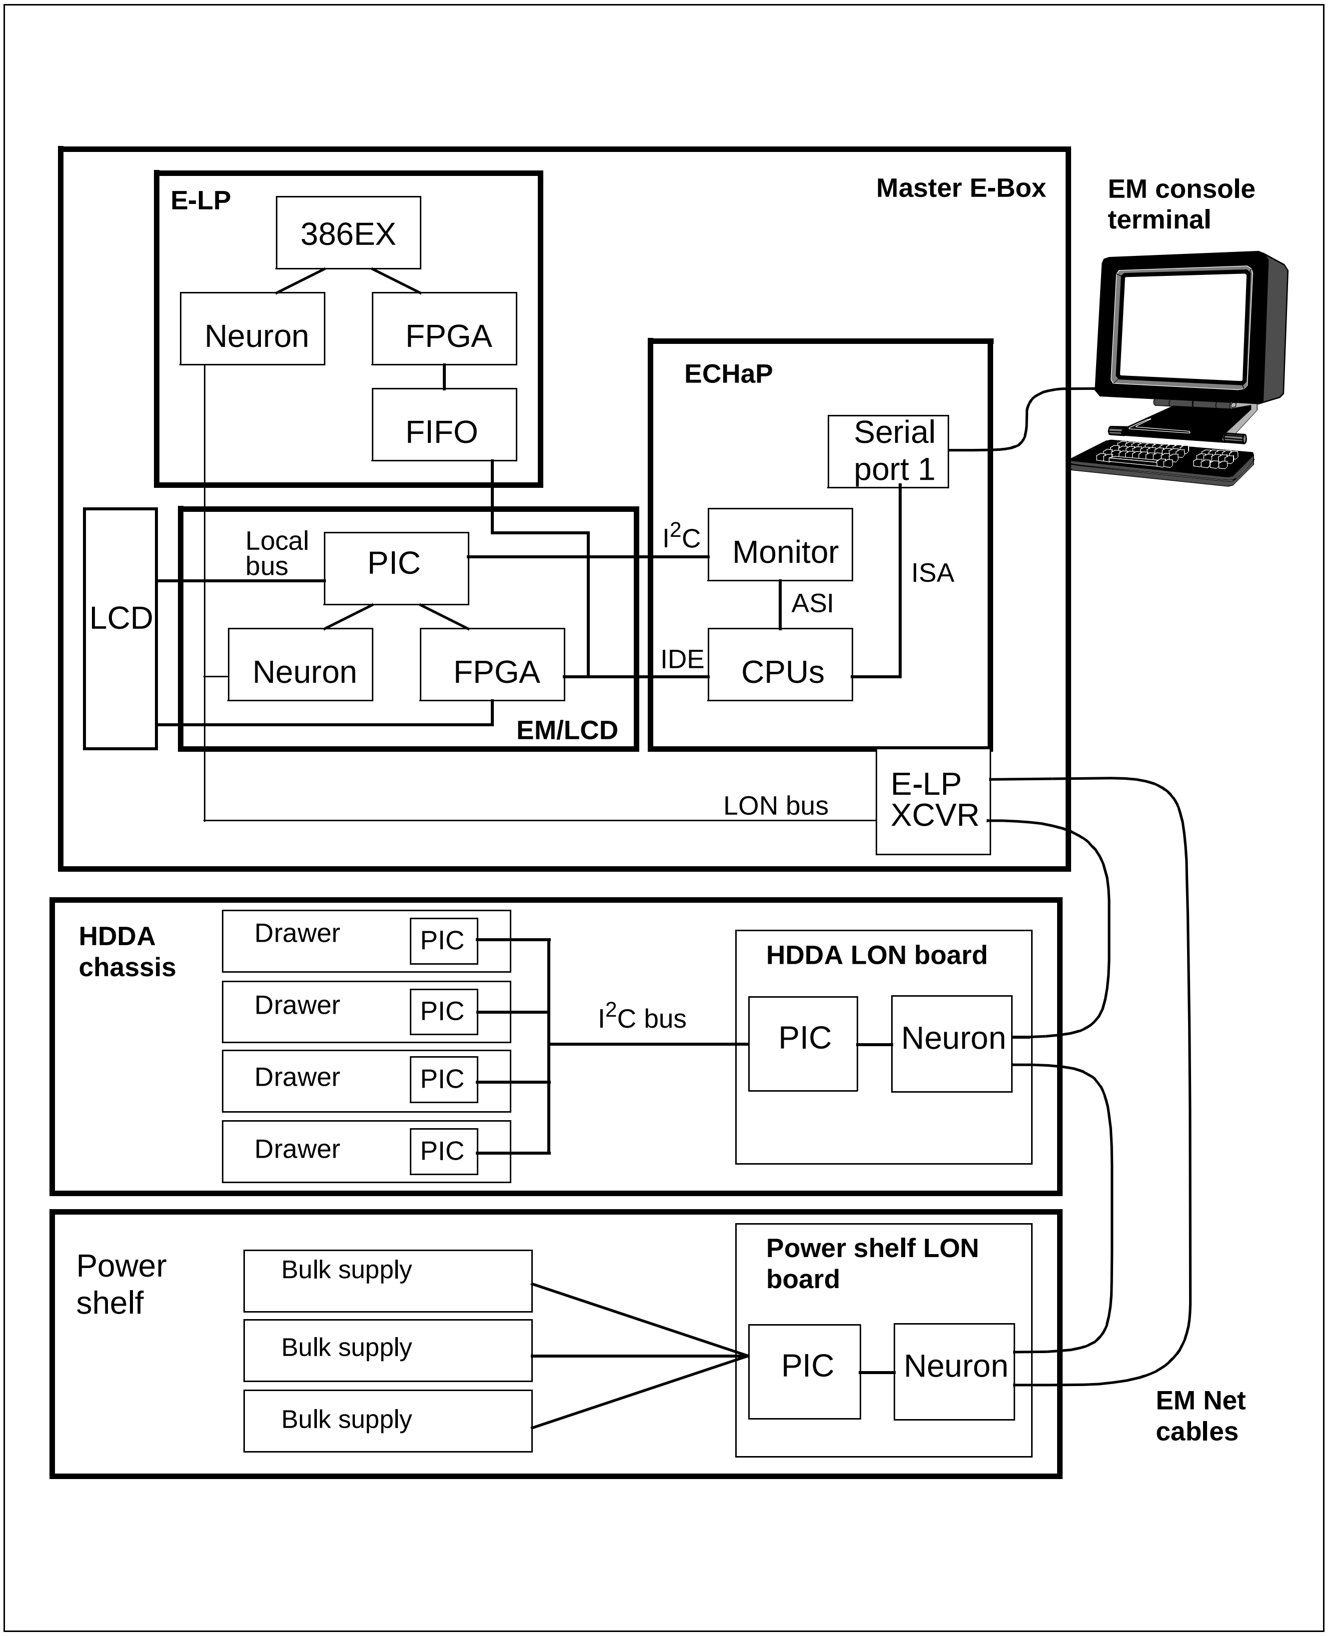

EM subsystem (base node, single HDDA)

EM Subsystem Diagrams and Interface Descriptions

The block diagram EM subsystem shows the major Environmental Monitor components and interfaces in a single-HDDA base node. The diagrams appearing later in this chapter show the E-LP, EM/LCD, and HDDA nodes in greater detail.

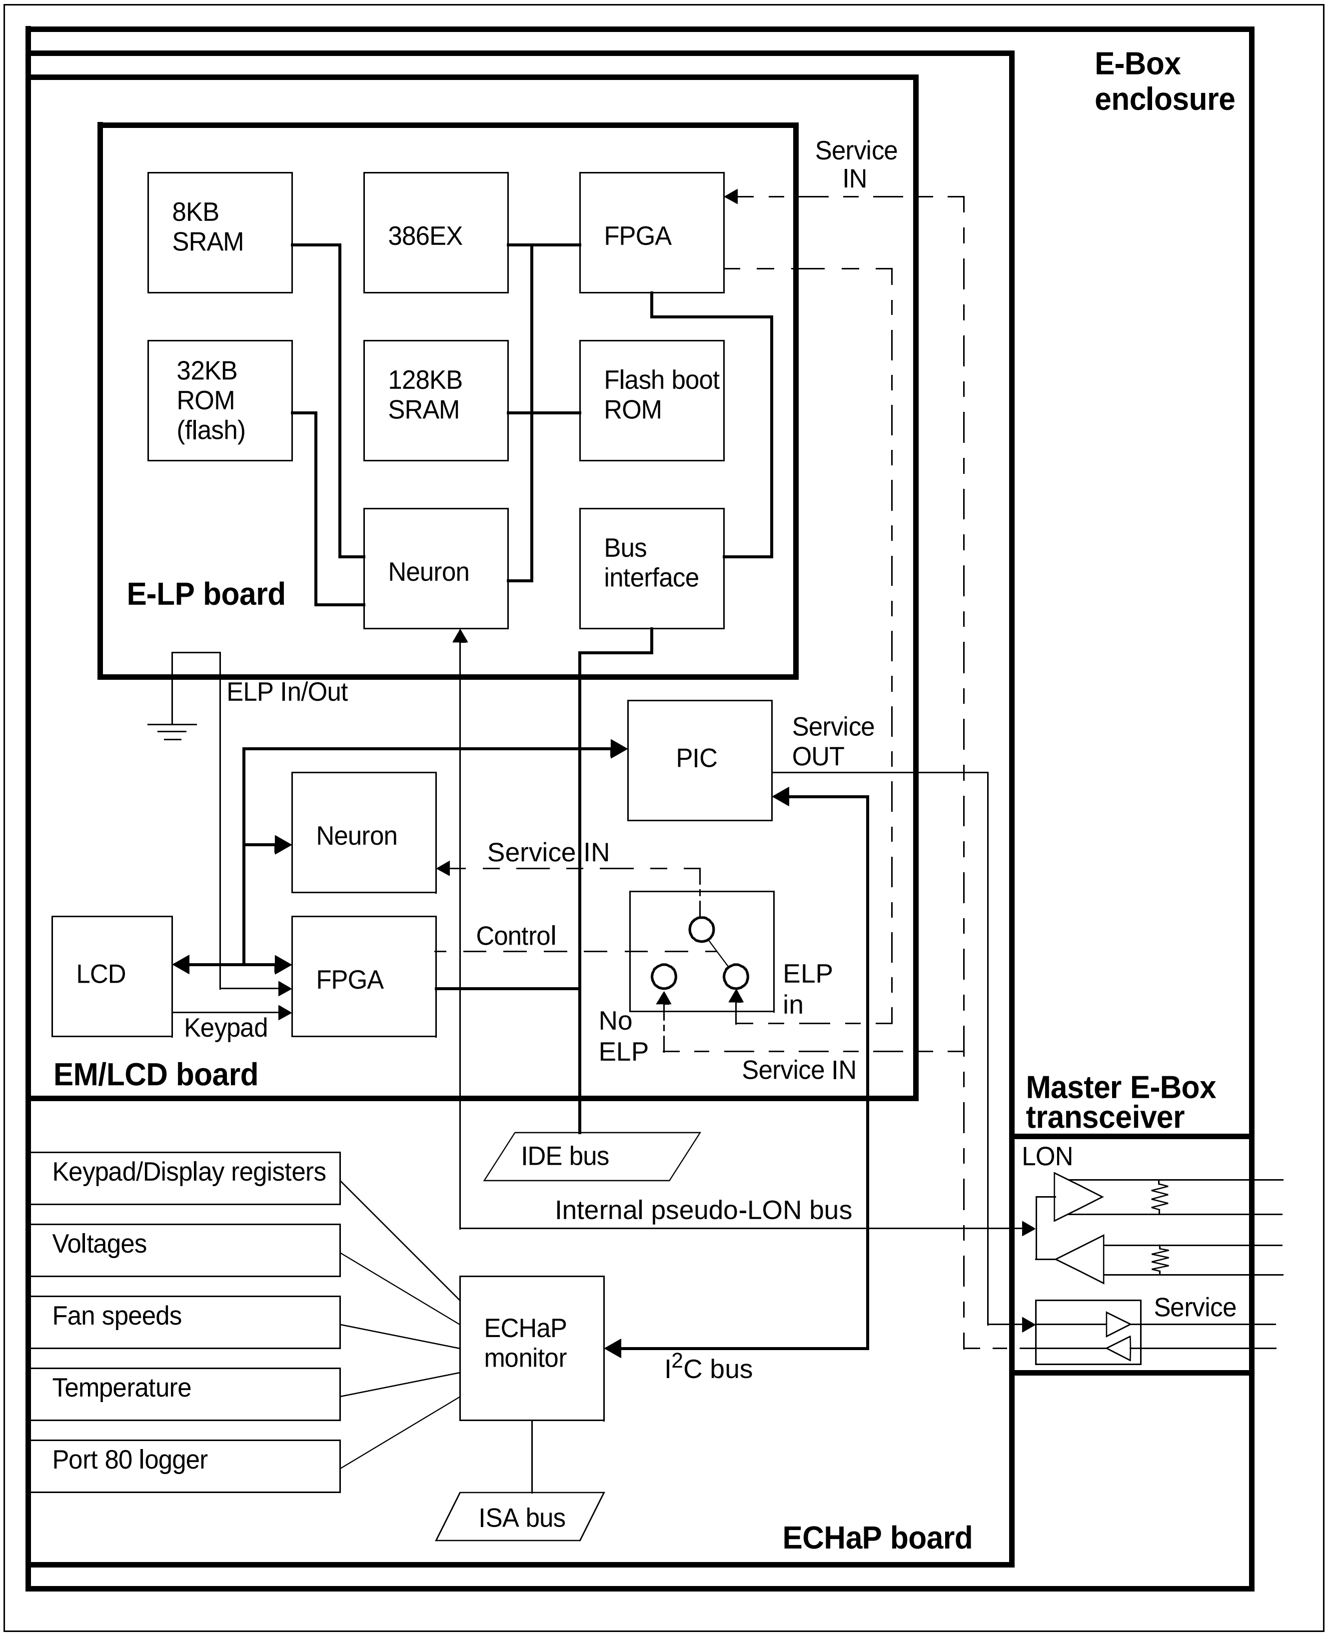

EM/LCD and E-LP nodes

Node and Interface Descriptions

E-LP Board and the ECHaP Interface

The E-Box LON Processor (E-LP) is a *smart card* with an embedded 386EX processor. that controls the EM subsystem. It is the master node of the EM Net, therefore, there is only one E-LP board per M2000. Physically, the E-LP is assembled as a daughter card onto the EM/LCD board in the master E-Box. The E-LP board does not perform any environmental monitoring, but collects data from the slave nodes and communicates with the CPUs in the master E-Box.

As with the EM/LCD board, the E-LP connects to the ECHaP with a pseudo IDE interface implemented by a Field Programmable Gate Array (FPGA).

Transmissions between the E-LP board and master E-Box CPUs occur asynchronously in each direction. Each transmission consists of 8-byte packets which are assembled into a complete message. One 8-byte packet is completely transferred before the next one can be started (in a given direction). This provides data stream flow control. Each packet header contains positional information about the message, but retries are not supported.

Messages are designed so that certain data from the nodes are spontaneously delivered to the CPUs. The master E-Box CPUs can also query the latest value received by the E-LP.

EM/LCD Node and LON Interface

The Environmental Monitor Liquid Crystal Display (EM/LCD) assembly provides a LON node for extracting ECHaP environmental information, controlling the ECHaP power supplies, and controlling the LCD. The EM/LCD connects to the ECHaP via a *pseudo* IDE interface (upper eight bits of the 16-bit IDE interface).

On the EM/LCD board, the Neuron chip is configured so the lower eight pins of the 11-pin I/O port function as an 8-bit bidirectional data bus. The

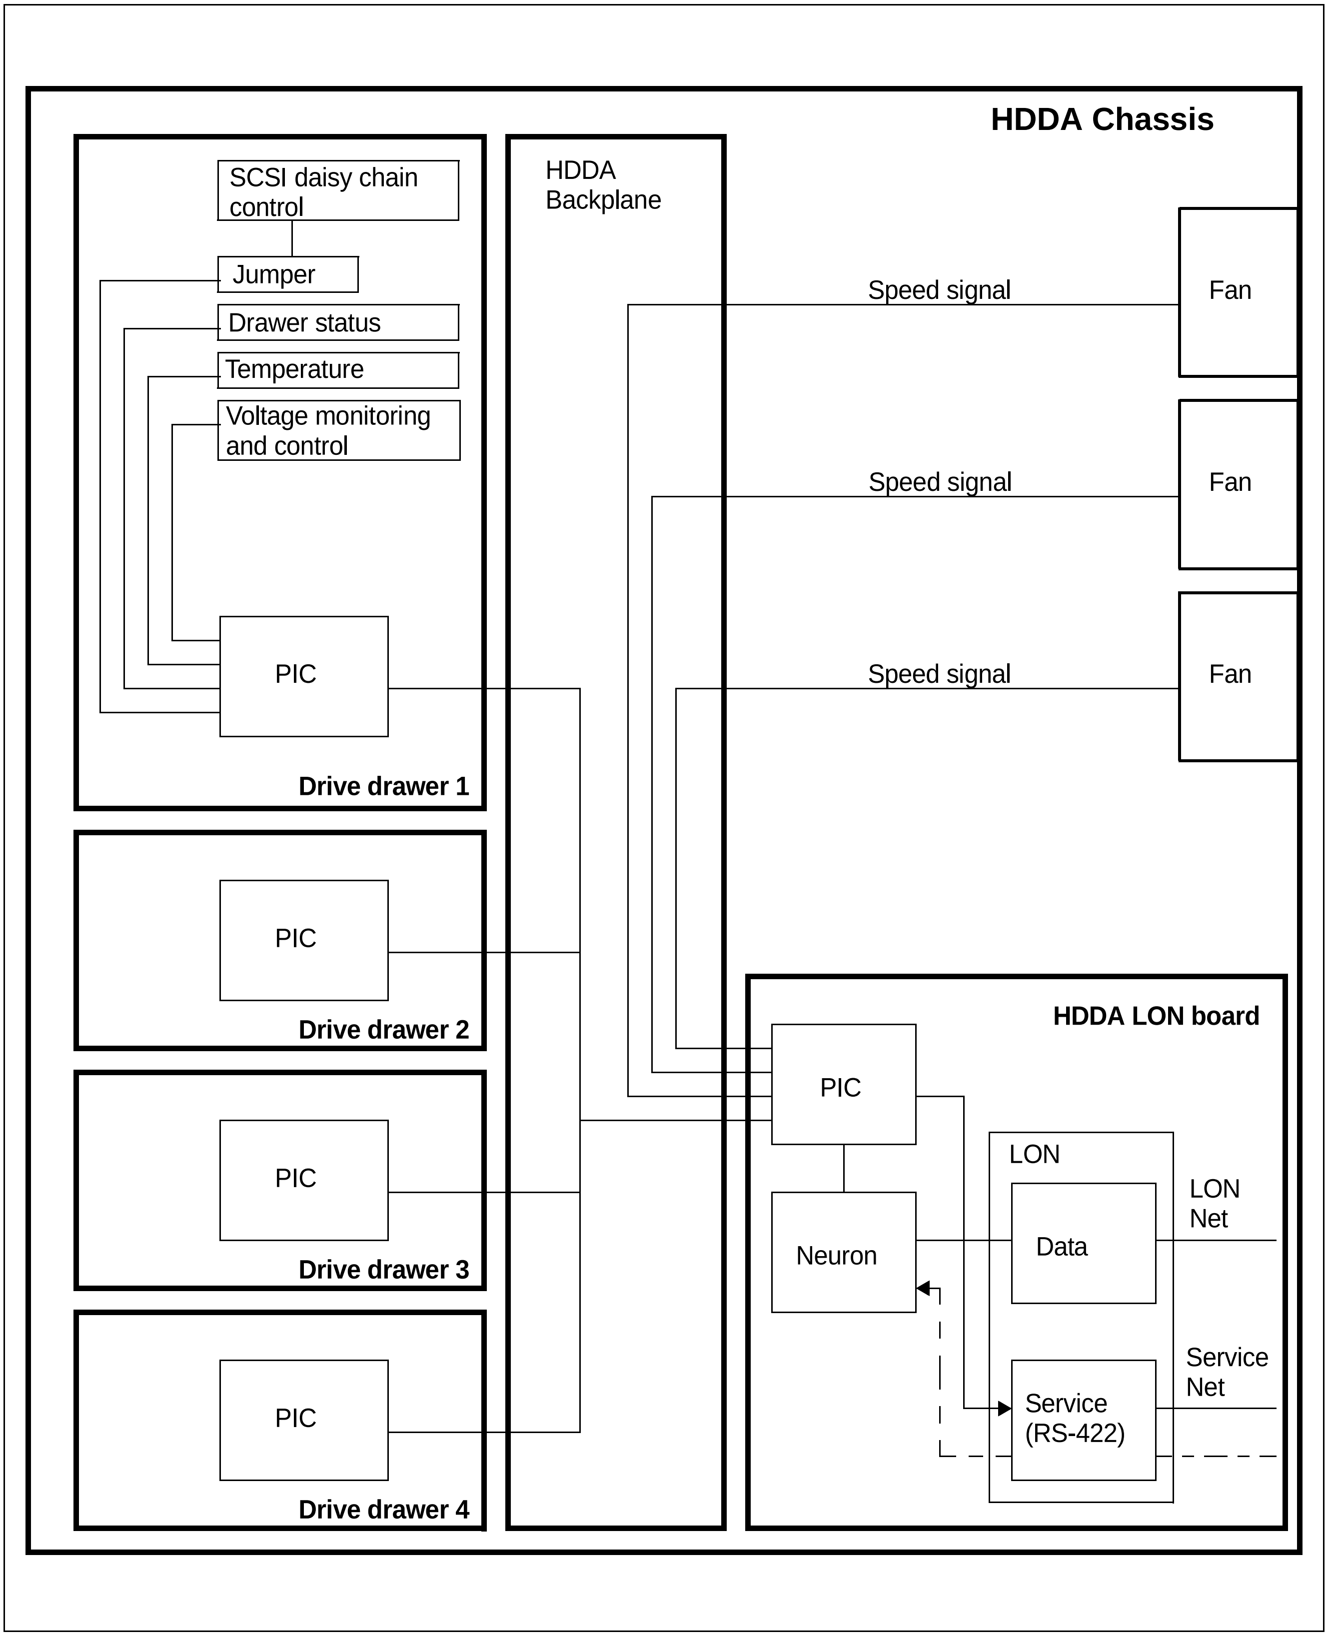
HDDA Chassis
SCSI daisy chain control
Jumper
Drawer status
Temperature
Voltage monitoring and control
PIC
Drive drawer 1
PIC
Drive drawer 2
PIC
Drive drawer 3
PIC
Drive drawer 4
HDDA Backplane
Speed signal
Fan
Speed signal
Fan
Speed signal
Fan
HDDA LON board
PIC
Neuron
LON
Data
Service (RS-422)
LON Net
Service Net

HDDA node

remaining 3 pins function as the chip select, the register select, and the I/O port direction control.

The Neuron chip appears as two registers to the PIC: an 8-bit read/write register, and a 1-bit, read-only status register. The 1-bit status register contains the handshaking bit which allows the PIC (master) to monitor the Neuron (slave) between every byte transfer, ensuring that both devices are ready for the byte to be transferred. The PIC must read the handshake bit before every read or write operation.

To eliminate the possibility of bus contention, a token-passing scheme is implemented in the Neuron chip. The token is owned by the PIC, which is the master after synchronization, and is passed between the PIC and the Neuron. After each data transfer is completed, the token owner writes an end-of-message (EOM) to indicate the transfer is complete.

Interpretation of the handshake bit depends on current token ownership. If the PIC owns the token, it waits for a low handshake level from the Neuron before writing data to the bus. If the Neuron owns the token, the PIC waits for a low transition of the handshake bit before reading the bus.

HDDA Node and LON Interface

Each HDDA drive drawer PIC gathers environmental data and communicates it to another PIC on the HDDA LON board. The PIC-to-PIC communication path is an Inter-Integrated Circuit bus (I^2C bus) that runs through the umbilical cable in the lower pan of the HDDA drive drawer assembly and connects to the HDDA backplane. From there, the I^2C bus continues on circuit traces to the LON port that supports the HDDA LON board.

On the HDDA LON, a *Neuron* chip is connected to the PIC. Each Neuron represents a logical node on the Environmental Monitor Network (EM Net). The HDDA LON Neuron passes *network variable* data arriving from the E-Box LON Processor board (E-LP board) directly to the PIC. It stores network variable data arriving from the local PIC as *output polled only* network variables. A packet is sent to the E-LP board, the EM Net master node, only when the local PIC issues a *propagate* command to the Neuron for a particular variable. This approach allows

Student notes

the E-LP board to obtain current HDDA data by polling the Neuron's network variables, but only announces HDDA data to the E-LP board when the PIC processor determines that the data has exceeded a reporting or a timing threshold.

The network variables represent status values on a particular node. There are a total of four readable network variables, one for each drawer, containing all the information about that drawer. Also, there are a total of six writable network variables, one for power control, one for setting trip thresholds for all variables on all drawers, and the remaining four variables for controlling each of the four drawers.

EM subsystem FRUs		
Part number	**Part Name**	**Description/Comments**
100158	HDDA LON board	LON node for the HDDA.
100165	Power Shelf LON board	LON node for the Power Supply Chassis.
500749	Master EM/LCD assembly	EM/LCD for the master E-Box (includes the E-LP board).
500866	Slave EM/LCD assembly	EM/LCD for a slave E-Box.
100166	Master E-Box (E-LP) transceiver	Use only with master E-Box.
100157	Slave E-box (ECHaP) transceiver	Use only with slave E-Box.
450330	Cable assembly, RJ45, 2 ft.	EM Net cable (daisy-chain).
450370	Cable assembly, RJ45, 14 ft.	EM Net cable (ring).

11 Exercise: Processing, Storage, Power, and EM Subsystems

Introduction

This exercise covers a variety of important topics from the last few chapters. It is not a test of what you remember from the lecture; it is practice for the beneficial field work habit of referencing information and subsequently applying it.

Instructions

Answer the following questions independently using the appropriate manual(s). Your instructor will discuss the answers in detail and ensure you are comfortable with the information before proceeding to new material. Good luck and have fun with the exercise.

1. What Host main memory configurations are supported? What are the Host module PCI card slot assignments?

2. How is SCI implemented for the M2000? Which modules are interconnected by SCI cables?

3. How many expansion slots are on the ECHaP primary PCI bus? On the secondary bus? What are the PCI card slot assignments?

4. What is the key function of the AFX8000? Which PCI slot does it occupy?

5. What is accomplished by inserting a jumper on jumper pad J4 on the HDDA drive drawer motherboard?

6. Where is the HDDA LON board physically located? What are its key functions?

7. Differentiate between point-to-point and daisy-chain SCSI cables.

8. Fill in the blanks with the appropriate power subsystem component.

   The _____________________________ provide 48 VDC to the Host, E-box, and HDDA.

   The _____________________________ regulate ECHaP CPU voltage.

   The _____________________________ converts the 48-volt input into the required operating voltages for the E-Box.

   The _____________________________ are located in each HDDA drive drawer.

9. The Power Shelves must be made independently redundant. Explain why this is true or false.

10. Distinguish between the LON and Service Network.

12 Laboratory Exercise

Introduction

This lab exercise allows you to apply classroom learning by performing a series of procedures on an M2000.

General Instructions

Unless otherwise directed by your instructor, perform the steps in the given sequence and complete the entire exercise. Fill in information and answer questions where called for. Refer to the *Auspex NetServer M2000 Field Service Guide* and any other documents as required to perform the exercise. Contact your instructor with any concerns or questions. Good luck and have fun with the exercise.

Objectives

▲ Identify, remove, and replace FRUs that constitute the Processor, Storage, Power, and Environmental Monitor subsystems.

Item: Power supply

Part No.:

Serial No.:

Host front panel

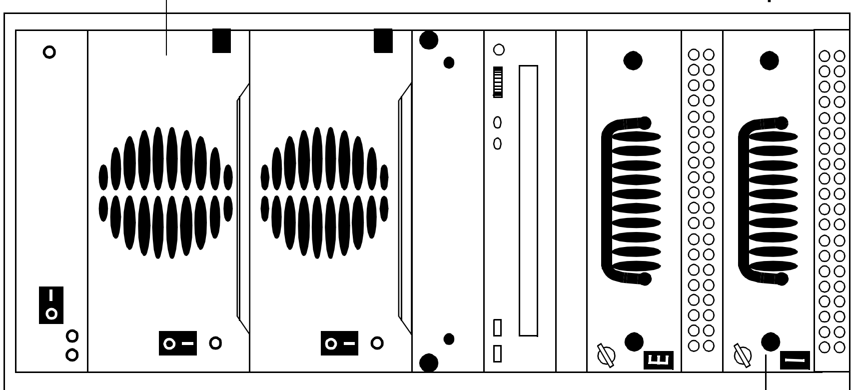

Item: Boot disk drive

Part No.:

Serial No.:

Section 1: Host FRUs

Front Panel FRUs

1. Referring to the figure on the facing page, remove each selected FRU per instructions in *Auspex NetServer M2000 Field Service Guide*.
2. Complete the requested information on the facing page, and reinstall the FRU. The CD-ROM is removed and replaced in the the subsection *Host FRUs: Internal FRUs.*
3. Continue with Lab Section 1, or follow the procedure outlined by your instructor.

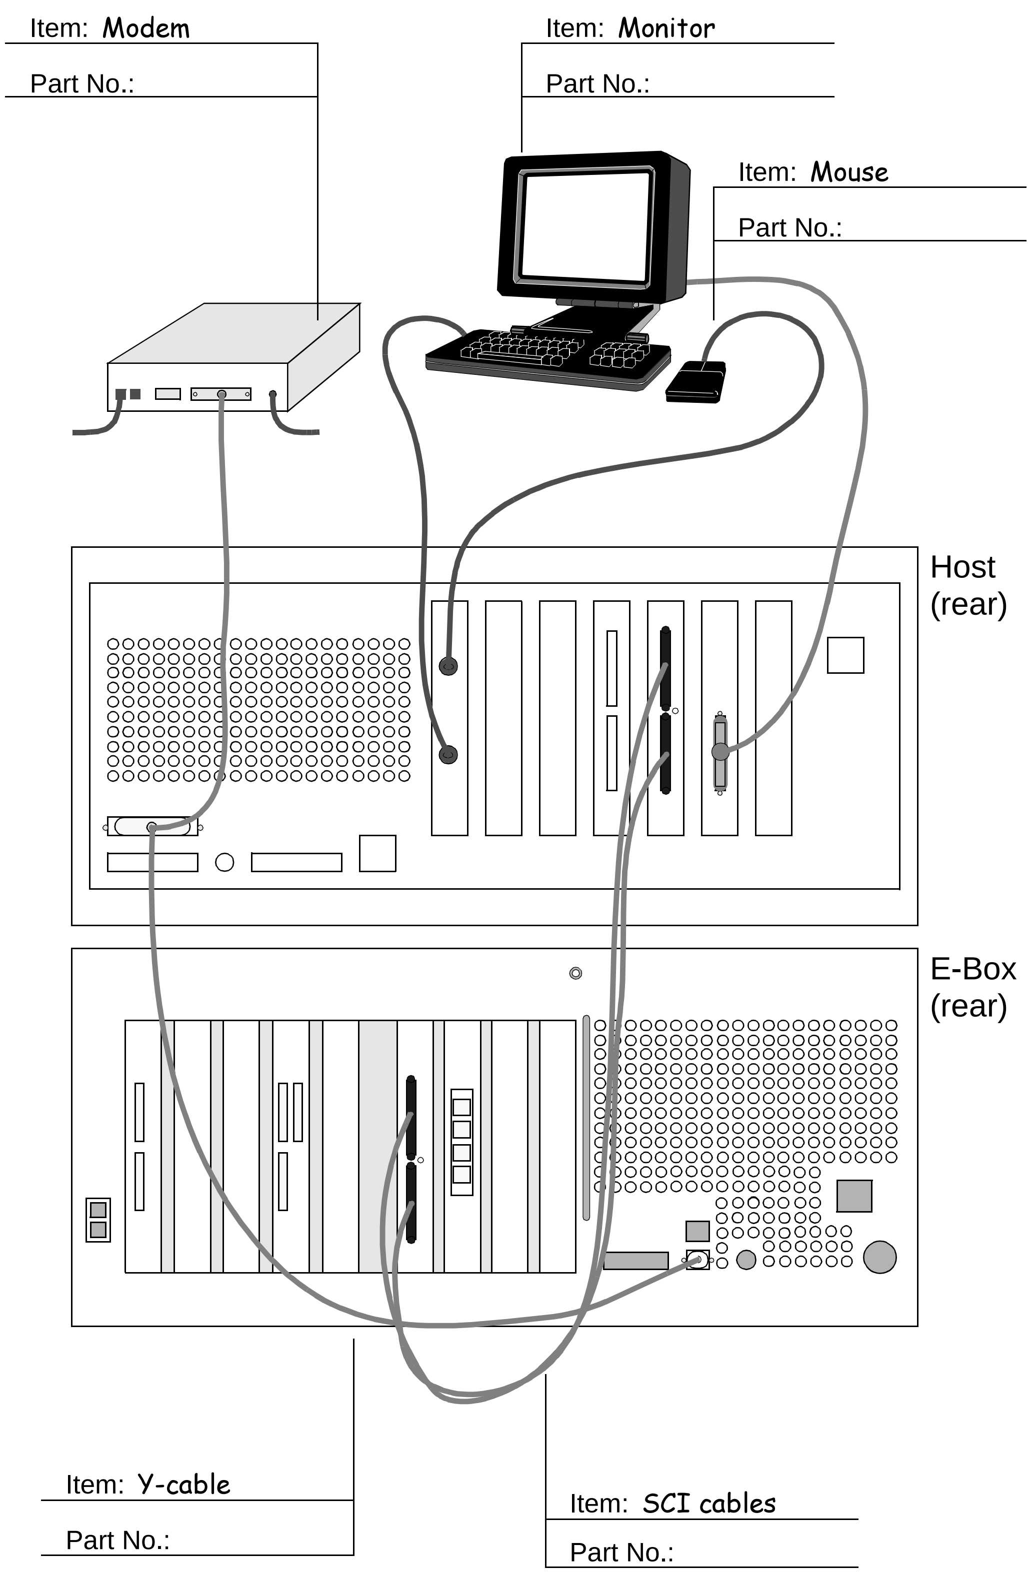
Item: Modem
Part No.:
Item: Monitor
Part No.:
Item: Mouse
Part No.:
Host
(rear)
E-Box
(rear)
Item: Y-cable
Part No.:
Item: SCI cables
Part No.:

Section 1: Host FRUs (continued)

Rear Panel FRUs

1. Referring to the figure on the facing page, remove each selected FRU per instructions in *Auspex NetServer M2000 Field Service Guide*.
2. Complete the requested information on the facing page, and reinstall the FRU.
3. Continue with Lab Section 1, or follow the procedure outlined by your instructor.

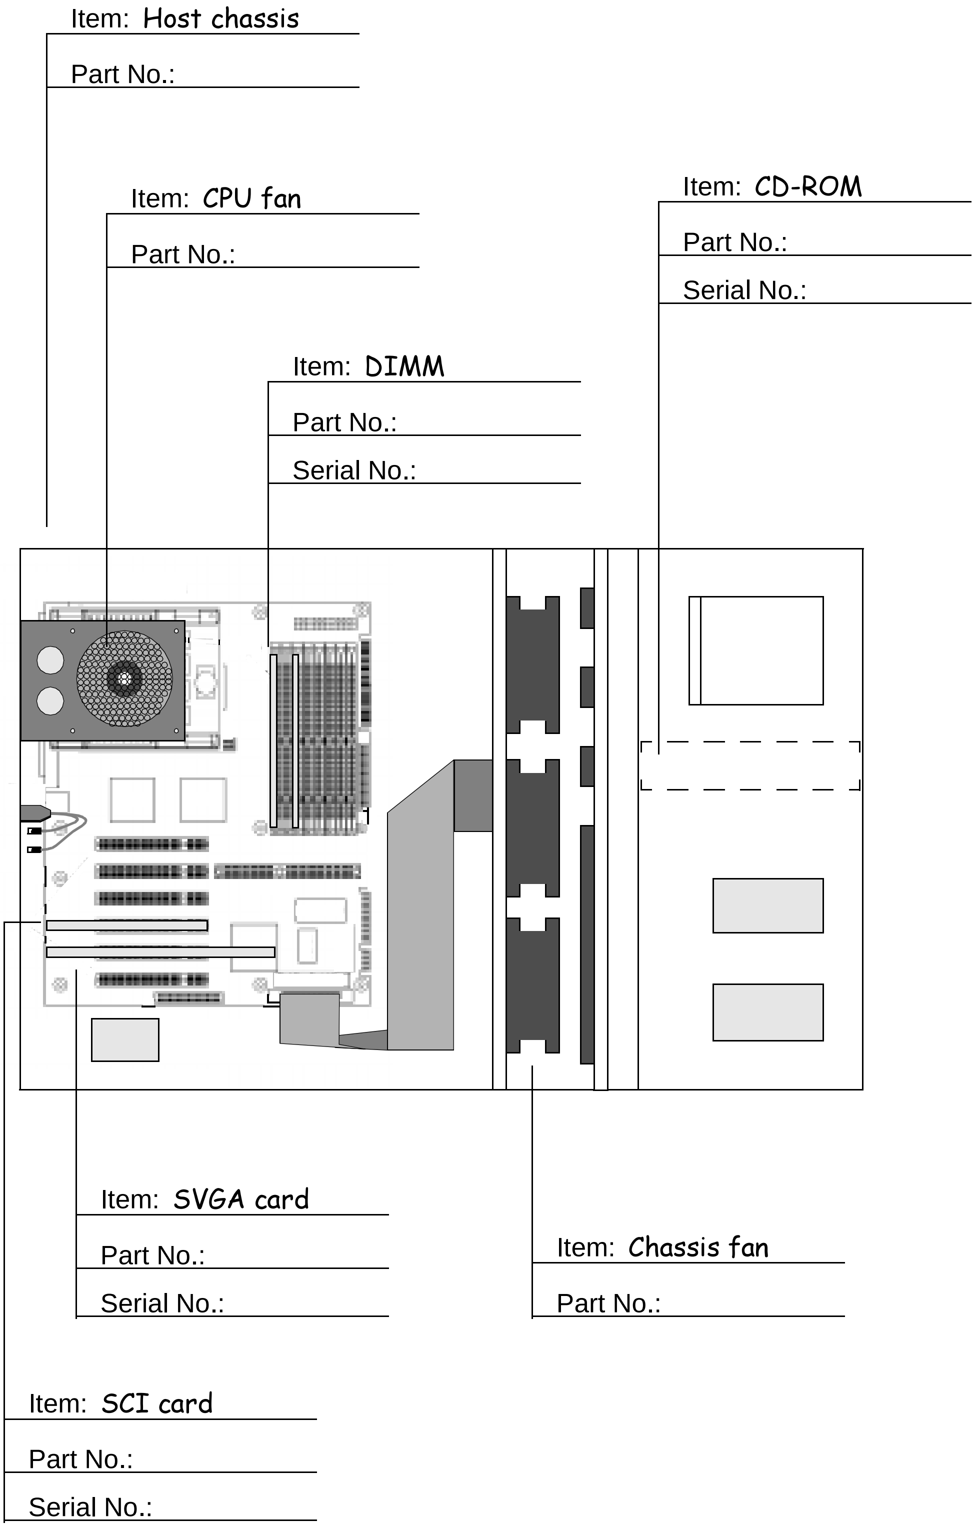
Item: Host chassis
Part No.:
Item: CPU fan
Part No.:
Item: CD-ROM
Part No.:
Serial No.:
Item: DIMM
Part No.:
Serial No.:
Item: SVGA card
Part No.:
Serial No.:
Item: Chassis fan
Part No.:
Item: SCI card
Part No.:
Serial No.:

Section 1: Host FRUs (continued)

Internal FRUs

1. Referring to the figure on the facing page, remove each selected component per instructions in *Auspex NetServer M2000 Field Service Guide*.
2. Complete the requested information on the facing page, and reinstall the FRU.
3. Move to the next section, or follow the procedure outlined by your instructor.

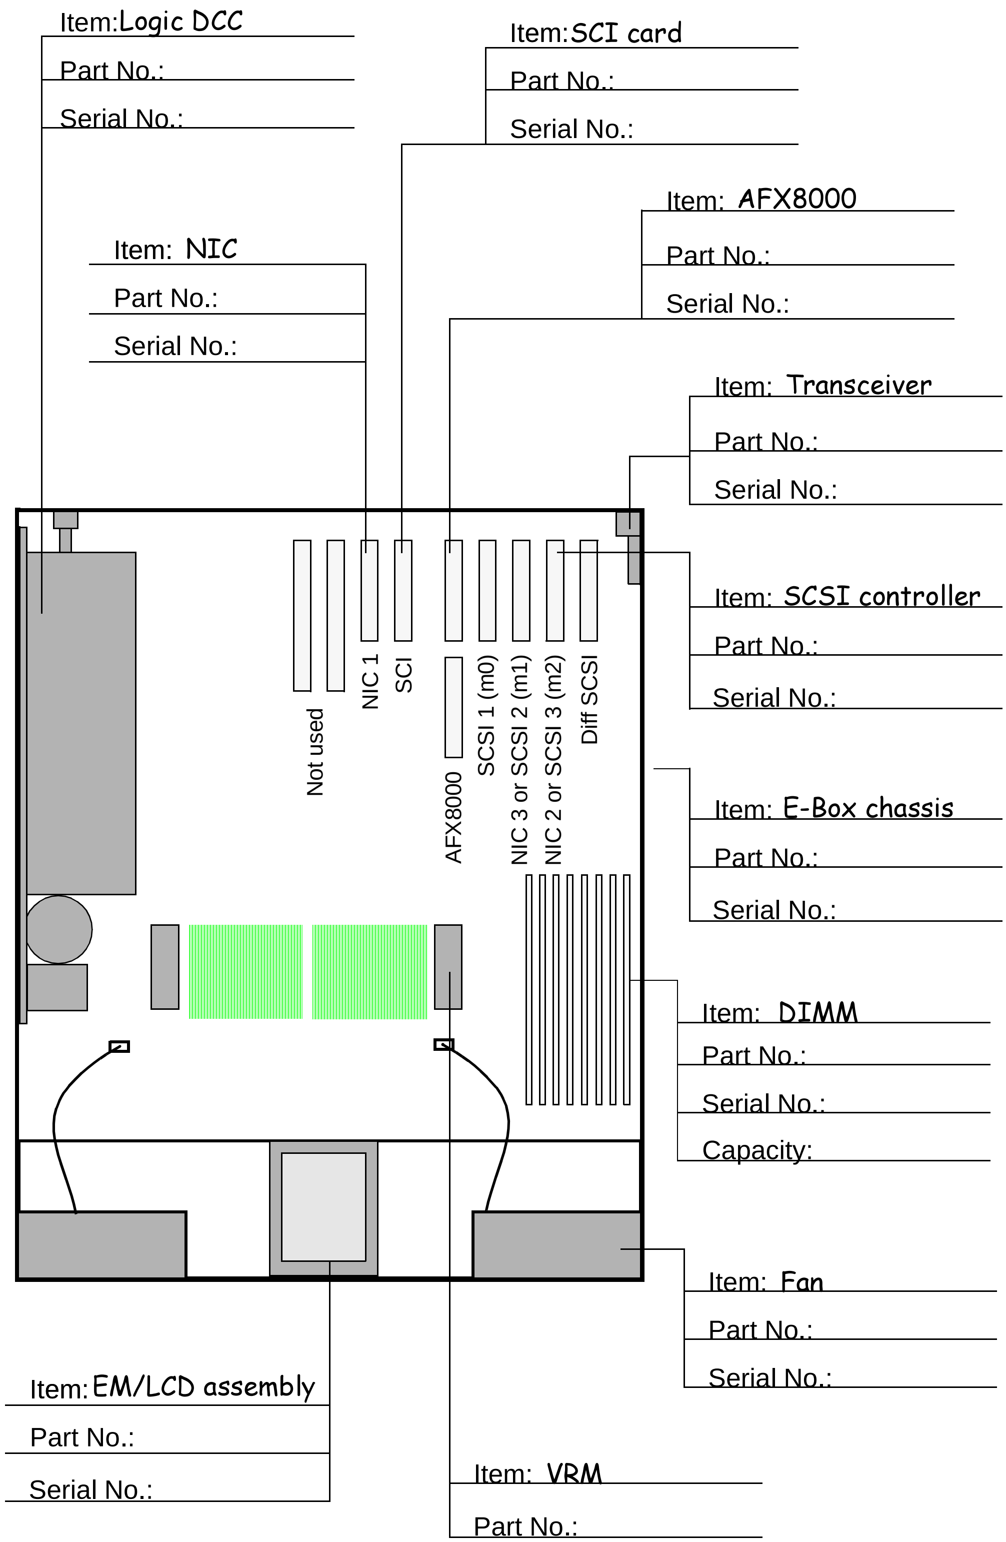
Item: Logic DCC
Part No.:
Serial No.:
Item: SCI card
Part No.:
Serial No.:
Item: AFX8000
Part No.:
Serial No.:
Item: NIC
Part No.:
Serial No.:
Item: Transceiver
Part No.:
Serial No.:
Not used
NIC 1
SCI
AFX8000
SCSI 1 (m0)
NIC 3 or SCSI 2 (m1)
NIC 2 or SCSI 3 (m2)
Diff SCSI
Item: SCSI controller
Part No.:
Serial No.:
Item: E-Box chassis
Part No.:
Serial No.:
Item: DIMM
Part No.:
Serial No.:
Capacity:
Item: Fan
Part No.:
Serial No.:
Item: EM/LCD assembly
Part No.:
Serial No.:
Item: VRM
Part No.:

Section 2: E-Box FRUs

Internal FRUs

1. Referring to the figure on the facing page, remove each selected FRU per instructions in *Auspex NetServer M2000 Field Service Guide*. Remove the E-Box chassis from the server prior to removing and reinstalling the internal FRUs.
2. Complete the requested information on the facing page, and reinstall each FRU in the E-Box chassis with the exception of the AFX8000. An NVRAM battery will be removed in the next subsection.
3. Continue with Lab Section 2, or follow the procedure outlined by your instructor.

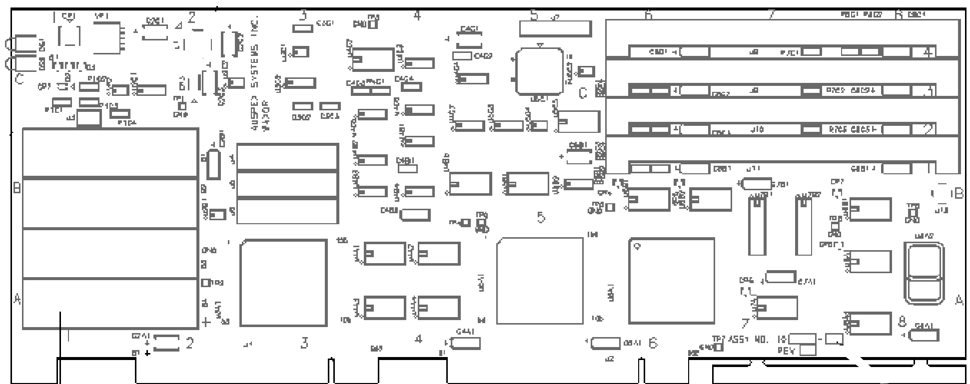

Item: Battery

Part No.:

Section 2: E-Box FRUs (continued)

AFX8000 Battery

1. If the AFX8000 is not already removed from the E-Box chassis, remove it now per instructions in *Auspex NetServer M2000 Field Service Guide*.
2. Referring to the figure on the facing page, remove an AFX8000 battery per instructions in *Auspex NetServer M2000 Field Service Guide*.
3. Complete the requested information on the facing page, reinstall the battery, and reinstall the AFX8000 in the E-Box.
4. Move to the next section, or follow the procedure outlined by your instructor.

Item: HDDA chassis

Part No.:

Serial No.:

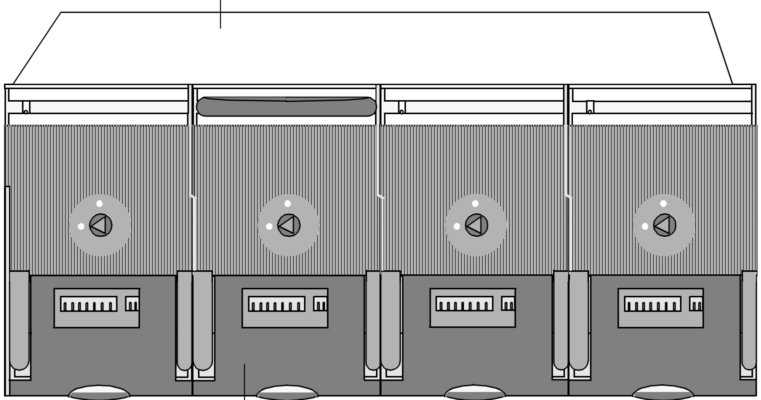

Item: HDDA drive drawer

Part No.:

Serial No.:

Section 3: Storage Subsystem FRUs

External FRUs, Front Panel

1. Referring to the figure on the facing page, remove each selected FRU per instructions in *Auspex NetServer M2000 Field Service Guide*.
2. Complete the requested information on the facing page and reinstall the FRU.
3. Continue with Lab Section 3, or follow the procedure outlined by your instructor.

Item: Fan

Part No.:

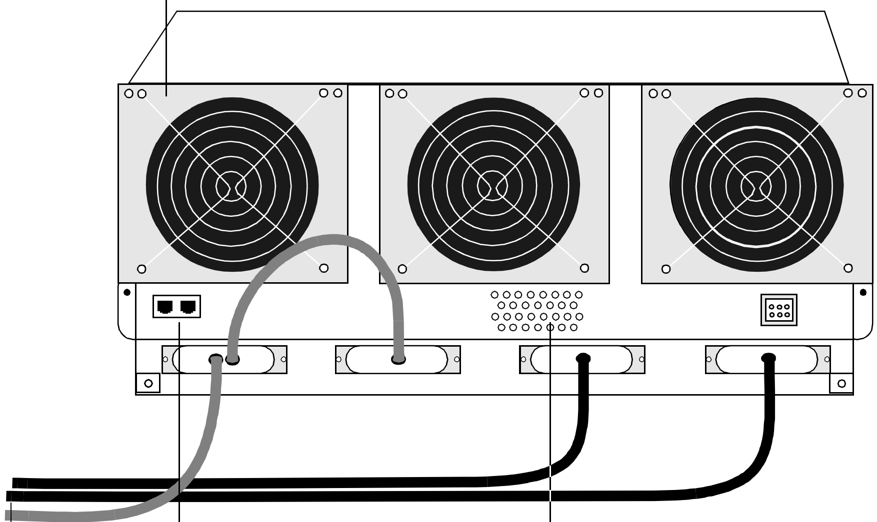

Item: Fuse (behind bracket)

Part No.:

Item: HDDA LON board (behind bracket)

Part No.:

Serial No.:

Item: SCSI cable (daisy-chain)

Part No.:

Item: SCSI cable (point-to-point)

Part No.:

Section 3: Storage Subsystem FRUs (continued)

External FRUs, Rear Panel

1. Referring to the figure on the facing page, remove each selected FRU per instructions in *Auspex NetServer M2000 Field Service Guide*.
2. Complete the requested information on the facing page and reinstall the FRU.
3. Continue with Lab Section 3, or follow the procedure outlined by your instructor.

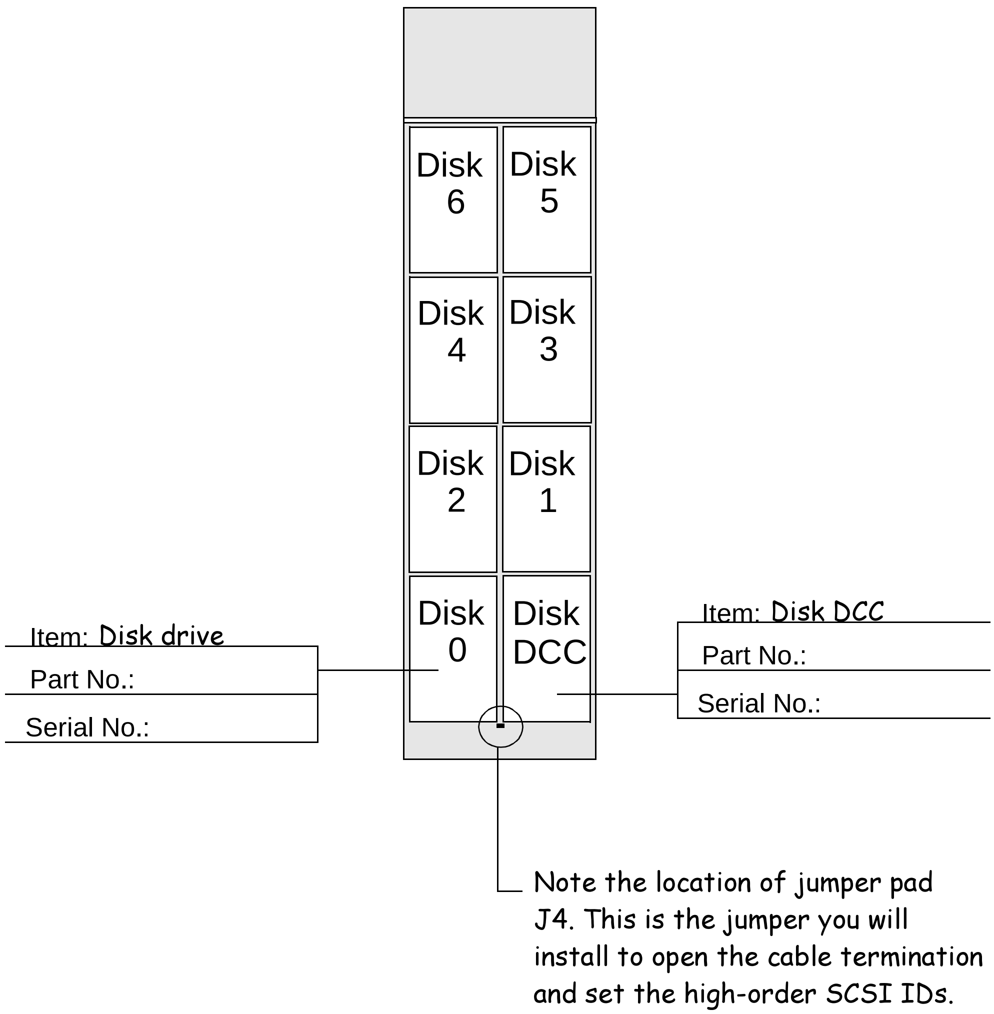
Disk 6
Disk 5
Disk 4
Disk 3
Disk 2
Disk 1
Disk 0
Disk DCC
Item: Disk drive
Part No.:
Serial No.:
Item: Disk DCC
Part No.:
Serial No.:
Note the location of jumper pad J4. This is the jumper you will install to open the cable termination and set the high-order SCSI IDs.

Section 3: Storage Subsystem FRUs (continued)

HDDA Internal FRUs

1. Referring to the figure on the facing page, remove each selected FRU per instructions in *Auspex NetServer M2000 Field Service Guide*.
2. Complete the requested information on the facing page and reinstall the FRU.
3. Move to the next section, or follow the procedure outlined by your instructor.

Item: Power Shelf

Part No.:

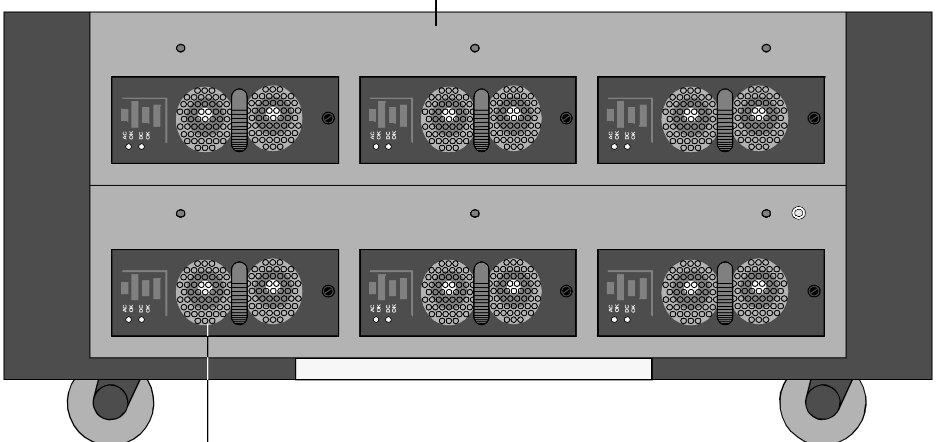

Item: Bulk power supply

Part No.:

Serial No.:

Section 4: Power Subsystem FRUs

Front Panel FRUs

1. Referring to the figure on the facing page, remove each selected FRU per instructions in *Auspex NetServer M2000 Field Service Guide*.
2. Complete the requested information on the facing page and reinstall the FRU.
3. Continue with Lab Section 4, or follow the procedure outlined by your instructor.

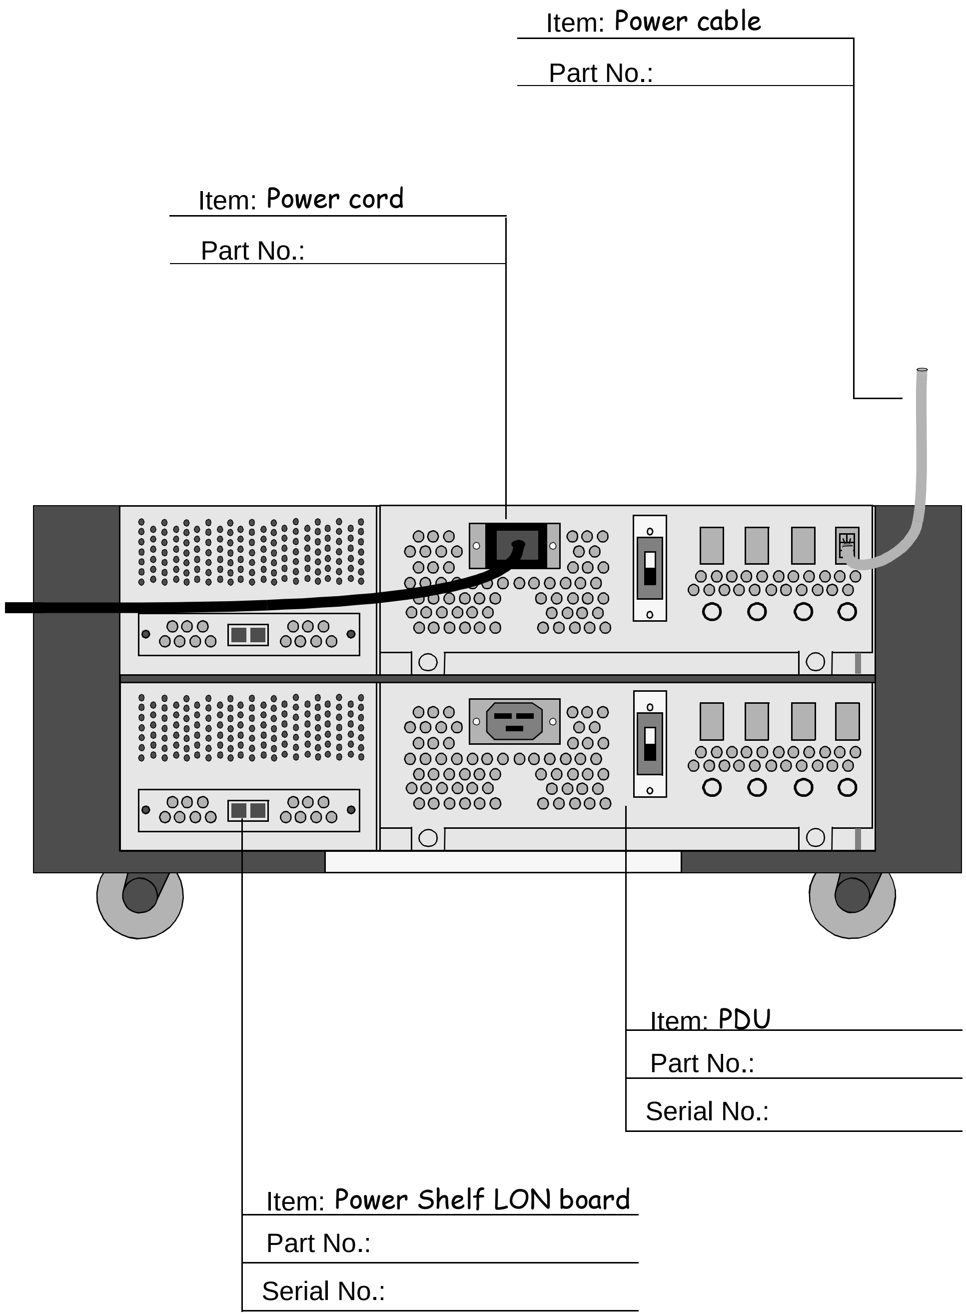

Item: Power cable
Part No.:
Item: Power cord
Part No.:
Item: PDU
Part No.:
Serial No.:
Item: Power Shelf LON board
Part No.:
Serial No.:

Section 4: Power Subsystem FRUs (continued)

Rear Panel FRUs

1. Referring to the figure on the facing page, remove each selected FRU per instructions in *Auspex NetServer M2000 Field Service Guide*.
2. Complete the requested information on the facing page and reinstall the FRU.
3. Move to the next section, or follow the procedure outlined by your instructor.

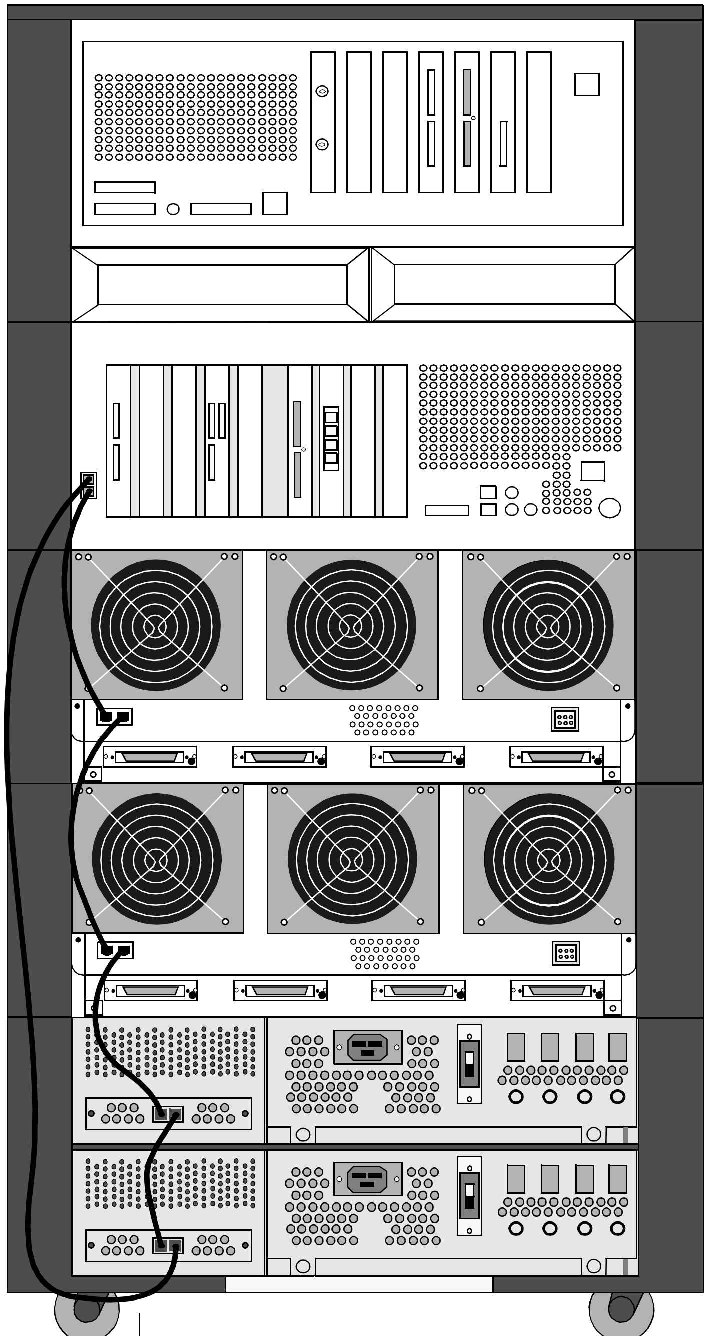

Item: EM Net cable

Part No.:

Section 5: EM Subsystem FRUs

Replacement procedures for other components of the EM subsystem are implemented in the earlier sections of this lab exercise:

- ▲ EM/LCD assembly and Transceiver in Section 2 (E-Box FRUs: Internal FRUs)
- ▲ HDDA LON board in Section 3 (Storage Subsystem FRUs: Rear Panel FRUs)
- ▲ Power Shelf LON board in Section 4 (Storage Subsystem FRUs: Rear Panel FRUs)

1. Remove one of the EM Net cables per *Auspex NetServer M2000 Field Service Guide*.
2. Complete the requested information on the facing page, and reinstall the cable.

This completes the hardware portion of the lab exercise.

Student notes

A Glossary

Introduction

This glossary defines and describes terms used in this manual. Italicized words and terms in the *Description* column are defined within this glossary.

Term	Description
AFX8000	An Auspex-designed, dual-port *PCI* card that serves as a nonvolatile staging buffer for writes to any single disk or array.
Bulk power supplies	The bulk power supplies are installed in the *Power Shelf* and convert AC to 48 VDC for the M2000 modules.
DataGuard	Auspex software product that allows UNIX to halt and reboot without disrupting NFS operations.
DCC	DC-to-DC converter. Converts one DC voltage (input) to another (output). Also see *Disk DCC* and *Logic DCC*.
DIMM	Dual In-line Memory Module. Main memory component of varying density used on the Host and *ECHaP* motherboards.
Disk DCC	DC-to-DC converter used in *HDDA* drive drawers; converts 48 VDC input to 5 VDC and 12 VDC for the disk drives in the drawer.
DRAM	Dynamic Random Access Memory. DRAM is the type of physical memory used in most computers.
E-Box	ECHaP Box. Module containing an *ECHaP* motherboard and other components; configured as a Network, File, and Storage Processor (NFSP).
ECHaP	Euclid Common Hardware Platform. Auspex-designed motherboard that serves as the base component for the *E-Box*.

ECHaP Transceiver	Also called *Slave E-Box Transceiver*. Small transceiver board providing *EM Net* connection. It is used in all *Expansion E-Boxes* and operates as a pass-through device (no terminations).
E-LP board	Environmental LON Processor board. The Environmental Monitor master *node*. Part of the *sled* which goes in the front of a *Master E-Box*.
E-LP Transceiver	Small transceiver board providing *EM Net* connection and termination in the *Master E-Box*. Also called *Master E-Box Transceiver.*
EM	Environmental Monitor or Environmental Monitor subsystem.
EM Console	A dumb terminal or terminal emulator software that provides access to environmental readings of the *E-Box*(es), *HDDA*(s), and *Power Shelf*(s).
EM Net	Environmental Monitor Network. A closed network dedicated to the *EM* subsystem. Connects a single *E-LP board*, with a slave *node* in each instrumented module. Contains a *LON* network for data transfer, and a service network for locating network faults and topology management.

EM/LCD assembly	Assembly of boards and mechanical brackets that mounts in the front of an *E-Box*. Provides a *Front Panel* user interface (LCD display, 3 input keys, 1 reset button). Also gathers status and generates control signals within the *E-Box*. The EM/LCD assembly in the *Master E-Box* has an *E-LP board* mounted to it as a daughter card.
EM/LCD board	Environmental Monitor and Liquid Crystal Display board. *EM node* for the *E-Box*. Gathers status and generates control signals within an *E-Box*. Part of the *EM/LCD* assembly (*sled*) which goes in the front of every *E-Box*.
EMON	ECHaP Monitor. A mini operating system used during boot-up of an *E-Box*. Supports applications such as diagnostics, and *NP/FSP* image download over the *SCI* bus.
Expansion E-Box	*E-Box* in M2000 expansion nodes. Distinguished from the *Master E-Box* in that it does not have an *E-LP board* on the *EM/LCD assembly*, and that it contains an *ECHaP Transceiver* versus an *E-LP Transceiver.* Also called *Slave E-Box*.
FPGA	Field Programmable Gate Array.
Front Panel	The user interface provided on the front of each *E-Box*. Consists of an LCD display, three push button keys, and one recessed button for reset.
FRU	Field Replaceable Unit. A component which may be replaced in the field for repair or upgrade.

FSP	File and Storage Processor. The CPU on the *ECHaP* that performs the file and storage processing.
GTL	Gunning Transceiver Logic. The bus that interconnects the *ECHaP* CPUs and memory components uses GTL.
HDDA	High Density Disk Array. Module containing up to four drawers, each with a *Disk DCC* and up to seven disk drives.
HDDA drawer motherboard	Motherboard in each *HDDA* drawer; contains some *EM* hardware.
HDDA LON board	*LON node* board which plugs into the back of an *HDDA* chassis. Communicates via an *I^2C* bus to *PICs* on each *HDDA drawer motherboard*. Gathers status and generates control signals within the *HDDA*.
I^2C	A bus and protocol standard for low speed serial communication among interconnected ICs.
I2I	ISA-to-ISA. On the ECHaP, the I2I *FPGA* is a device that bridges two *ISA* buses and has a small FIFO for transferring commands and data between the single-board computer (SBC) host and the *ECHaP*. It is unlikely this slot will ever host an SBC, but the functionality is available.
IDE	Integrated Drive Electronics. Bus interface for mass storage devices.

ISA	Industry Standard Architecture. The bus architecture used in the IBM PC/XT and PC/AT. The AT version of the bus is called the AT bus and became a de facto industry standard. Starting in the early 90s, *PCI* local bus began to replace the *ISA* bus architecture. The *ECHaP* includes an *ISA* bus for slower devices and two *PCI* buses for devices that need better performance.
JTAG	Joint Test Action Group. An industry standard serial interface used as a test and programming interface for in-system programmable logic devices.
LCD board	Contains mounting for LCD display and push buttons for the front panel interface. Part of the *EM/LCD assembly* (*sled*) which goes in the front of every *E-Box*.
Logic DCC	The DC-to-DC converter mounted to the *ECHaP* motherboard that accepts 48 VDC and converts it to the required operating voltages for the *E-Box*.
LON	Local Operating Network. A standards based, robust communications network designed for status and control. Echelon is the company licensing the technology. The LON is part of the *EM Net* which connects EM components throughout the M2000 system.
Master E-Box	*E-Box* in the base *node*. Distinguished from the *Expansion E-Box* in that it has an *E-LP board* on the *EM/LCD assembly*, and that it contains an *E-LP Transceiver* versus an *ECHaP Transceiver.*

Master E-Box Transceiver	See *E-LP Transceiver.*
N+1 redundancy	N+1 redundancy means that for a specific server configuration, the minimum number of power supplies are installed plus one more added for redundant operation. This allows a single power supply to fail and be replaced without disrupting server operation.
Network variable	A construct defined within the *LON* portion of the *EM Net*. Several status and control parameters from a module are grouped into each network variable. Network variables are the communication unit of transfer over the *EM Net*.
Neuron	A microcontroller chip designed to handle handshake protocol of the *LON* network. Used on the *LON* boards (*EM/LCD*, *HDDA*, and *Power Shelf LON boards*).
NFSP	Network, File, and Storage Processor. Another name for the *E-Box*.
NIC	Network Interface Card.
Node	In the context of the M2000 system, a single *E-Box* with associated storage and network connections, *HDDAs*, and *Power Base* that forms the functional stack structure. In the context of the *EM Net* or *LON*, an addressable device on the *LON*.
NP	Network Processor. The CPU on the *ECHaP* that performs network processing.

NVRAM	Nonvolatile Random Access Memory. A type of memory that retains its contents when power is turned off. Typically, NVRAM is SRAM that is made nonvolatile by connecting it to a constant power source such as a battery
PCI	Peripheral Component Interconnect. A local bus standard developed by Intel Corporation. Most modern PCs include a PCI bus in addition to a more general *ISA* expansion bus.
PDU	Power Distribution Unit. The M2000 PDU plugs into the *Power Shelf* and directs the AC input voltage to the *bulk power supplies*.
PIC	Programmable Integrated Controller. A single chip microcontroller optimized for analog and digital signal manipulation. Used on the *LON* boards.
Port 80	A standard address in the address space of most PCs that is used for posting boot-up and diagnostic messages.
Power Base	Crate structure at the base of the M2000 that supports one or two *Power Shelves*.
Power Shelf	Module containing up to three *bulk power supplies*, and a *PDU*. One or two Power Shelves are assembled into the *Power Base* of the M2000.
Power Shelf LON board	*LON node* board which plugs into the back of the *Power Shelf*. Gathers status and generates control signals within the *Power Shelf*.

RAID	Redundant Array of Independent Disks. A disk drive strategy that employs multiple disks in combination for performance and fault tolerance. RAID is implemented in different forms called **RAID levels**.
SCI	Scalable Coherent Interface. SCI is an distributed interconnection standard that provides the expected computer bus services but avoids traditional bus limitations by using point-to-point links and a packet protocol.
SCSI	Small Computer System Interface. A parallel interface standard used by various PCs, and many UNIX systems for attaching peripheral devices.
Slave E-Box	See *Expansion E-Box*.
Slave E-Box Transceiver	See *ECHaP Transceiver.*
Sled	See *EM/LCD assembly.*
SVGA	Super Video Graphics Adapter. A set of graphics standards designed to offer greater resolution than VGA (video graphics array).
VHDCI	Very High Density Cable Interconnect.
VRM	Voltage Regulator Module. On the *ECHaP*, the VRMs regulate the input voltage to the CPUs.
XCVR	Abbreviation for Transceiver.

Student notes

B Course Evaluation

Evaluation Goals

The Auspex Customer Services Department provides global technical support and dispatch service for domestic service calls. As such, we constantly receive feedback from our customers, distributors and service partners. This feedback helps us shape our product and service. Please feel free in providing any suggestions to help us improve our training. We want you to be well prepared to deal with the challenge of field service.

Instructions

Please complete the course evaluation form that follows and submit it to your instructor. Thank you for participating in the evaluation and the course.

Final thanks and reminders:

Thank you for participating in the course and the evaluation process. Your feedback, now and in the future, is needed and appreciated.

Make sure you take the appropriate documentation on every service call. This may include:

- Field Service Guide
- Field Service Guide Supplement(s)
- Storage Peripherals Manager’s Guide or Peripheral Devices Manager’s Guide
- Hardware Manual or Installation Guide
- FRU Information documents
- FCO Rework Instructions

When in doubt, call Auspex Technical Support at 1-800-3AUSPEX.

Field Engineer Technical Education: Course Evaluation

Course: FE Technical Education Supplement, M2000

Student name (optional): ____________________

Date: ____________________

	Excellent	Good	Average	Fair	Poor
Overall course quality					
Instructor name: ____________________					
Subject knowledge					
Presentation skills					
Responsiveness to questions					
Pace of class					
Course content					
Lecture					
Classroom exercises					
Laboratory exercises					
Documentation					
Training manual					
Field Service Guide					
Education Facility					
Provided all needed services					
Comfortable					

NCR Customer Engineers and Auspex field support personnel: please continue the evaluation on the reverse side

Are you comfortable with your knowledge of the M2000?	Yes	No
Comments:		

Do you feel prepared to replace FRUs in the M2000?	Yes	No
Comments:		

How can we change this course to help you become more effective?

Thank you for your feedback.

C New Products

Introduction

The material in this appendix describes M2000 hardware and software products released after the publication of this manual. Typically, the product descriptions are not detailed; they are intended only to keep students apprised of the changes in Auspex's product line.

This page intentionally left blank.

D Errata

Introduction

This appendix lists and corrects errors discovered in the text of the manual. The errors listed are technical rather than typographical.

This page intentionally left blank.