hp AlphaServer ES47/ES80/GS1280

User Information

Version 3.0

This document is intended for those who manage, operate,
or service hp AlphaServer ES47, ES80 and GS1280 systems.

Legal notices and Regulatory notices

June 2004

Regulatory Notices for Series ET2003

FCC Notice

Part 15 of the Federal Communications Commission (FCC) Rules and Regulations has established Radio Frequency (RF) emission limits to provide an interference-free radio frequency spectrum. Many electronic devices, including computers, generate RF energy incidental to their intended function and are, therefore, covered by these rules. These rules place computers and related peripheral devices into two classes, A and B, depending upon their intended installation. Class A devices are those that may reasonably be expected to be installed in a business or commercial environment. Class B devices are those that may reasonably be expected to be installed in a residential environment (i.e., personal computers). The FCC requires devices in both classes to bear a label indicating the interference potential of the device as well as additional operating instructions for the user. The rating label on the device shows which class (A or B) the equipment falls into. Class B devices have an FCC ID on the label. Class A devices do not have an FCC ID on the label. Once the class of the device is determined, refer to the following corresponding statement.

Class A Equipment

This equipment has been tested and found to comply 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 instructions, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at personal expense.

Class B Equipment

This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.
- Consult the dealer or an experienced radio or television technician for help.

Modifications

The FCC requires the user to be notified that any changes or modifications made to this device that are not expressly approved by Hewlett-Packard Company may void the user's authority to operate the equipment.

Cables

Connections to this device must be made with shielded cables with metallic RFI/EMI connector hoods in order to maintain compliance with FCC rules and regulations.

Declaration of Conformity for products marked with the FCC logo – United States only

This device complies with Part 15 of the FCC rules. Operation is subject to the following two conditions: (1) This device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

For questions regarding your product, contact:

Compaq Computer Corporation

P. O. Box 692000, Mail Stop 530113

Houston, Texas 77269-2000 or, call

1-800- 652-6672 (1-800-OK COMPAQ)

For questions regarding this FCC declaration, contact:

Compaq Computer Corporation

P. O. Box 692000, Mail Stop 510101

Houston, Texas 77269-2000 or, call (281) 514-3333

To identify this product, refer to the part, series, or model number found on the product.

Taiwanese Notice

警告使用者:

這是甲類的資訊產品，在居住的
環境中使用時，可能會造成射頻
干擾，在這種情況下，使用者會
被要求採取某些適當的對策。

Japanese Notice

この装置は，情報処理等電波障害自主規制協議会（VCCI）基準に基づくクラス A 情報装置です。この装置を家庭環境で使用すると電波障害を引き起こすことがあります。この場合には，使用者が適切な対策を講じるよう要求されることがあります。

Canadian Notice **(Avis Canadien)**

Class A Equipment

This Class A digital apparatus meets all requirements of the Canadian Interference-Causing Equipment Regulations. Cet appareil numérique de la classe A respecte toutes les exigences du Règlement sur le matériel brouilleur du Canada.

Class B Equipment

This Class B digital apparatus meets all requirements of the Canadian Interference-Causing Equipment Regulations.Cet appareil numérique de la classe B respecte toutes les exigences du Règlement sur le matériel brouilleur du Canada.

European Union Notice

Products with the CE marking comply with the EMC Directive (89/336/EEC) and the Low Voltage Directive (73/23/EEC) issued by the Commission of the European Community and if this product has telecommunication functionality, the R&TTE Directive (1999/5/EC).

Compliance with these directives implies conformity to the following European Norms (in parentheses are the equivalent international standards and regulations):

- EN55022 (CISPR 22) – Electromagnetic Interference
- EN55024 (IEC61000-4-2, 3, 4, 5, 6, 8, 11) – Electromagnetic Immunity
- EN61000-3-2 (IEC61000-3-2) – Power Line Harmonics
- EN61000-3-3 (IEC61000-3-3) – Power Line Flicker
- EN60950 (IEC60950) – Product Safety

Introduction

The AlphaServer ES47, ES80, and GS1280 are systems built around HP's Alpha chip technology. The latest version of this chip now includes inter-processor ports, an I/O port, and two memory controllers. With this design, it is possible to build machines without a system bus or switch because processors can communicate directly to other processors in a mesh of processors. An I/O chip with four I/O ports was developed to form the bridge between the CPU and three PCI/PCI-X buses and an AGP bus

These building block components are placed on building block modules which are placed in building block drawers. Two CPUs are placed on a dual processor module. The I/O chip, know as the IO7 chip, is placed on an I/O riser module. The dual-processor module is placed in either a 2P drawer or an 8P drawer. The I/O riser module is place in either a standard I/O drawer or a high-performance I/O drawer.

It is the drawers and modules that make it possible to build systems with from 2 to 128 processors. When 2P drawers are used, a system with up to 8 processors can be built. When 8P drawers are used, a system with up to 64 processors can be built.

NOTE: When you unpack your system, be sure to save and store all shipping brackets, pallets, and packing material. You will need this material to repack the system, if you should decide to relocate it.

ES47/ES80 Systems

ES47/ES80 System Overview

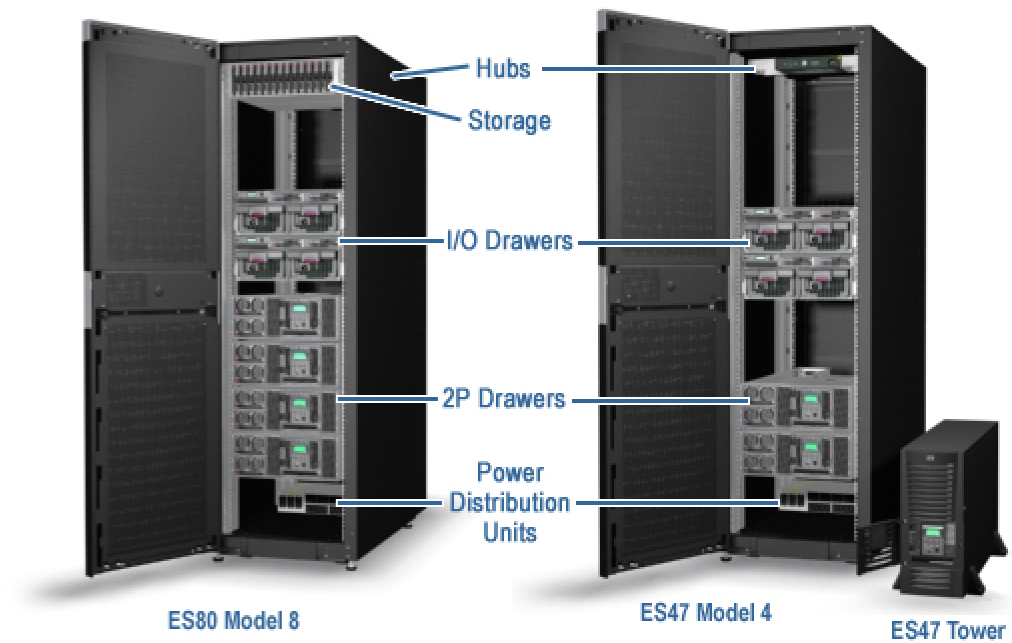

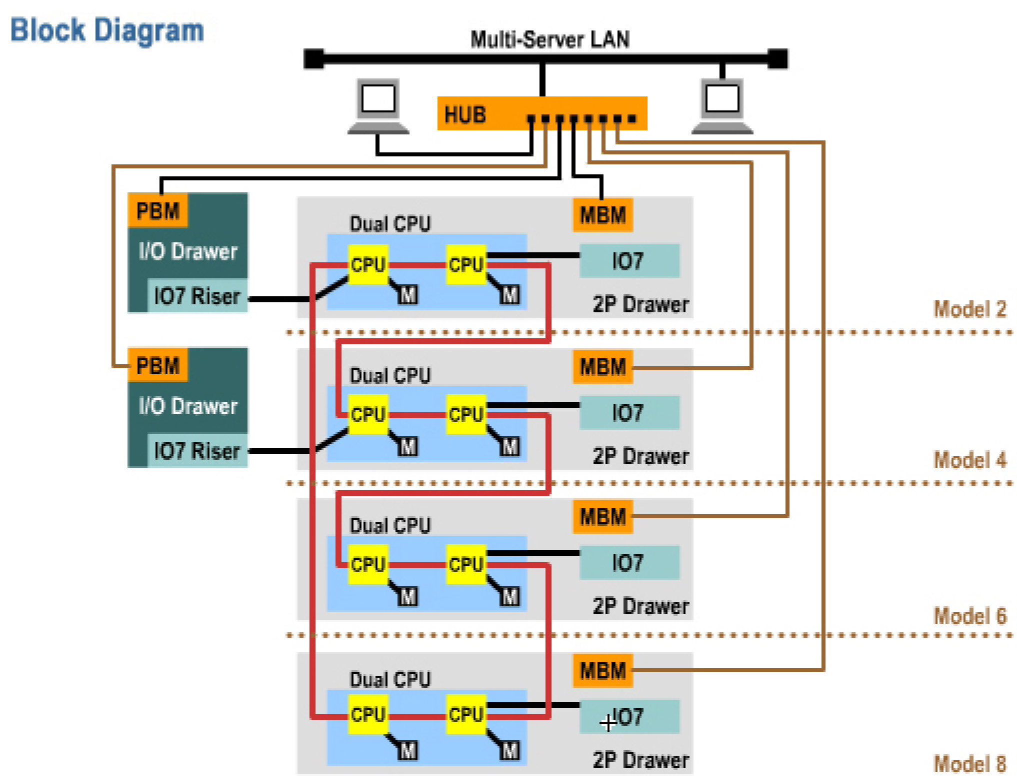

A typical ES80 system may contain four 2P drawers, I/O drawers, storage shelves, and a LAN management HUB, plus additional storage and I/O in an adjoining storage cabinet. An ES47 system supports up to two 2P drawers, two I/O drawers, storage, and the LAN management HUB. The ES47 tower contains one 2P drawer.

CPUs, Memory, and I/O slots	ES47 Tower	ES47	ES80 Model 8
Maximum CPUs supported	2	4	8
Maximum memory supported	16 GB	32 GB	64 GB
Maximum PCI/PCI-X slots supported	5	32	64
Maximum AGP slots supported	1	4	8

At A Glance

- 1000-MHz / 1150 MHz EV7 Alpha 21364 processors supported
- Advanced on-chip memory controllers and switch logic capable of providing 12.3 GB/s of memory bandwidth per processor
- Choice of memory options; up to 8 GB of RDRAM memory per CPU supported
- Redundant features providing maximum uptime - N+1 Voltage Regulator Modules (VRMs); hot- plug redundant power supplies; cooling provided by hot-plug redundant system fans; dual AC input is standard
- 5 PCI-X/PCI slots and one AGP slot in each 2 Processor Building Block Drawer
- Optional RAID memory support
- Optional Standard I/O Drawer with 11 configurable PCI-X/PCI slots and one AGP slot; hot-swap power supplies
- Optional High-performance I/O Drawer with eight PCI-X slots @133 MHz; hot-swap power supplies
- Enhanced reliability with ECC-protected memory, processor cache, and system data paths
- Tru64 UNIX or OpenVMS factory installed software (FIS); optional high availability support with Tru64 UNIX and OpenVMS cluster solutions

ES47 Tower System

This system contains one 2P drawer.

2P Drawer

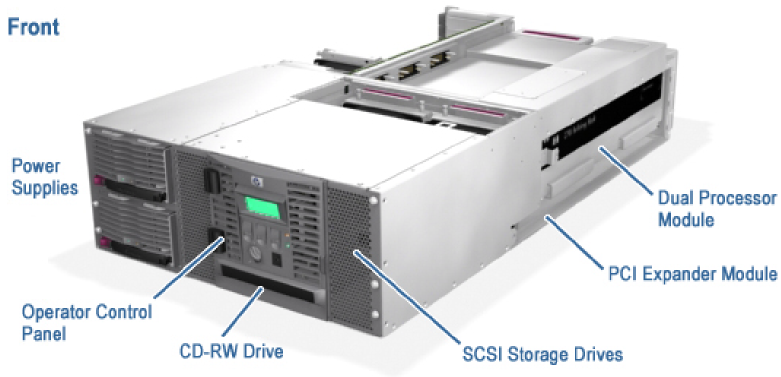

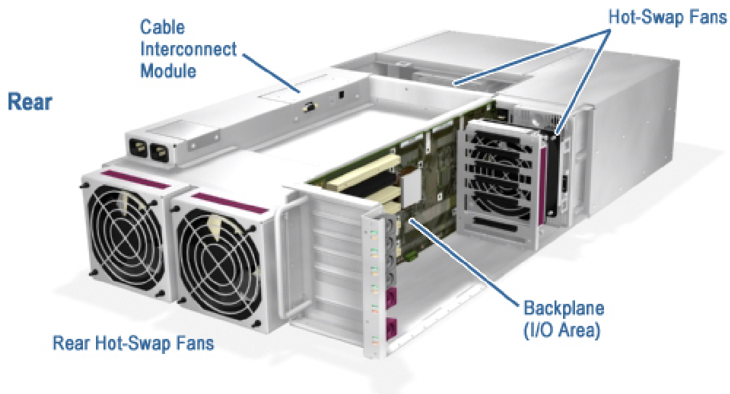

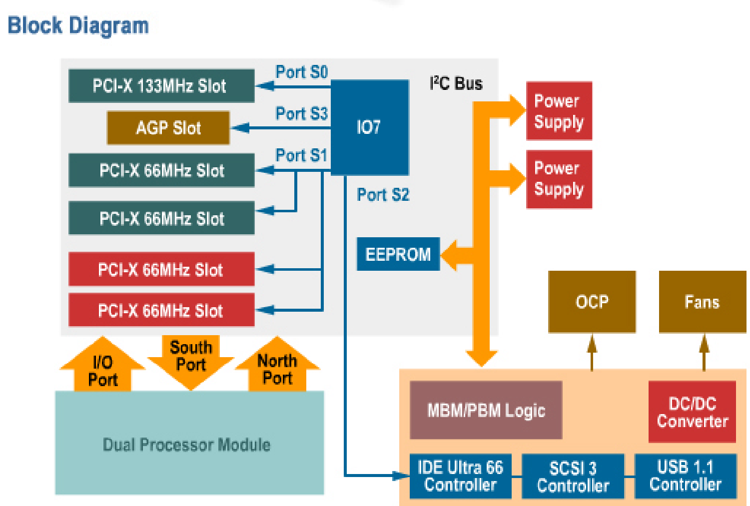

The 2P drawer is the building block for the ES47/ES80 cabinetand pedestal systems. A cabinetsystem can have up to four 2P drawers; a pedestal system contains one 2P drawer.

The 2P drawer has:

- A dual processor module
- A backplane with an IO7 chip, five I/O slots, and an AGP slot
- Two power supplies
- An expansion I/O module with a backplane manager, and connectors for a SCSI drive, CD-ROM, and USB

2P Drawer Control Panel

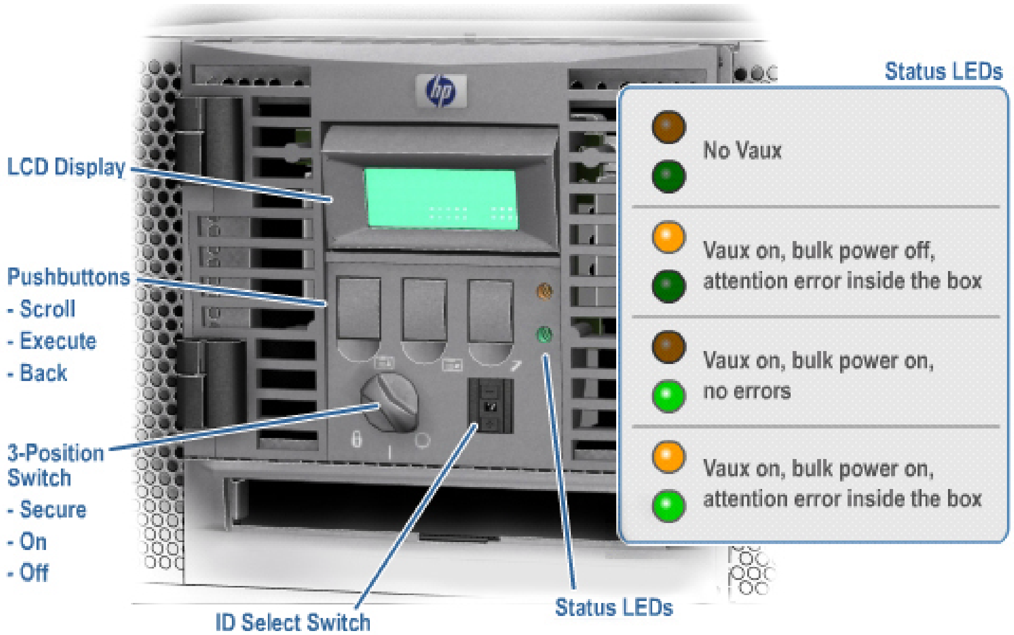

The 2P drawer control panel has an LCD display for menu and system status information. The control panel also has:

- SCROLL, SELECT, and EXECUTE pushbuttons for navigating system menus and executing commands
- A 3-position switch to turn the system on and off, and to secure the drawer
- A FAULT LED and an OK LED
- A pushbutton wheel switch used to set the drawer ID number

Green Power LED	Amber FAULT LED	Description
Off	Off	No Vaux (or no connection to the MBM/PBM
Off	On	Vaux On, bulk power Off, attention error inside drawer
On	Off	Vaux On, bulk power On, no errors
On	On	Vaux On, bulk power On, attention error inside drawer

2P Dual Processor Module

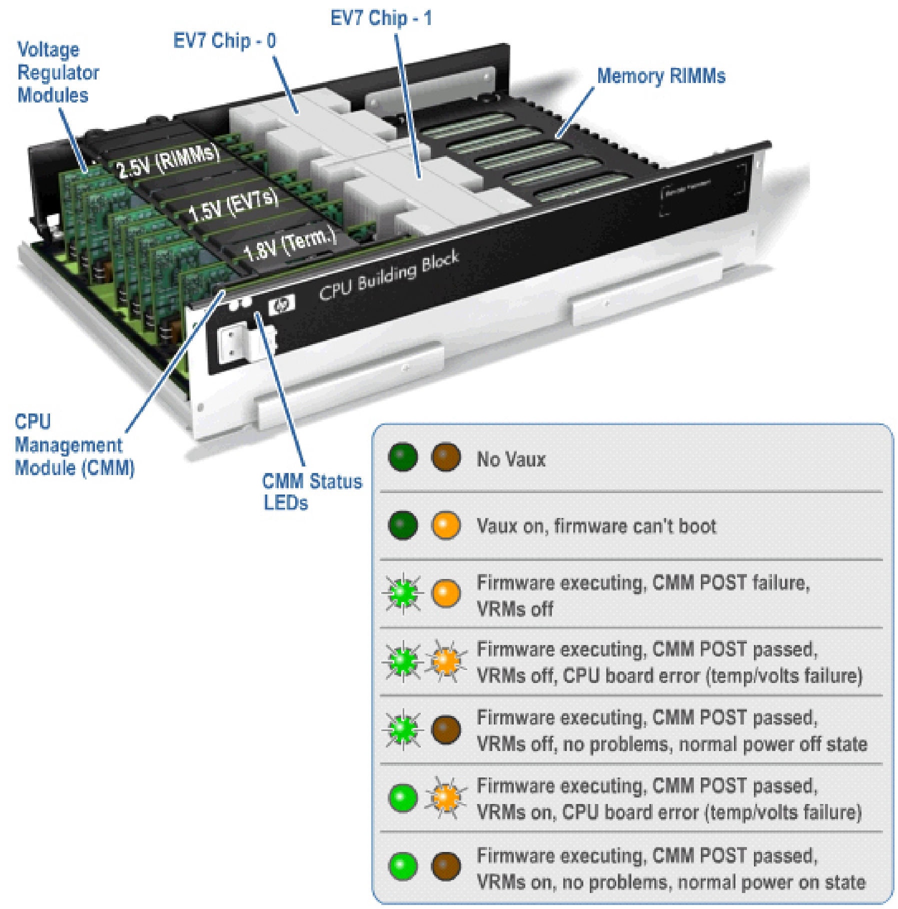

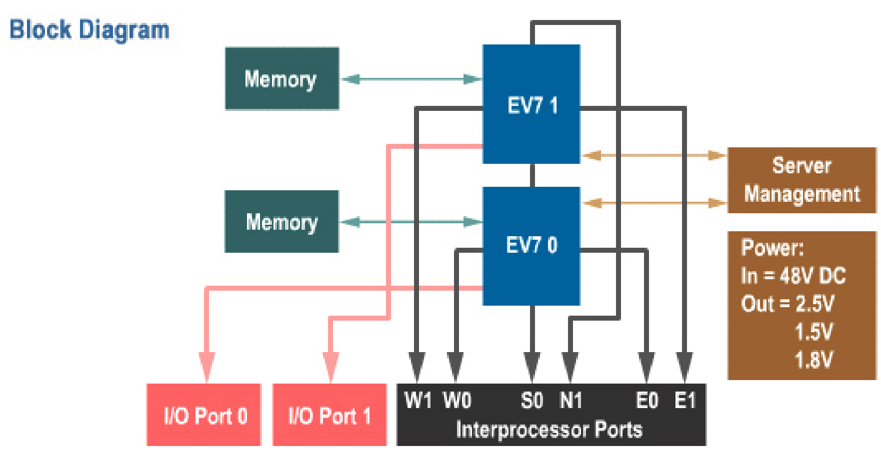

The dual processor module has two EV7 system chips, each with its own memory controllers and I/O port. The module also contains:

- a CPU management module (CMM)
- 20 RIMM slots, supporting up to 16 GB of Rambus memory
- 12 slots for voltage regulator modules (VRMs)

NOTE: A CPU must have local memory if its I/O port is in use.

2P Backplane

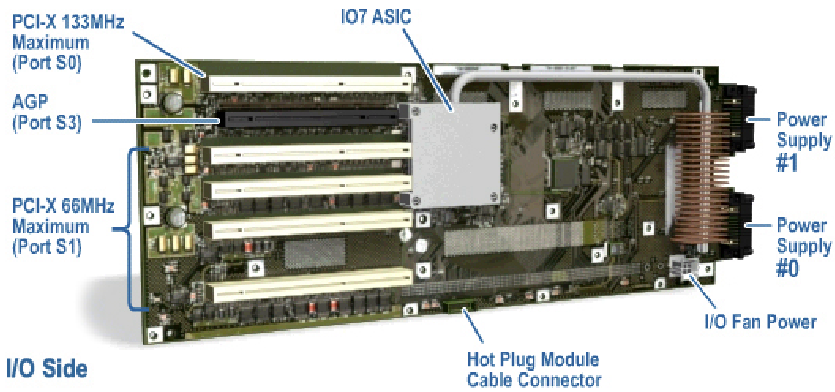

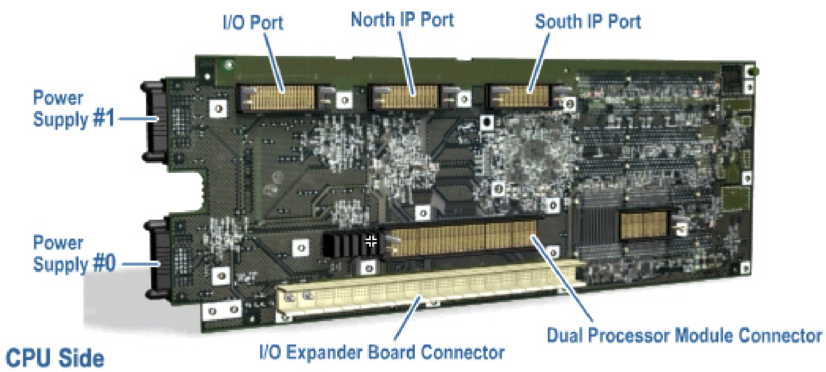

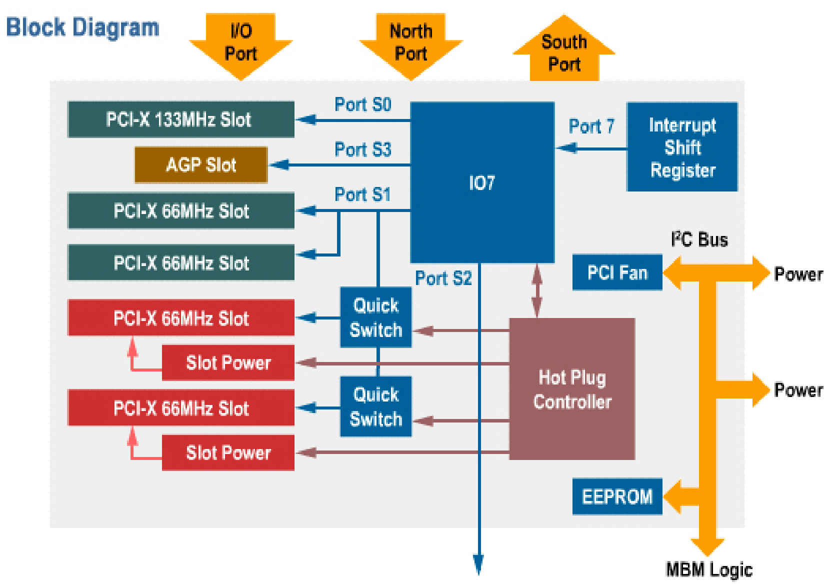

The 2P backplane contains the following:

- Connectors for a dual processor module
- Connectors for a north port, a south port, and an I/O port
- I/O slots for a 1-slot high performance PCI-X bus, two 2-slot PCI-X buses, and an AGP bus.
- An I/O expander module

2P Cable Interconnect Module

External Connectors

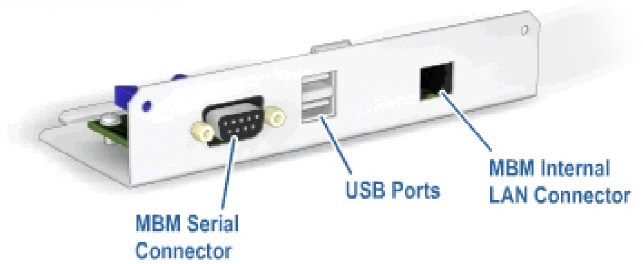

Internal Connectors

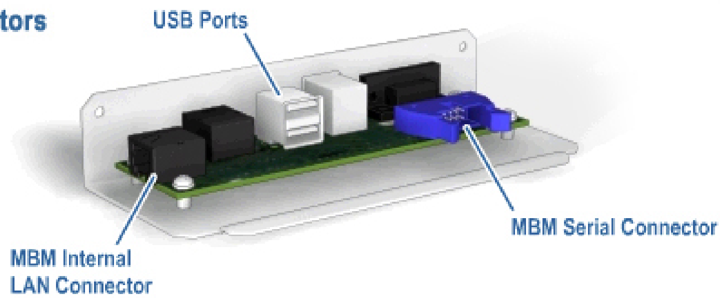

Block Diagram

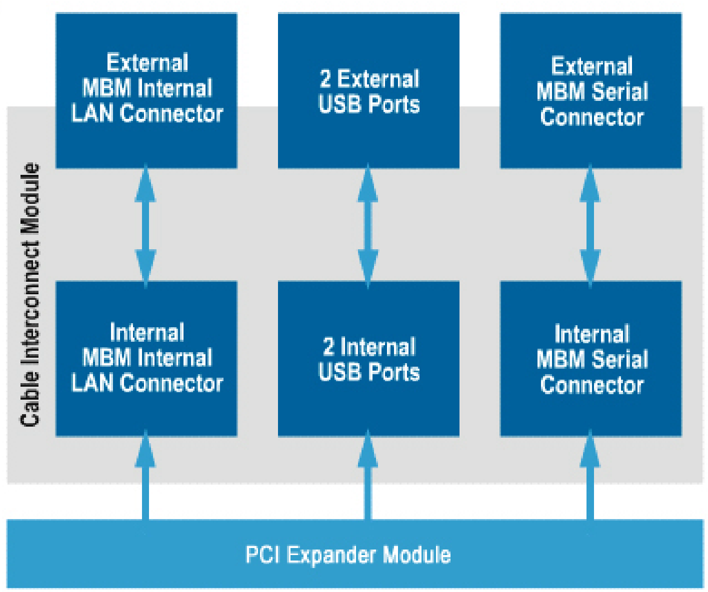

This module connects to the modem port, the USB port, and the LAN connector on the I/O extender card inside the 2P drawer and provides external connections to these ports.

2P I/O Expander Module

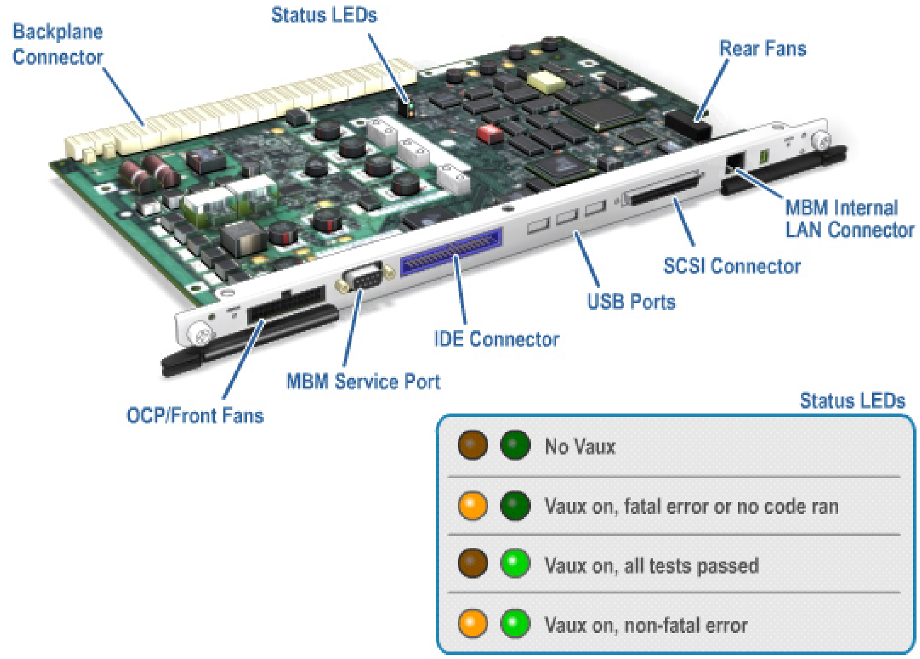

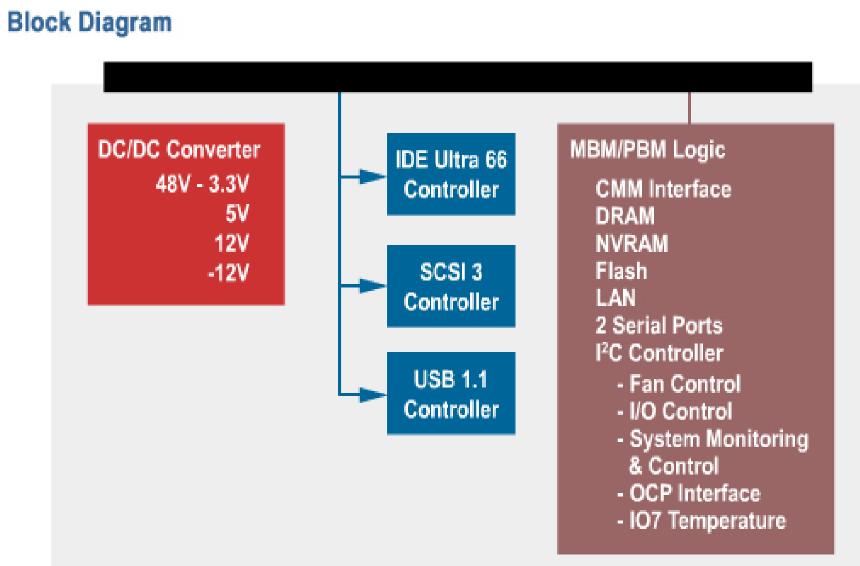

The I/O expander module contains the backplane manager logic for the system. It has an IDE controller for a DVD/CD-ROM, a SCSI controller for storage disks, and three USB connectors for a keyboard, mouse, and modem. All connections to this module are internal to the drawer. The cable interconnect module is used for connections made externally to the drawer.

2P Power Supply

Front

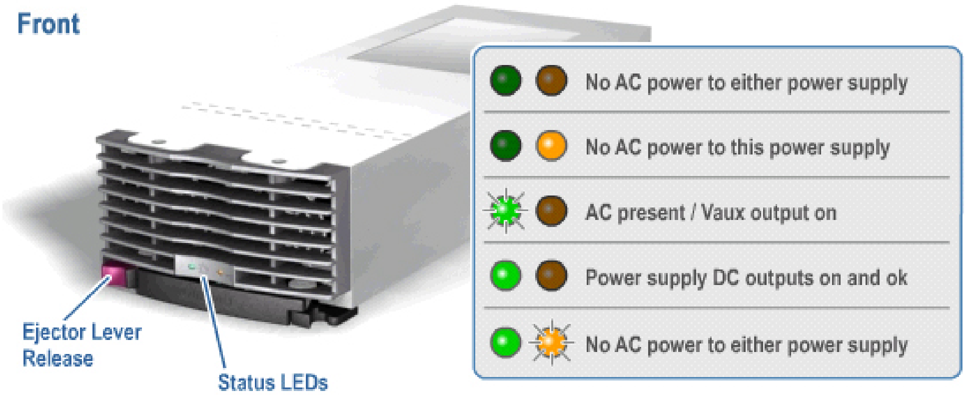

Rear

Block Diagram

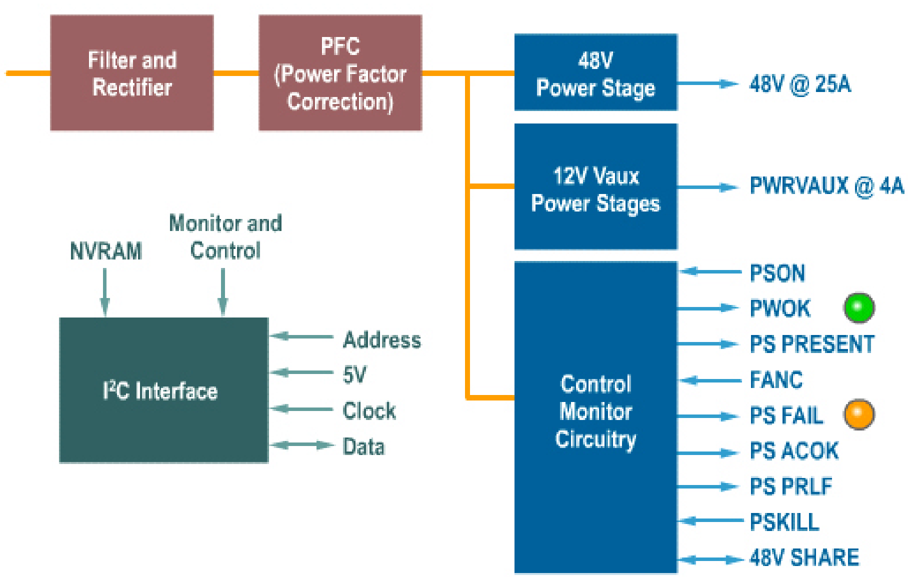

- The power supply has a microprocessor that monitors its environment and function. The power supply can be hot-swapped.
-

2P Configurations by Cabinets

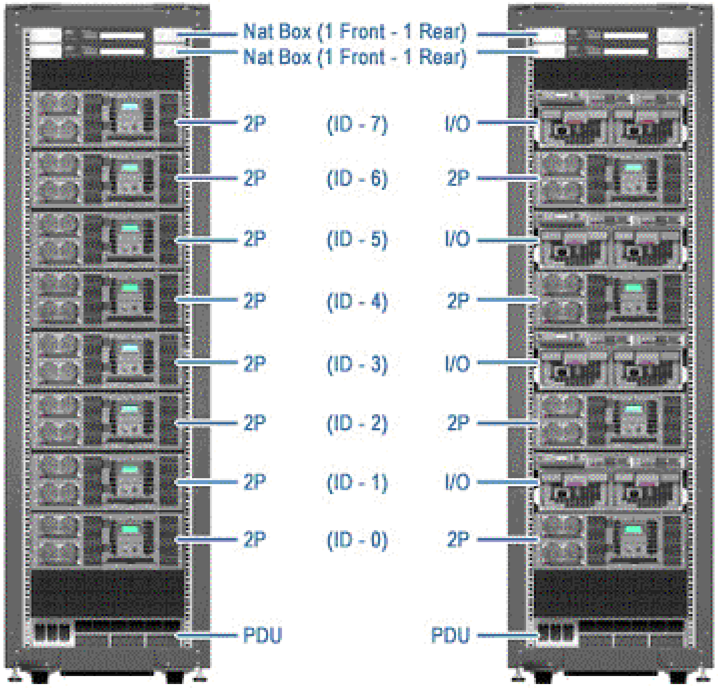

The 41U and 42U cabinets holds a maximum of eight drawers (2P and I/O), and the 34U cabinet holds a maximum of six drawers. For both the ES47 and ES80 systems, I/O drawers must be in the same physical cabinet as the system drawers.

Additionally, you may have more than one system in a cabinet. Each system must have a dedicated NAT box.

For example, in the drawing above, the cabinet on the left could be four M2 systems, each with their own NATbox. Or it could be one M6 and one M2 system, or two M4 systems. The cabinet on the right could be four M2 systems, each with a dedicated I/O drawer.

Another example: you may have an ES47 M4 (2 drawers) and an ES80 M2 (1 drawer) system in the same cabinet, leaving 5 drawers for I/O or future upgrades. Or you may have two ES47 systems in the same cabinet, for example an M4 (2 drawers) and an M6 (3 drawers) with 3 drawers available for I/O or future upgrade.

- Each system must be connected to a NAT box dedicated to that system only.

You may have up to 4 NAT boxes installed in a cabinet.

2P Configurations per NAT Box

- You may have more than one ES47 or ES80 system in a cabinet.
- Your possible configurations are dependent on the number of NAT boxes present in your system, to connect each system in the cabinet. And if you have I/O drawers installed, this will also limit the number of spaces available for system drawers.

Here is a list of 2P configurations showing the combination of models you can put in a cabinet, sorted by the number of NAT boxes required.

# of NAT boxes	Models	# of 2P drawers
1	1 M2	1
1	1 M4	2
1	1 M6	3
1	1 M8	4
2	2 M2	2
2	1 M2 + 1 M4	3
2	1 M2 + 1 M6	4
2	2 M4	4
2	1 M2 + 1 M8	5
2	1 M4 + 1 M6	5
2	2 M2 + 1 M8	6
2	2 M6	6
2	1 M6 + 1 M8	7
2	2 M8	8
3	3 M2	3
3	2 M2 + 1 M4	4
3	2 M2 + 1 M6	5
3	1 M2 +2 M4	5
3	1 M2 + 1 M4 + 1 M6	6
3	3 M4	6
3	1 M2 + 1 M4 + 1 M8	7
3	2 M4 + 1 M6	7
3	1 M2 + 1 M6 + 1 M8	8
3	1 M4 + 2 M6	8
3	2 M4 + 1 M8	8
4	4 M2	4
4	3 M2 + 1 M4	5
4	1 M2 + 2 M4	6
4	3 M2 + 1 M6	6
4	1 M2 + 3 M6	7
4	3 M2 + 1 M8	7
4	2 M2 + 1 M4 + 1 M6	7
4	4 M4	8
4	2 M2 + 1 M4 + 1 M8	8
4	2 M2 + 2 M6	8
4	1 M2 + 2 M4 + 1 M6	8

NAT Box Configurations, by Drawers

When adding a NAT box, you may choose to re-configure your current 2P drawers. The following chart shows you the options you have, how many NAT boxes the configuration requires, and it is sorted by the number of drawers available.

M2 systems connected to an I/O drawer require a NAT box. Standalone M2 systems (i.e., without an I/O drawer) are usually managed by a serial port connection. While you can attach a NAT box, serial port connection is recommended for standalone M2s, not NAT box.

# of 2P drawers	Models	# NAT boxes
1	1 M2	1
2	1 M4	1
	2 M2	2
3	1 M6	1
	1 M2 + 1 M4	2
	3 M2	3
4	1 M8	1
	1 M2 + 1 M6	2
	2 M4	2
	2 M2 + 1 M4	3
	4 M2	4
5	1 M2 + 1 M8	2
	1 M4 + 1 M6	2
	2 M2 + 1 M6	3
	1 M2 +2 M4	3
	3 M2 + 1 M4	4
6	2 M2 + 1 M8	3
	2 M6	2
	1 M2 + 1 M4 + 1 M6	3
	3 M4	3
	1 M2 + 2 M4	4
	3 M2 + 1 M6	4
7	1 M6 + 1 M8	2
	1 M2 + 1 M4 + 1 M8	3
	2 M4 + 1 M6	3
	1 M2 + 3 M6	4
	3 M2 + 1 M8	4
	2 M2 + 1 M4 + 1 M6	4
8	2 M8	2
	1 M2 + 1 M6 + 1 M8	3
	1 M4 + 2 M6	3
	2 M4 + 1 M8	3
	4 M4	4
	2 M2 + 1 M4 + 1 M8	4
	2 M2 + 2 M6	4
	1 M2 + 2 M4 + 1 M6	4

GS1280 System

GS1280 System Overview

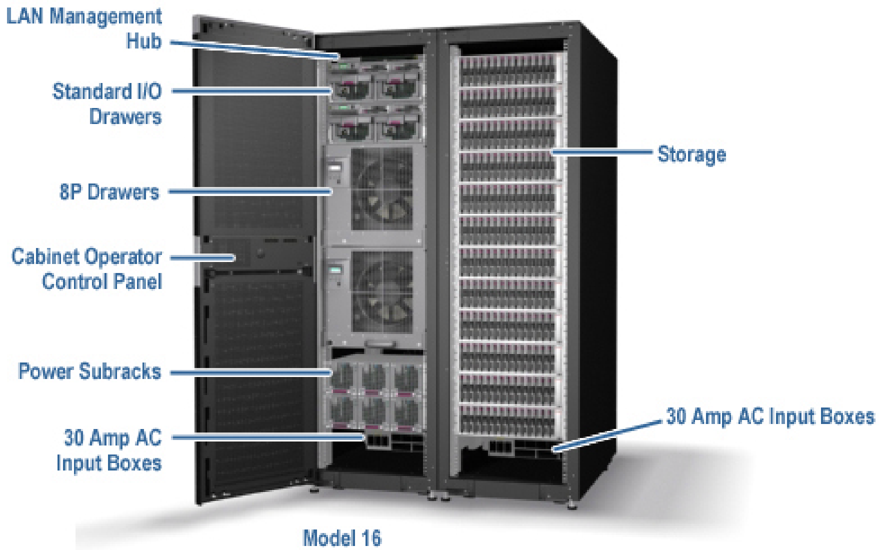

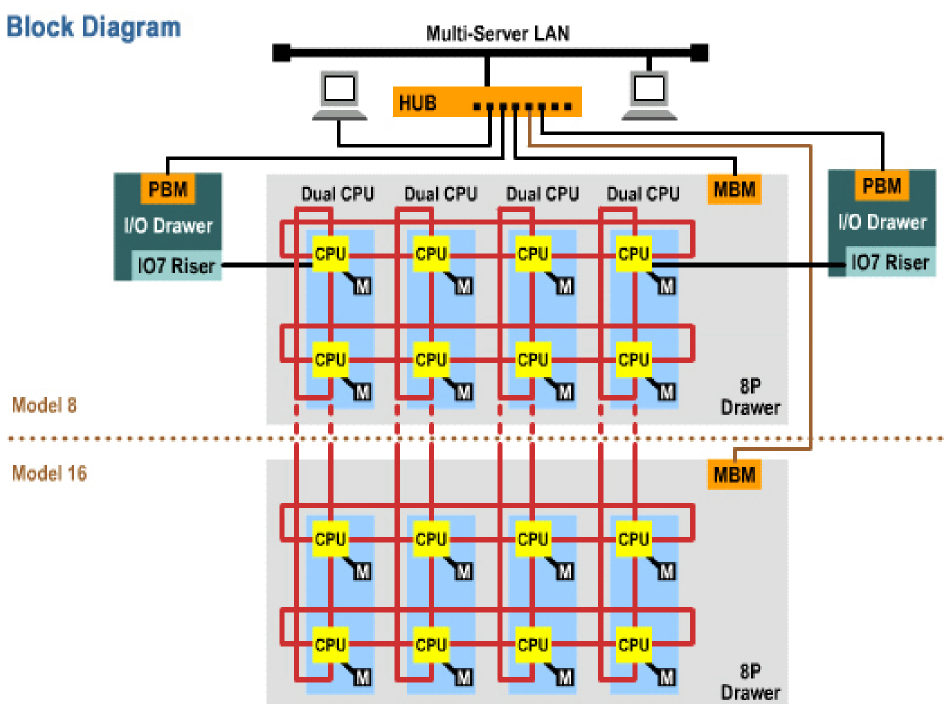

A typical system may have two 8P drawers, power subracks, I/O drawers, and a LAN management HUB in the system cabinet, plus additional storage and I/O in an adjoining storage cabinet.

GS1280 Model 32

Model 32

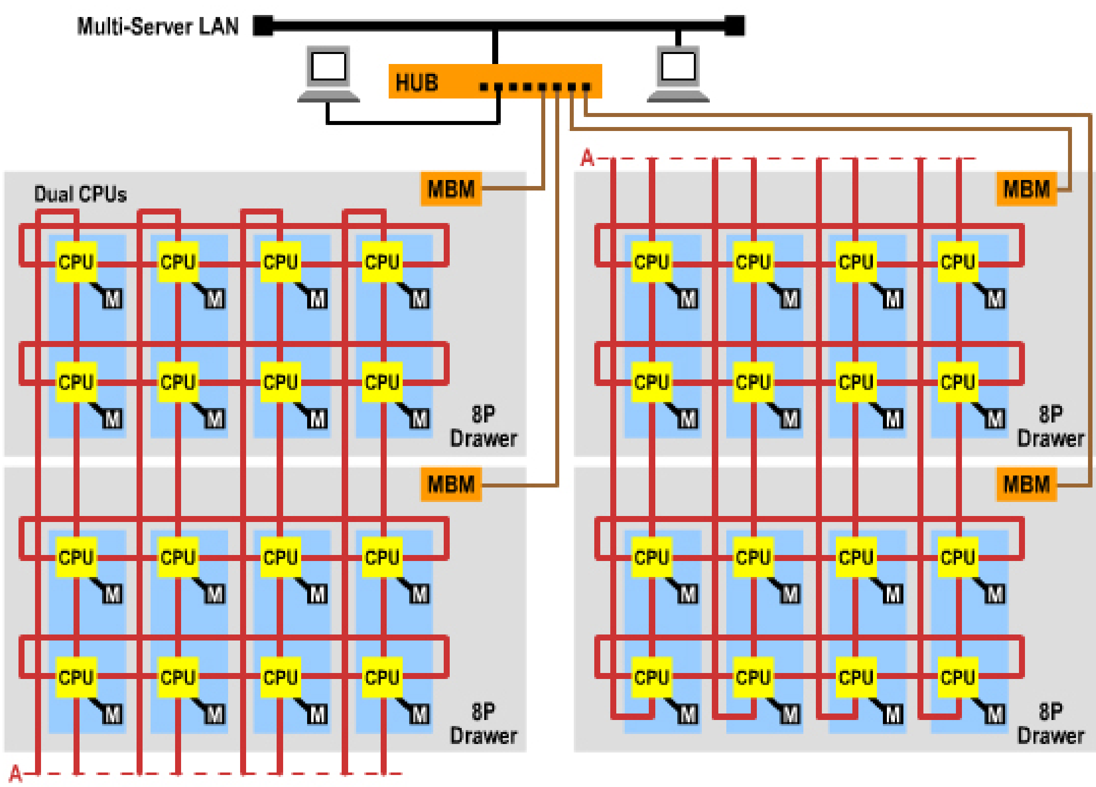

GS1280 Model 64

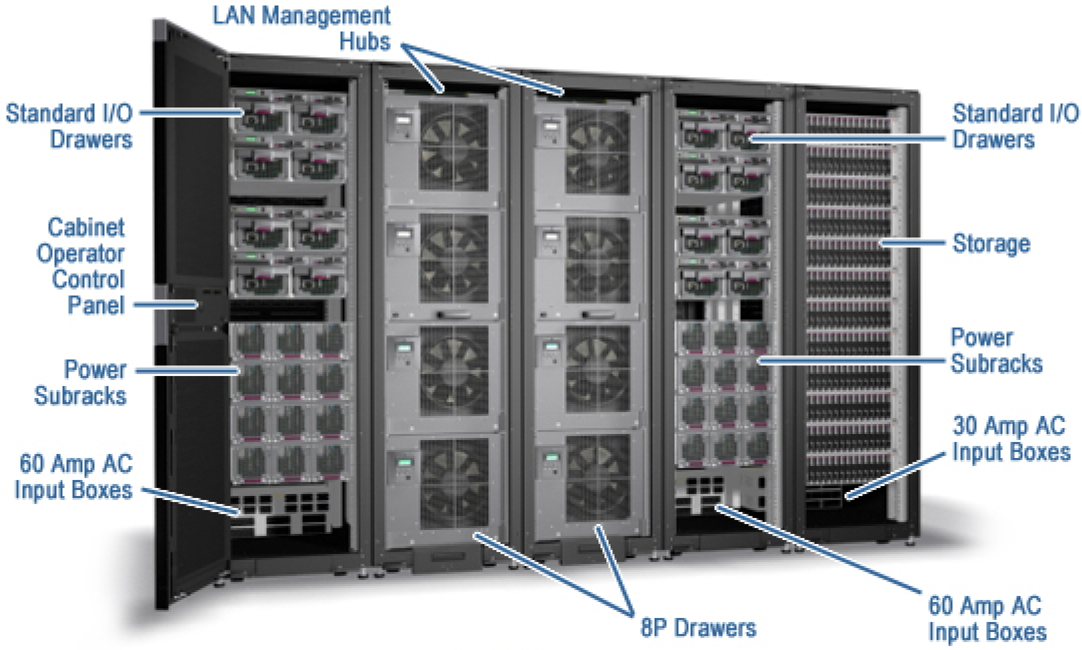

Model 64

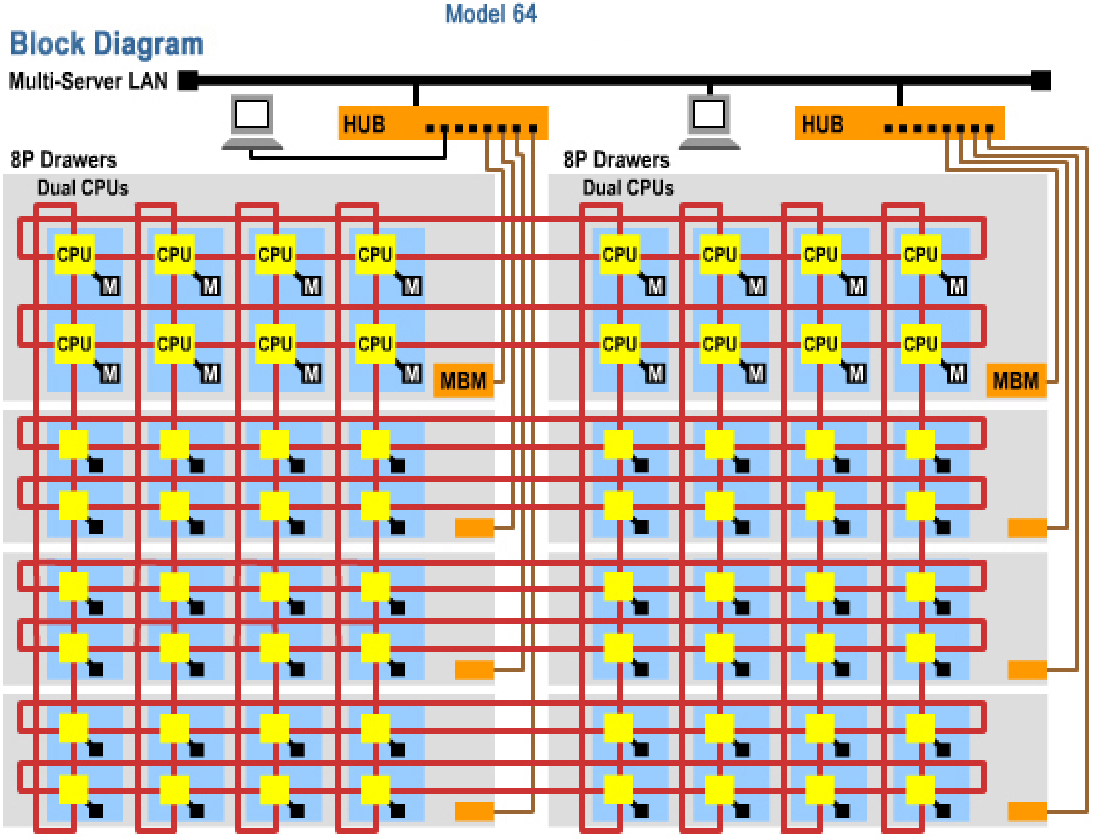

CPUs, Memory, and I/O slots	GS1280 Model
Maximum CPUs supported	64 per hard partition for Tru64 UNIX, 32 per hard partition for OpenVMS
Maximum memory supported	512 GB
Maximum PCI-X/PCI slots supported	704 (176 per hard partition)
Maximum AGP slots supported	64 (16 per hard partition)

At A Glance

- 1150-MHz / 1300 MHz Alpha 21364 processors supported (up to 64)
- Advanced on-chip memory controllers and switch logic capable of providing 12.3 GB/s of memory bandwidth per processor
- Choice of memory options; up to 8 GB of RDRAM memory per CPU supported (512 GB total for a 64P system)
- Redundant features providing maximum uptime - N+1 Voltage Regulator Modules (VRMs); hot- plug redundant power supplies; cooling provided by hot-plug redundant system fans; dual AC input is standard
- Optional RAID memory support
- Standard I/O Drawer with 11 configurable PCI-X/PCI slots and one AGP slot; hot-swap power supplies
- High-performance I/O Drawer with eight PCI-X slots @133 MHz; hot-swap power supplies
- Enhanced reliability with ECC-protected memory, processor cache, and system data paths
- Tru64 UNIX or OpenVMS factory installed software (FIS); optional high availability support with Tru64 UNIX and OpenVMS cluster solutions

8P Drawer

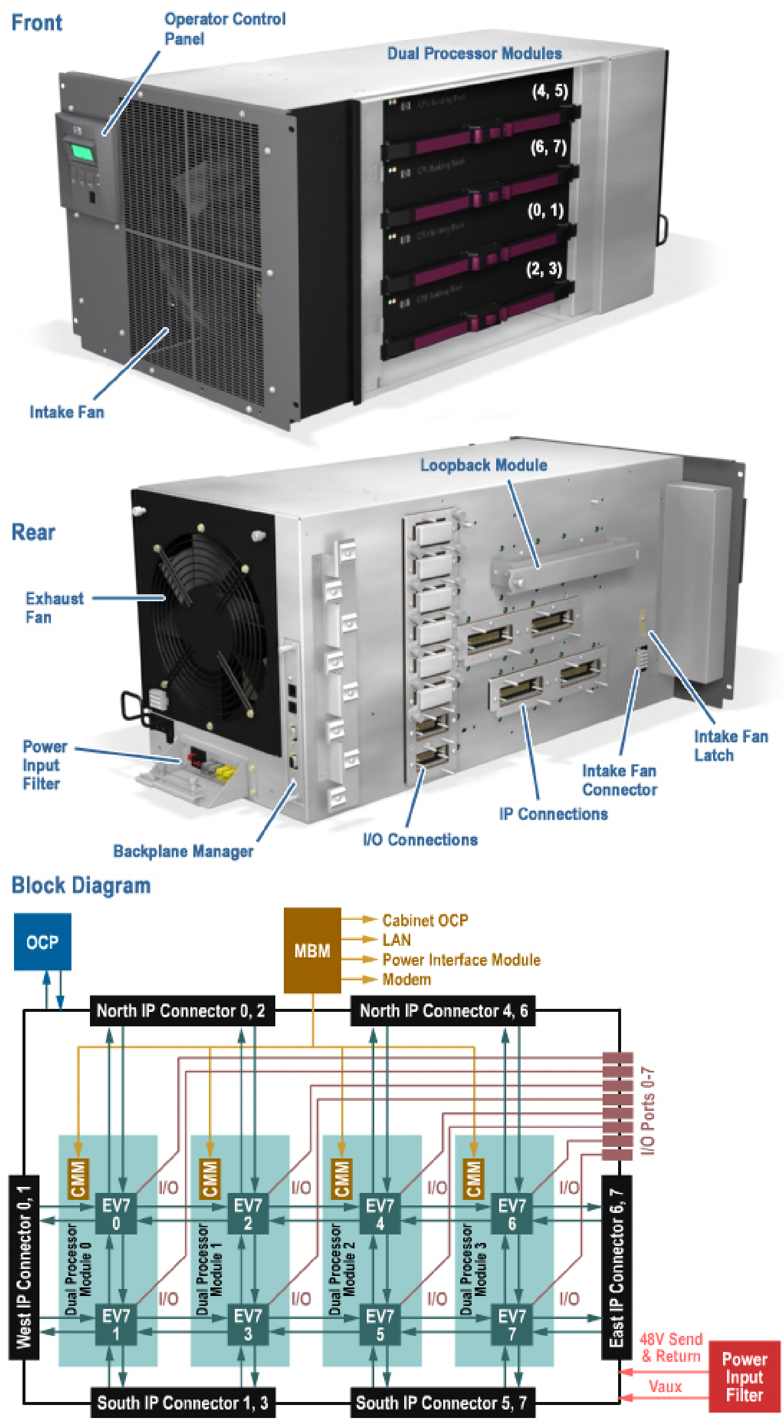

The 8P drawer is the building block for GS1280 systems. It contains up to four dual-processor modules, a backplane manager (MBM) module that monitors and manages the system, and a control panel. Each 8P drawer supports up to eight I/O drawers.

8P Drawer Control Panel

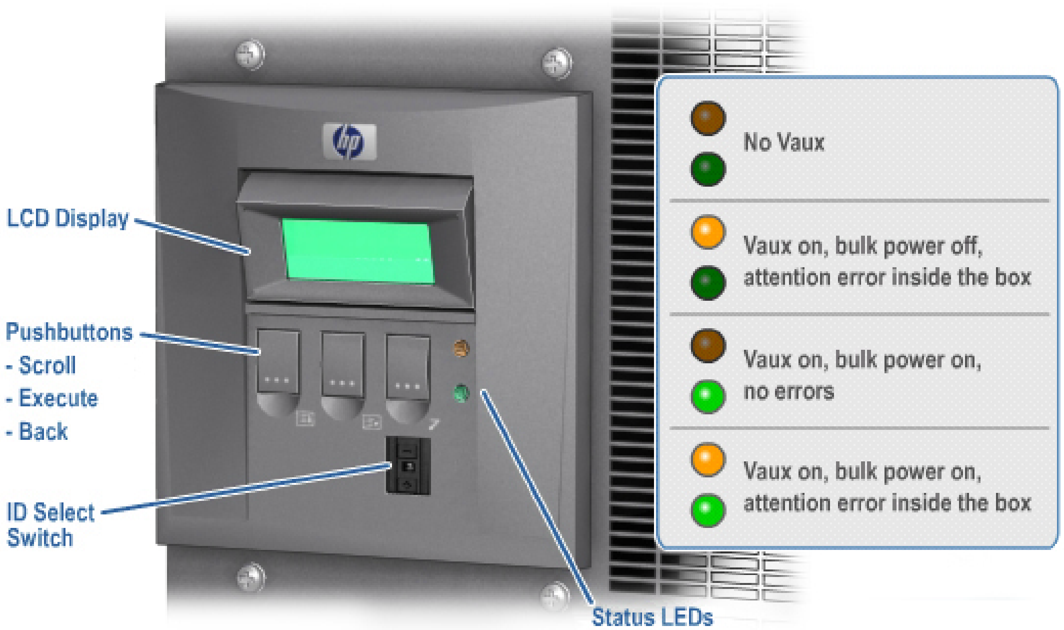

GS1280 / 8P cabinet systems have 2 control panels: The 8P drawer has a control panel and the 8P cabinet also has a cabinet control panel. Note that 2P cabinet systems only have control panels on the drawers, none for the cabinet itself.

The 8P drawer control panel has an LCD display for menu and system status information, and also has the following:

- SCROLL, EXECUTE, and BACK pushbuttons for navigating system menus and executing commands
- A FAULT LED and a POWER LED
- A push-wheel used to set the drawer ID number

Control Panel Pushbuttons

Scroll	Use this button to locate the command you wish to execute.
Execute	Pressing this button executes the command or selects the menu you wish to use.
Back	Use this button to clear information or to return to the beginning of a process or menu tree.

Control Panel LEDs

Green Power LED	Amber FAULT LED	Description
Off	Off	No Vaux (or no connection to the MBM/PBM
Off	On	Vaux On, bulk power Off, attention error inside drawer
On	Off	Vaux On, bulk power On, no errors
On	On	Vaux On, bulk power On, attention error inside drawer

8P Dual Processor Module

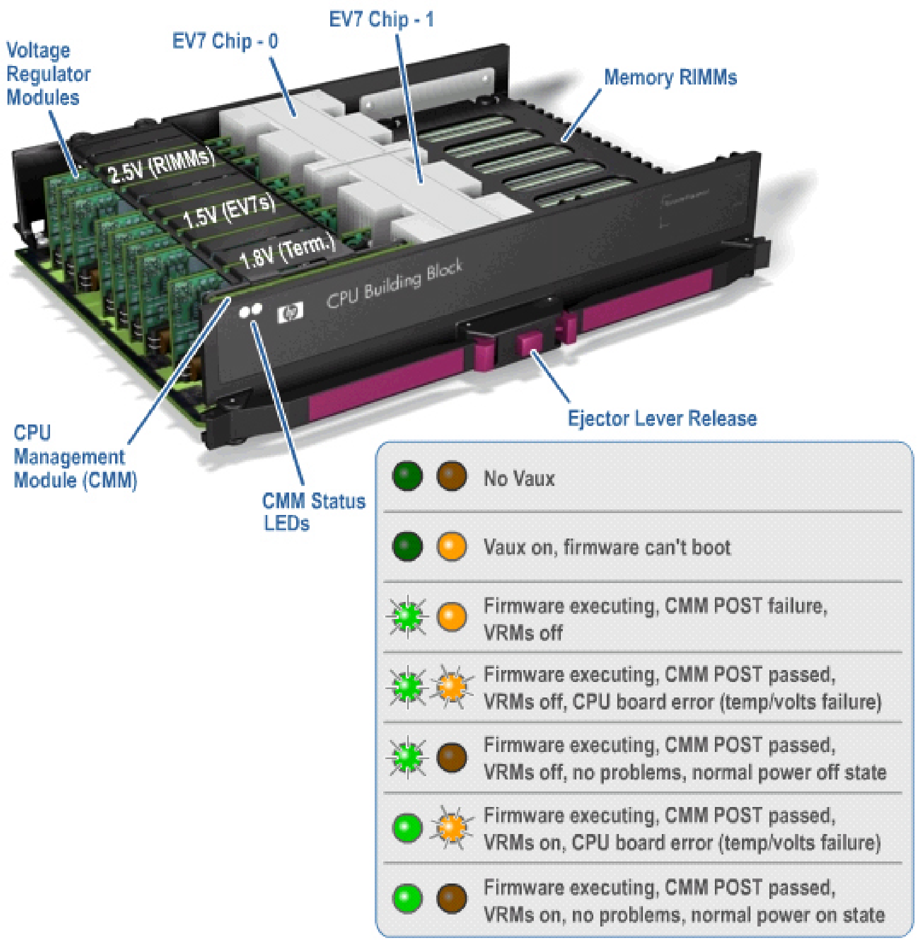

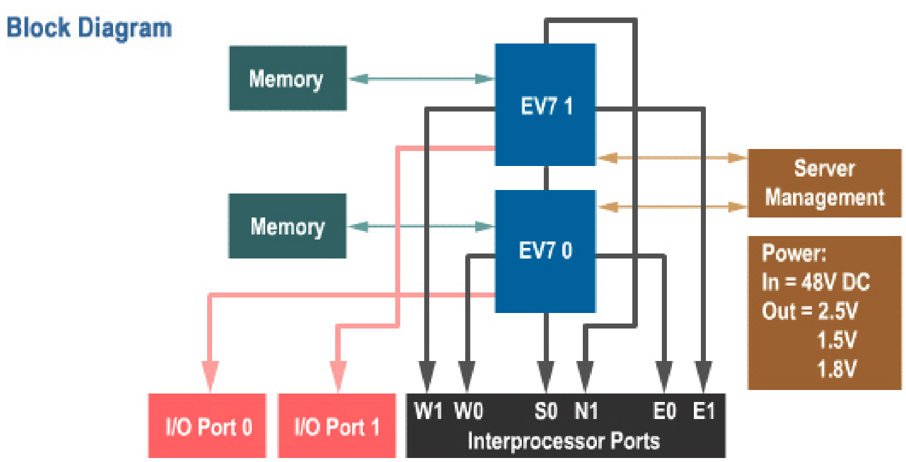

The dual processor module has two EV7 system chips, each with its own memory controllers and I/O port. The module also contains:

- a CPU management module (CMM)
- 20 RIMM slots, supporting up to 16 GB of Rambus memory
- 12 slots for voltage regulator modules (VRMs)=

NOTE: A CPU must have local memory if its I/O port is in use.

8P Backplane

External Connectors

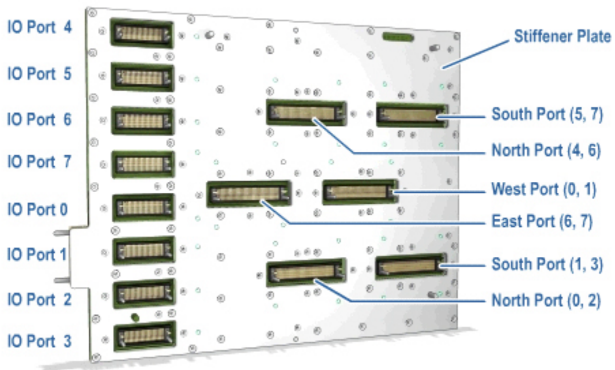

Internal Connectors

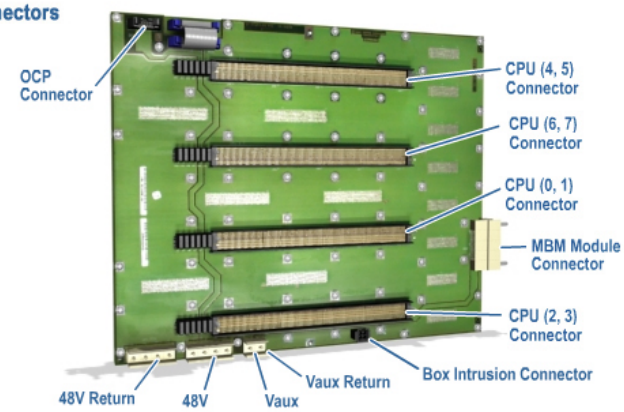

Block Diagram

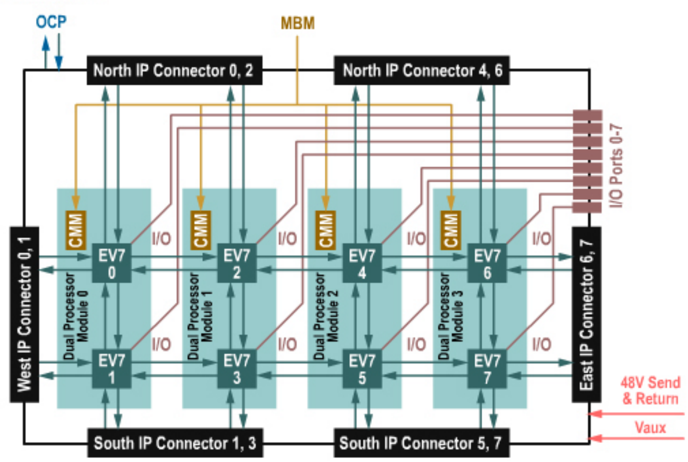

The 8P backplane distributes system power and interconnects the dual processor modules present in the drawer. It has

- eight external I/O connectors
- six external interprocessor connectors, plus
- power connectors

The backplane manager (MBM) module and OCP module connect to the backplane.

8P Backplane Manager

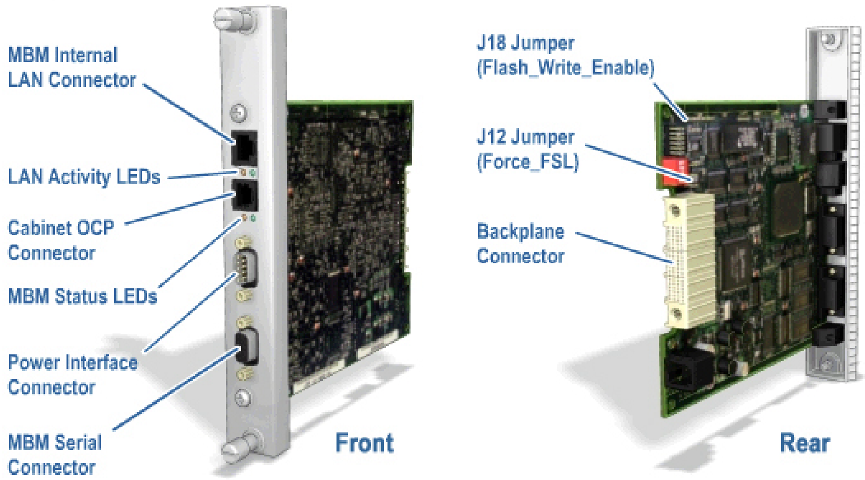

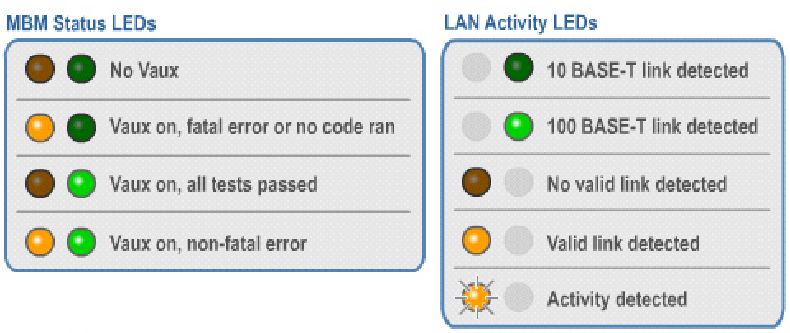

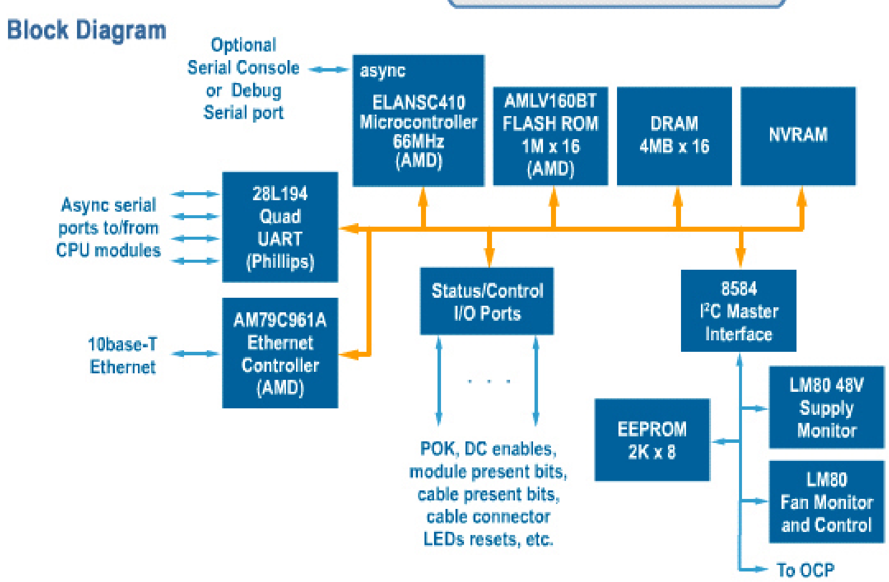

The backplane manager (MBM) controls the CPU management modules (CMMs) and contains logic to monitor and control environmental conditions in the 8P drawer. The MBM connects to the server management LAN. Other functions of the MBM include:

- Storing system configuration and partition information
- Loading the console and initializing the system
- Allowing hot-swapping of the dual processor modules and I/O drawers

8P Configuration Guidelines

The first expansion cabinet should be located to the right of the system cabinet; the second expansion cabinet, to the left of the system cabinet.

Setting ID Switches

Cabinets, 8P drawers, I/O drawers, and power subracks are assigned identification (ID) numbers. Use the following guidelines:

- Set the ID switch to 0 on the first cabinet containing 8P drawers.
- If a second cabinet contains an 8P drawer, set the ID switch to 1.
- 8P drawer ID switches are set from bottom to top, starting with a 0 setting.
- After assigning ID numbers to the 8P drawers, set ID numbers for other components (I/O drawers, and power subracks). Starting with the lowermost component; set the ID switch to the next consecutive number, and proceed to the topmost component in the cabinet. For example, a 16P system containing two power subracks and two I/O drawers has the following bottom-to-top ID number assignments: 2, 3, 0, 1, 4, 5.

Cabinets

Cabinet Controls and Indicators

The 8P cabinet has a control panel which includes:

- an On/Off/Secure switch
- a pushwheel for setting the cabinet ID
- Power, Fault, and Cabinet Status LEDs, and a cabinet ID display

The 2P cabinets do not have a control panel. You will find additional control panels on the 8P drawer and on the 2P drawers themselves.

8P Cabinet Control Panel

Outside

Inside

Block Diagram

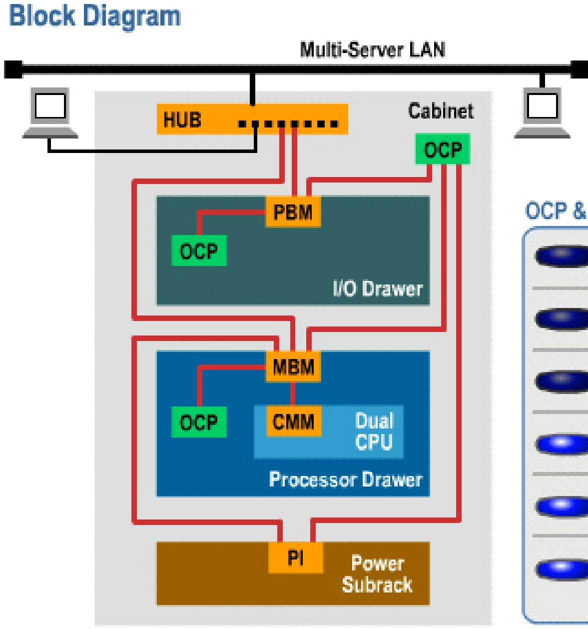

OCP & Cabinet Status LEDs

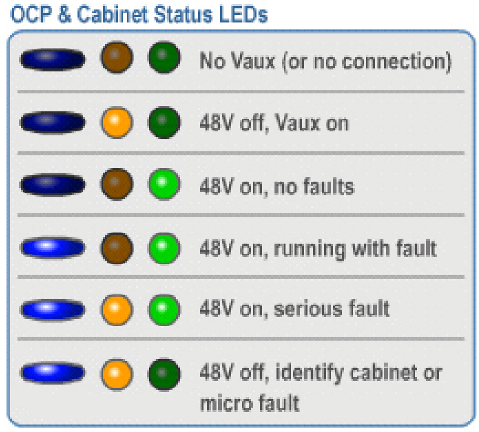

On/Off/Secure Switch

In configurations having more than one system cabinet, one switch is designated (Cabinet 0 keyswitch) to function. All other cabinet switches are inoperable.

Switch Position	System State
On	• All partitions are powered on. • Components not in a partition will not be powered on. • If main AC fails and then returns, all partitions are powered up, regardless of the "soft state" at the time of the power failure.
Off	All partitions and components are powered off. Power on commands issued via the LAN, control panel, or the MBM CLI are prevented and receive an error response. If main AC power fails and returns, the system will remain off.
Secure	All partitions are powered on. Commands issued via the LAN, control panel, or the MBM CLI which change the state of the system are prevented and receive an error response. If main power fails and returns, the system will power up all partitions, regardless of its soft state at the time of the power failure.

Power, Fault, and Cabinet Status LEDs

- When the green Power LED is on, power is on to at least one system drawer in the cabinet.
- When the blue Cabinet Status LED is on, either service is required, or, the cabinet is being identified from server management.
- When the amber Fault LED is on, a failure has occurred, or, tests have not been run.

Green Power LED	Amber Fault LED	Blue Cabinet Status LED	Status
Off	***Off***	***Off***	Vaux is off, or there is no connection
Off	***On***	***Off***	Vaux is on; 48v power is off
On	***Off***	***Off***	48v power is on, no faults
On	***Off***	***On***	48v power is on, running with fault
On	***On***	***Off***	48v power is on, serious fault (probably not running)
Off	***On***	***On***	48v power is off, identify cabinet or micro fault

Cabinets

Cabinets (rack and expansion) come in 3 widths: 41-, 34- and 42 U.
Dimensions are given below.

Dimensions (mm)	
41U Cabinet (H9A45)	
Height	2000
Width	600
Depth	1200
Service Clearance	
Front	800
Sides	0
Rear	1000
Shipping Dimensions	
Height	2170
Width	813
Depth	1219

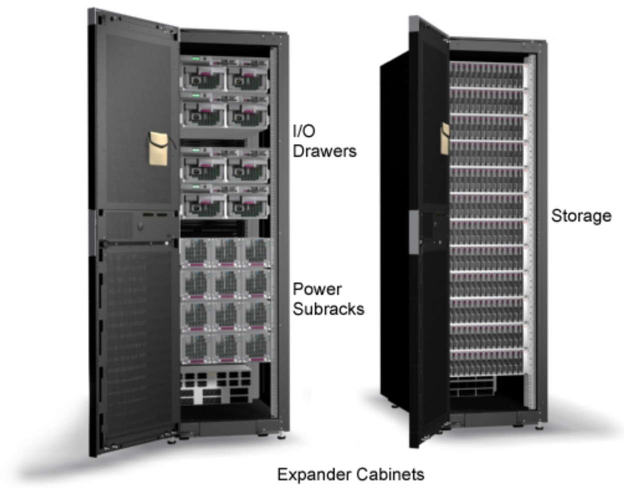

34U Cabinet (H9A40)	
Height	1700
Width	600
Depth	1200
Service Clearance	
Height	800
Width	0
Depth	1000
Shipping Dimensions	
Height	1897
Width	813
Depth	1219

42U Cabinet (10642)	
Height	2000
Width	610
Depth	1008
Service Clearance	
Height	800
Width	0
Depth	1000
Shipping Dimensions	
Height	2190
Width	813
Depth	1219

Connecting the Cabinet

Grounding cabinets together is important in creating a stable electrical environment throughout the system

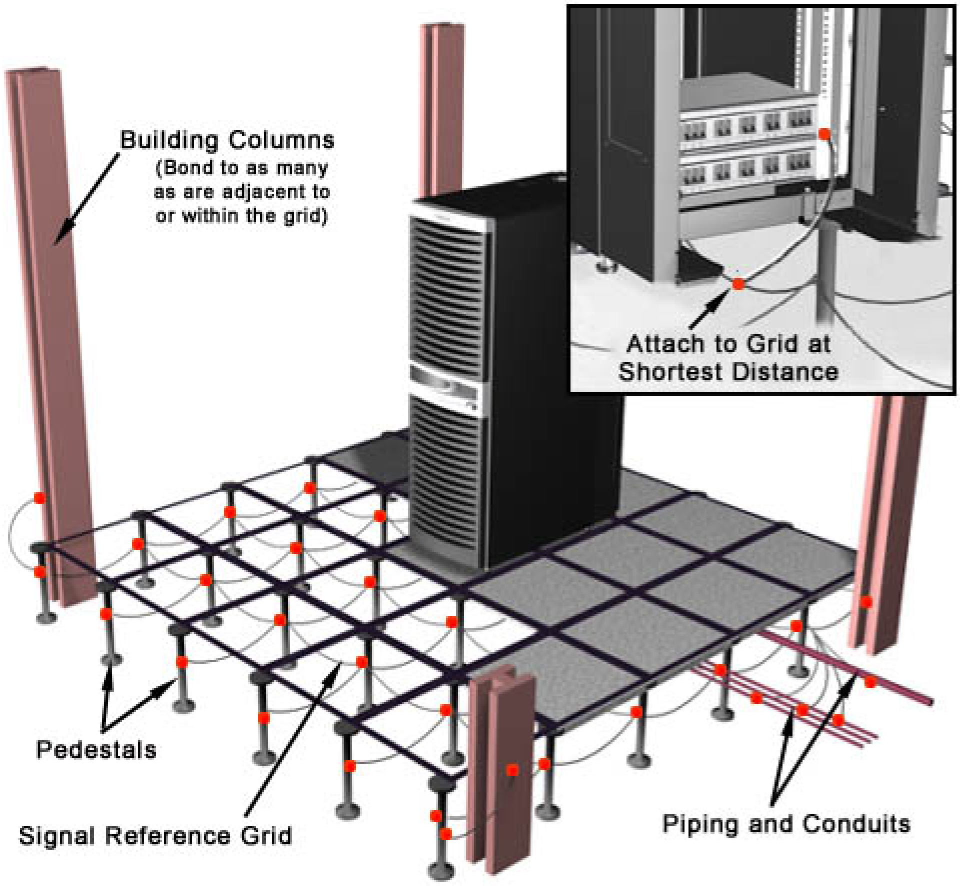

Memory

Configuration Guidelines

- All the RIMMs in a memory option must be the same size and the same DRAM technology.

For example: You cannot mix a 128Mbyte RIMM using 128Mbit Rdrams (8 devices per RIMM) with a 128Mbyte RIMM using 256Mbit Rdrams. (4 devices per RIMM)

- In a dual-processor CPU module, each processor can have unequal amounts of memory, but the two processor controllers (called ZBOXes) associated with a single CPU must control equal amounts of memory using the same memory option part number.

- When adding memory, Zbox-0 or memory controller 0 must be populated before Zbox-1. Memory that is populated only in Zbox-1 will not be utilized.

- Each CPU supports one or two memory options

- Each 2P drawer must have (1) memory option per DUO
  Each 8P drawer must have (2) memory options per DUO

- Memory option per attached I/O required for all ES47/80 and GS1280 systems.

- For ES47 and ES80 systems, any I/O Expansion Drawer must have at least one memory option for each CPU in the CPU Building Block Module connected to the I/O Drawer.

- The RAID option must be the same size and density memory RIMM and the memory option.

- Each memory option consists of four RDRAM Inline Memory Modules (RIMMs). An optional fifth RIMM (RAID option) may be selected for redundancy that will allow uninterrupted operation in case of the loss of an entire RIMM. RAID options must be selected for each memory option on one CPU, but RAID options do not have to be selected for all CPUs.

For more information:

- Up to date configuration rules and memory options are listed in the QuickSpecs for the ES47/ES80 and the GS1280 systems.

Adding Memory

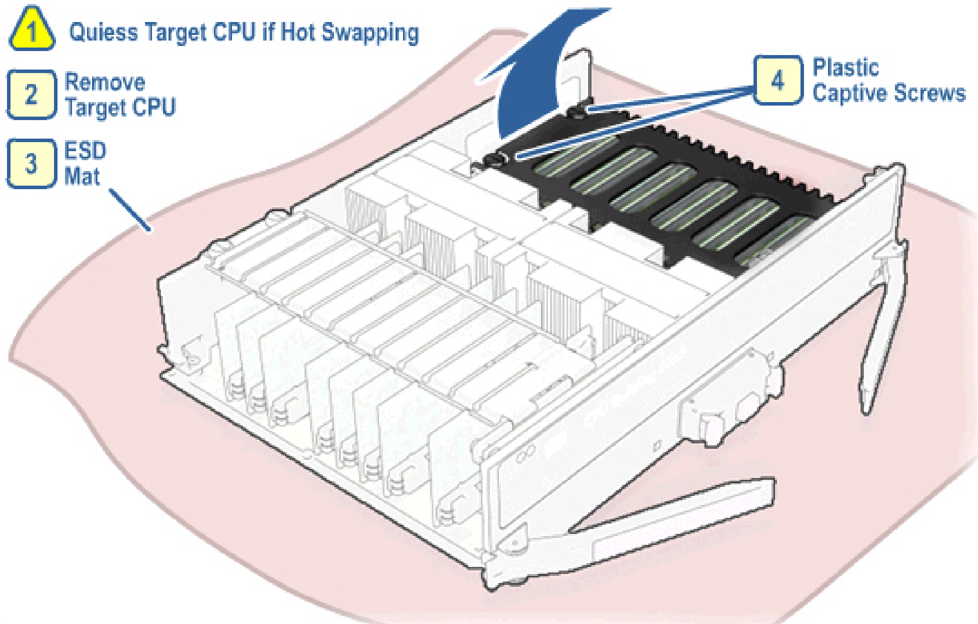

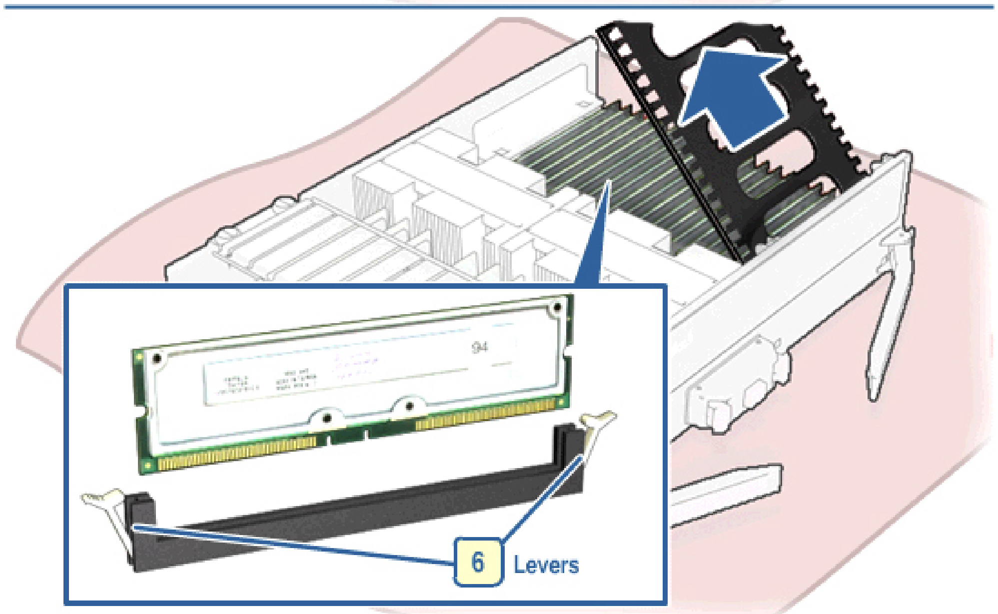

You may view and print out the instructions for installing memory by clicking here.
Observe static precautions at all times.

Removal

1. If you intend to hot-swap or add memory, be sure to quiess the CPU you remove.
2. Remove the CPU containing the target memory.
3. Place the CPU module on an ESD mat.
4. To get the RIMM hold-down bracket out of the way:
   1. Loosen the two plastic cap screws that release the far end of the bracket.
   2. Swing the free end up and pull the other end off its plastic stanchions.
5. Identify which memory RIMM you intend to replace.
6. There are locking levers on the end of each RIMM connector. Open the levers and gently pull the RIMM from the connector.

Replacement

1. When replacing a RIMM in an array, make sure that the size of the RIMM matches the other RIMMs in the array.
2. Insert the RIMM and press it into the connector. The two locking levers on the sides of the connector should close.
3. Secure the hold-down bracket.
4. Replace the CPU.

Verification

Power up the system.

Memory Map

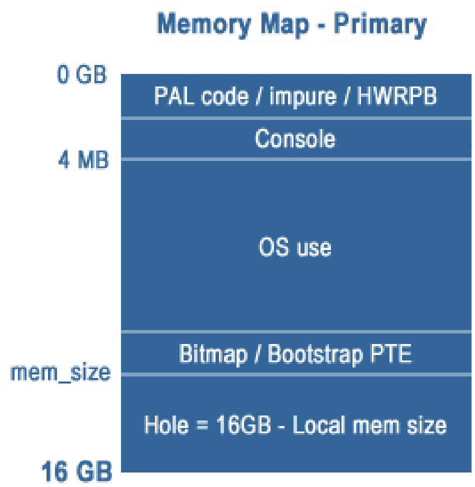

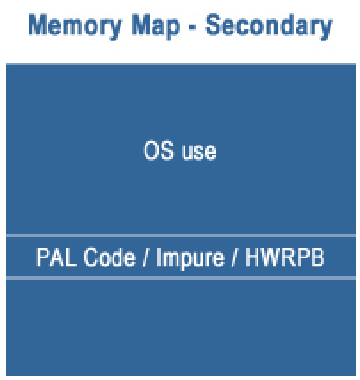

System Memory Maps

Not all CPUs have memory, however, those that do have their own predetermined memory contents described by the maps above. Important to notice is that each CPU, whether a primary or secondary, has its own copy of PALcode. This makes access to PALcode efficient. When a CPU does not have local memory it relies on the PALcode in its neighbor's memory.

Primary CPU Memory Map

Memory local to a primary CPU in a partition contains:

- A copy of PAL code/an impure area/the HWRPB
- a copy of the SRM console code
- space for operating system use
- Bitmaps used to define good/bad memory and bootstrap page table entries (PTEs)

Secondary CPU Memory Map

Memory local to a secondary CPU in a partition contains:

- space for operating system use
- Bitmaps used to define good/bad memory and a copy of PAL code/an impure area

Memory Address

The following elements are considered when the console determines memory addresses.

- Each CPU is always granted a minimum of 16GB of memory address space.
- A CPU may or may not have local memory.
- The size of a CPU's local memory is determined by RIMM size and, when memory is present, can range between 512MB and 32 GBs.

The architecture supports two addressing modes:

1. One supports a maximum of 256 processors each with 16GBs of memory
2. The other supports a maximum of 128 processors each with 32 GBs of memory

Each partition's virtual memory starting address is 0, though its physical address is determined by the location of the CPU to which the memory belongs in the mesh.

I/O

The hp AlphaServer ES47/ES80/GS1280 family has both standard and high performance I/O drawers.
These can be configured as a master I/O or an expansion I/O drawer. I/O drawers can only be connected to a CPU that has dedicated memory on the board.

1. A Master I/O Drawer consists of a Standard or High Performance I/O Drawer plus a Combination Adapter in one slot. The Combination Adapter provides an Ultra3 SCSI connection for two disks and a CD-RW drive in the drawer; plus a USB connection for keyboard, mouse, and monitor. Its components include:

- 10 PCI/PCI-X slots
- One AGP slot
- An N+1 redundant power system
- A CD-RW drive
- Two Ultra3 SCSI hot-plug disk drive bays (the disk drives must be ordered separately).

2. An expansion I/O drawer contains:

- 11 PCI/PCI-X slots
- One AGP slot
- An N+1 redundant power system

An I/O master drawer is a mandatory option for the GS1280.
A separately ordered I/O cable is required to connect an I/O drawer to a CPU drawer.

AlphaServer Model	Max I/O Drawer	I/O Drawer Type
ES47 M2	1	Expansion
ES47 M4	2	Expansion
ES80 M2	1	Expansion
ES80 M4	2	Expansion
ES80 M6	3	Expansion
ES80 M8	4	Expansion
GS12808 M8	8	Master or Expansion
GS1280 M16	16	Master or Expansion
GS1280 M32		Master or Expansion
GS1280 M64		Master or Expansion

-

Standard I/O Drawer

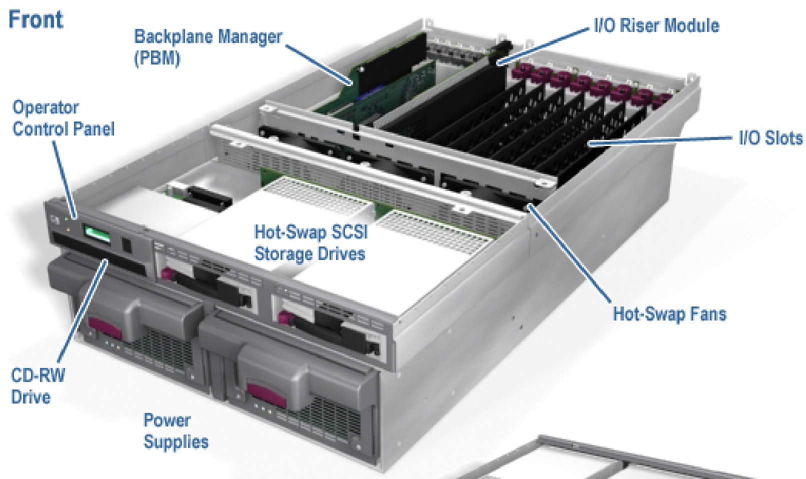

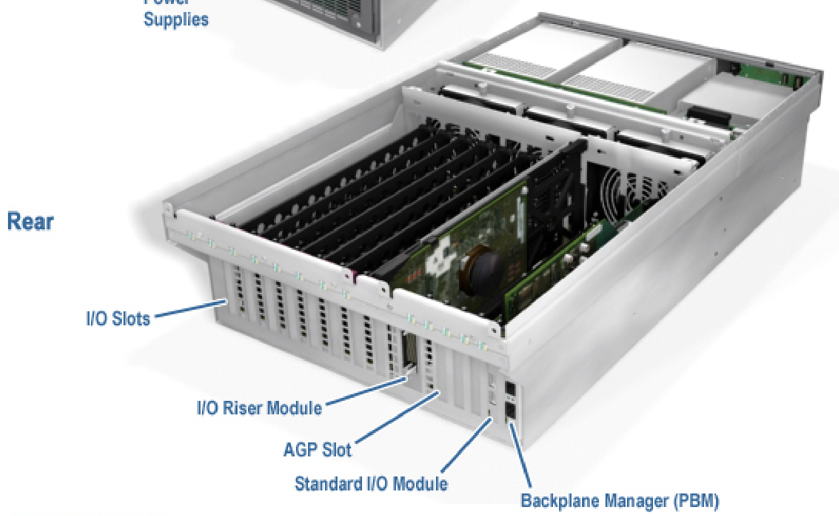

Block Diagram

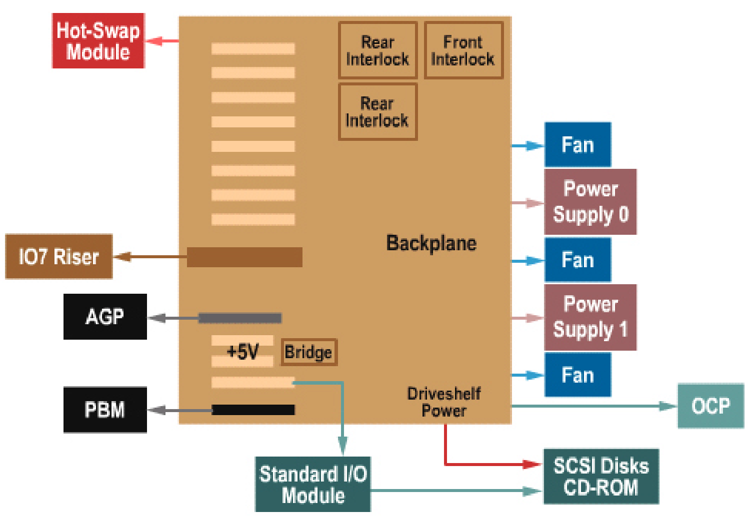

The standard I/O drawer contains an I/O riser module connected by cable to the EV7 system chip. The IO7 chip on the riser module controls a six-slot and a two-slot PCI/PCI-X bus, a three-slot PCI bus, and an AGP slot. Space is available for two SCSI storage drives and a CD-ROM; these components require a standard I/O module installed in the drawer.

Standard I/O Module

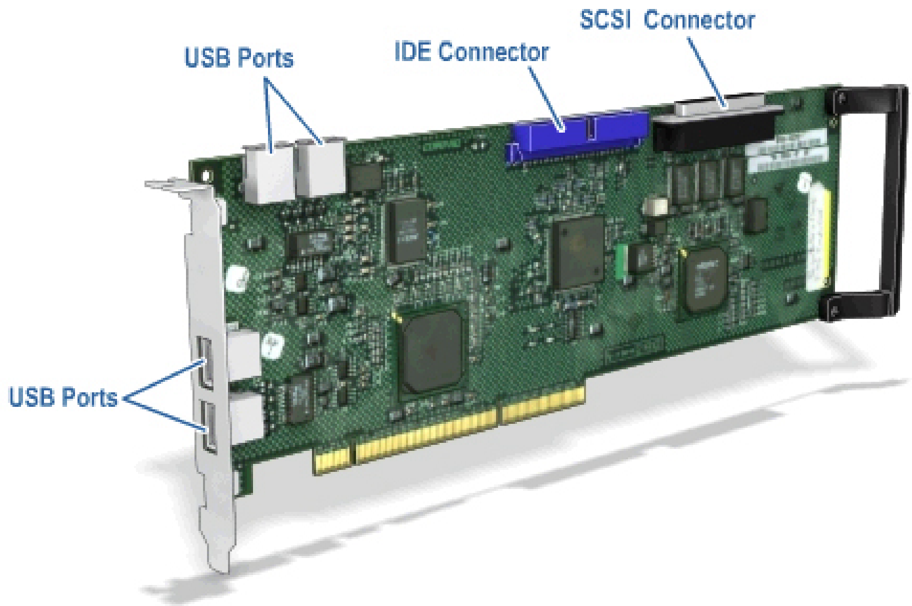

Block Diagram

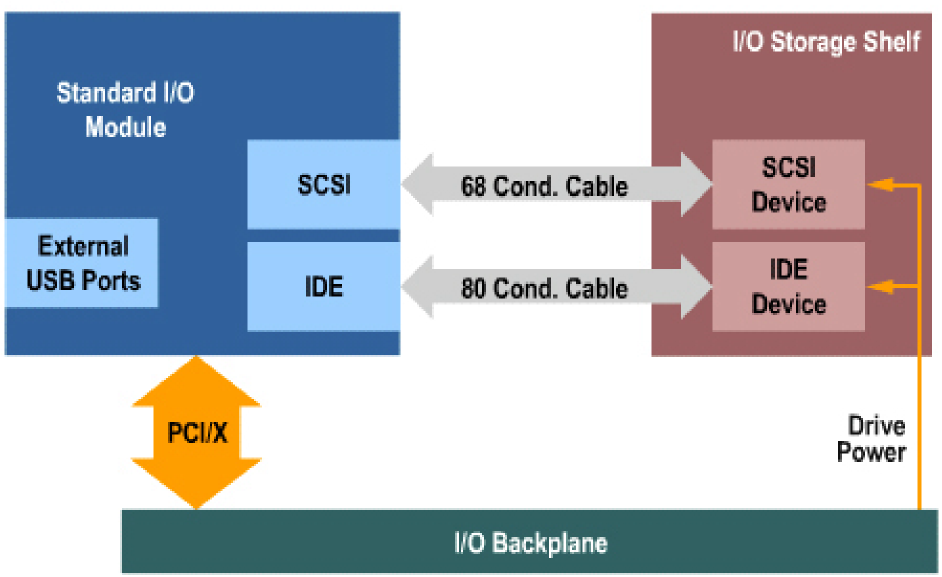

The standard I/O (SIO) module is an option available with the standard I/O box. It contains a SCSI controller, an IDE controller, and a USB controller. With the SIO, two SCSI drives and a CD-ROM may be used in the I/O drawer. Four USB connections, two internal and two external to the drawer become available for USB devices.

I/O Riser Module

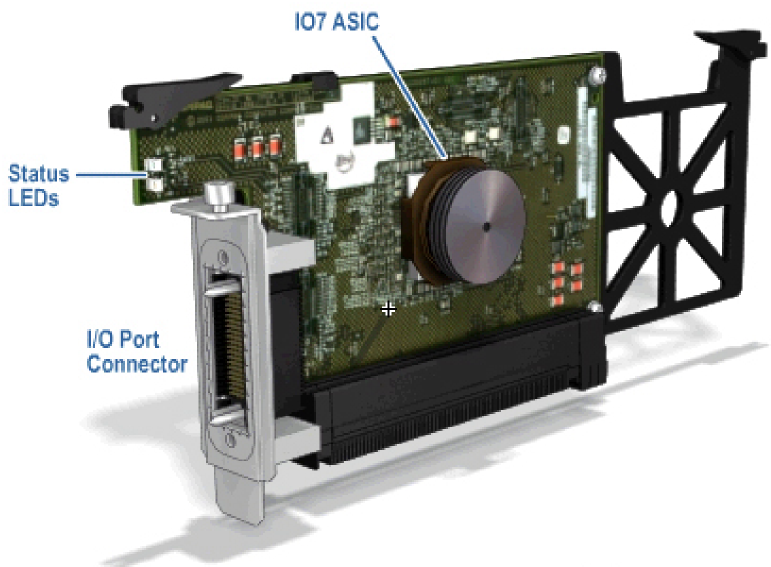

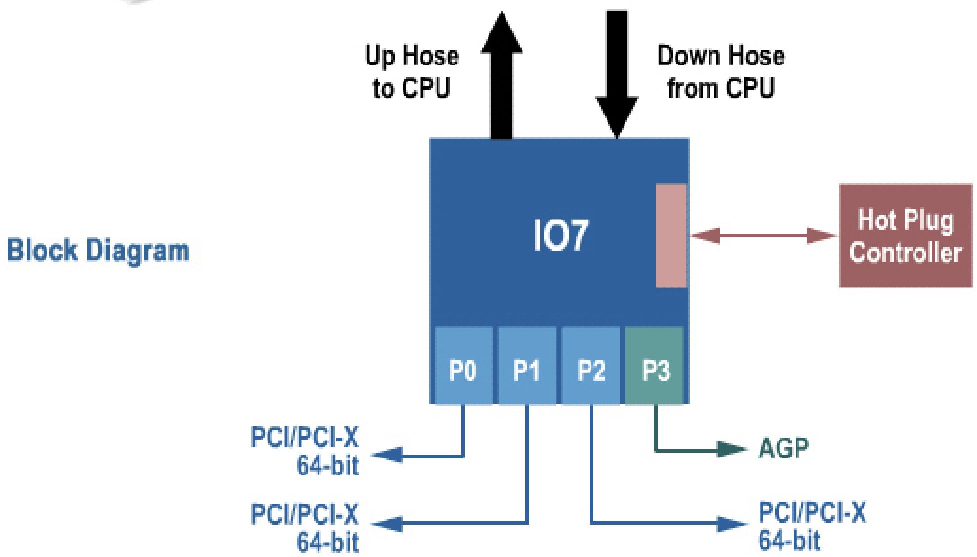

The I/O riser module connects by cable to an EV7 I/O port. The module functions as the interconnect between the EV7 chip and the I/O buses and the AGP bus. The module plugs into the standard I/O backplane.

I/O Backplane

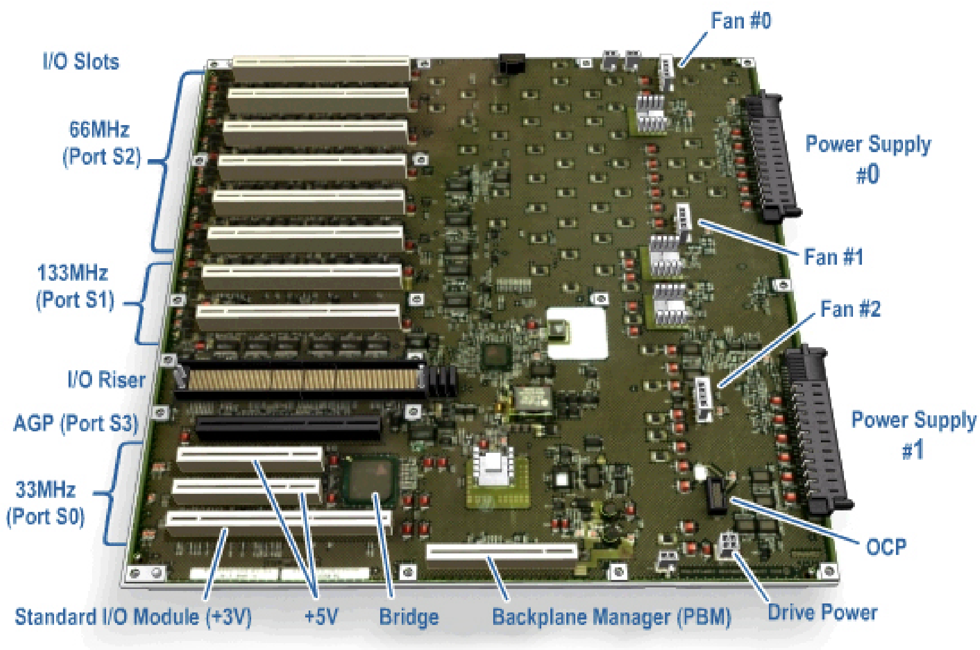

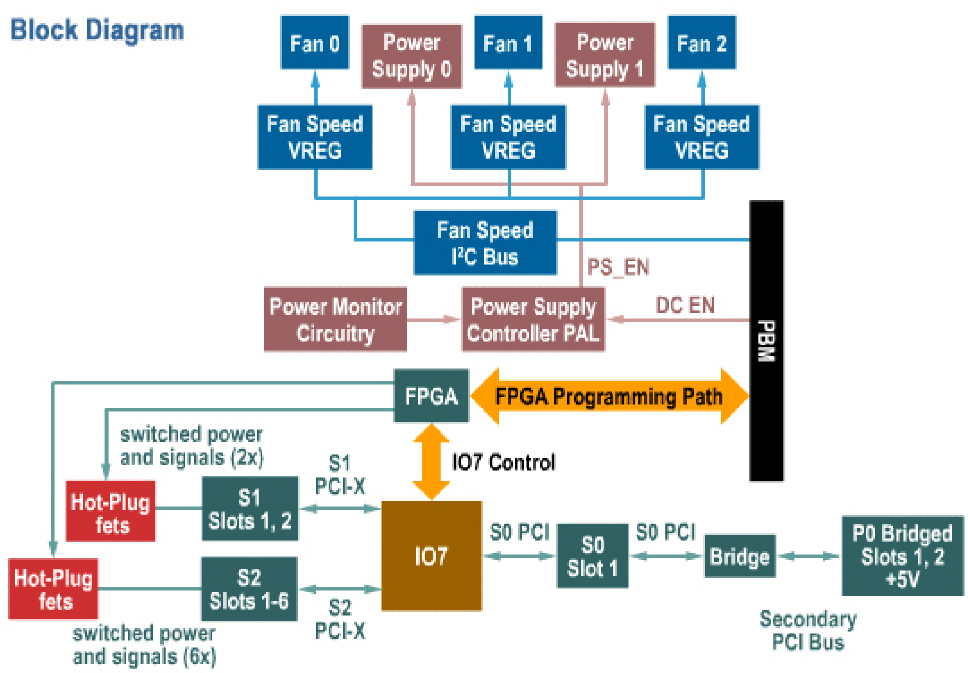

The standard I/O backplane has slots for the I/O backplane manager (PBM) module, an IO7 riser module, an accelerated graphics port (AGP), a two-slot and six-slot PCI/PCI-X bus, and a three-slot PCI bus. Bus speeds vary depending on the speed of the options and the size of the bus. The right-most slot of the three-slot PCI bus may be filled with a standard I/O module, the other two slots are dedicated to memory channel and CIPCA devices.

High Performance I/O Drawer

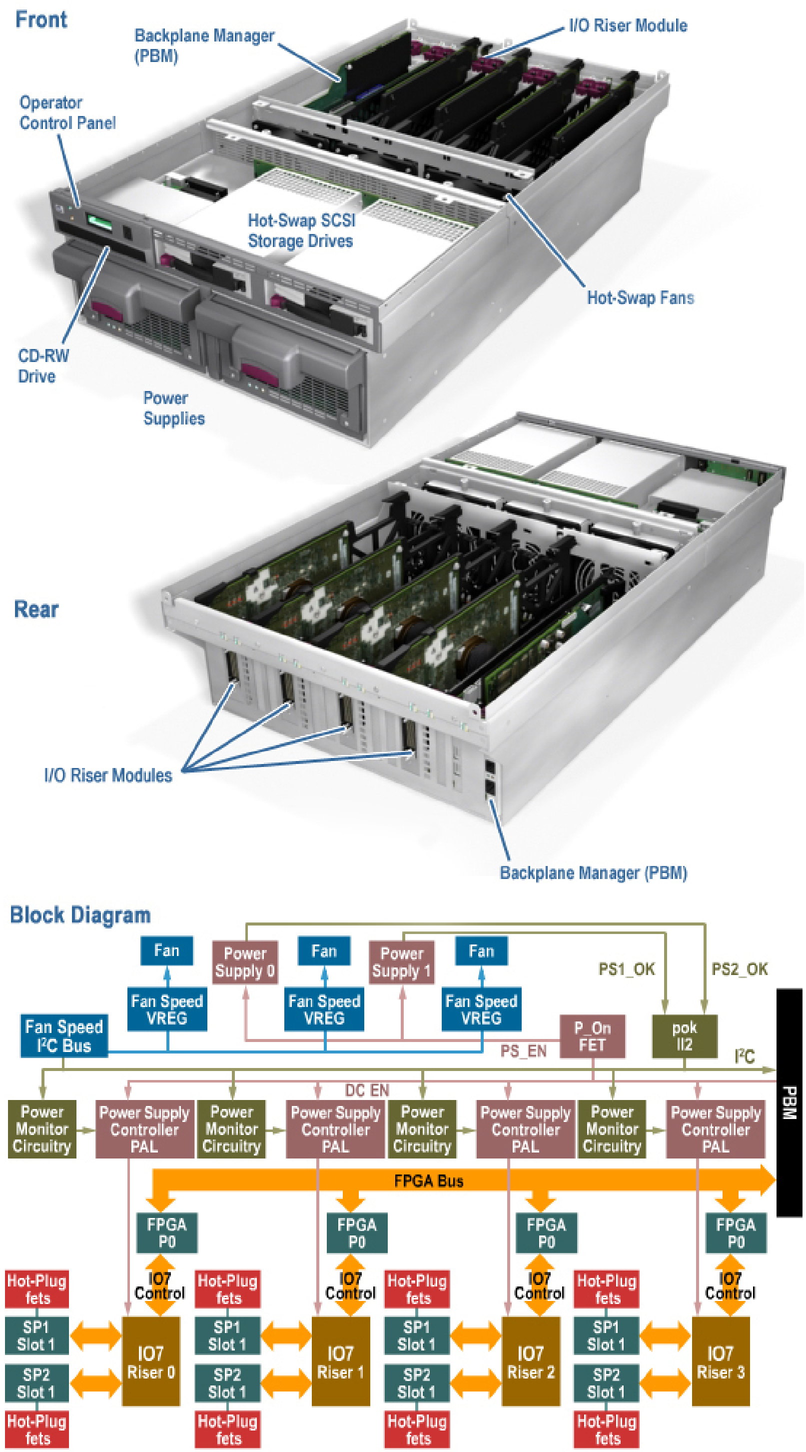

The high performance I/O drawer contains up to four I/O riser modules connected by cable to the I/O port of EV7 system chips. The IO7 chip on each riser controls two high speed single-slot PCI-X buses. The drawer, therefore, contains up to eight PCI-X buses

High Performance I/O Backplane

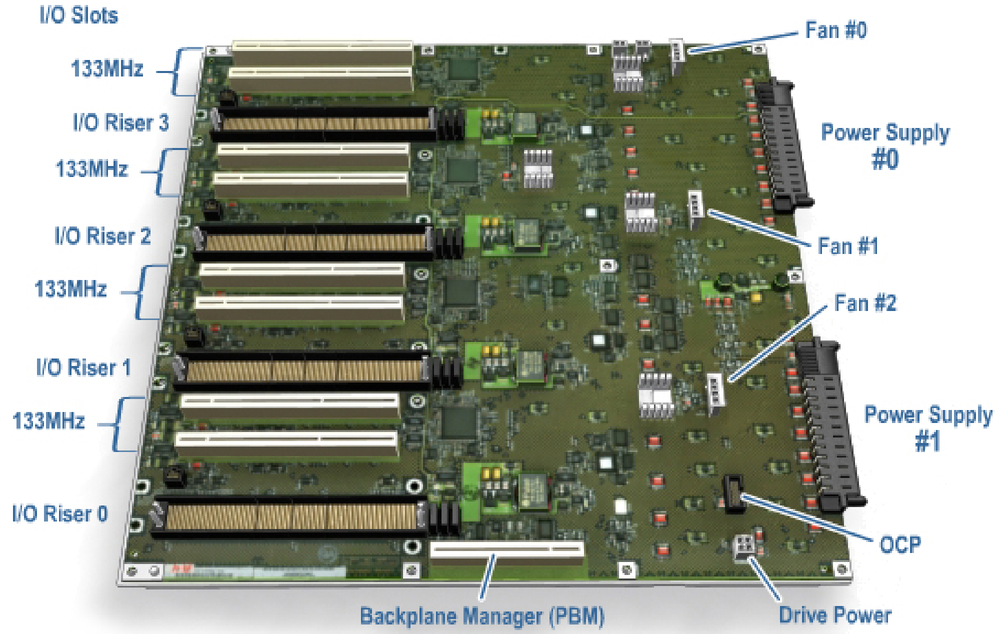

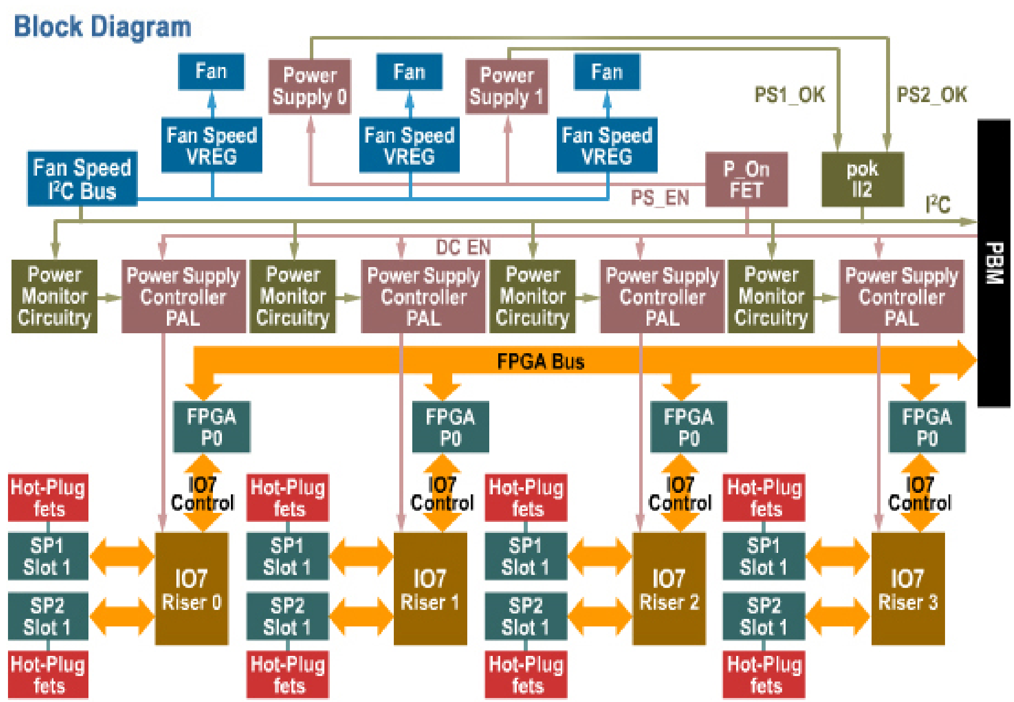

The high performance I/O drawer backplane has slots for the I/O backplane manager (PBM) module, four IO7 riser modules, and eight single slot PCIX busses running at 133MHz. Each IO7 supports two PCIX busses.

I/O Drawer Control Panel

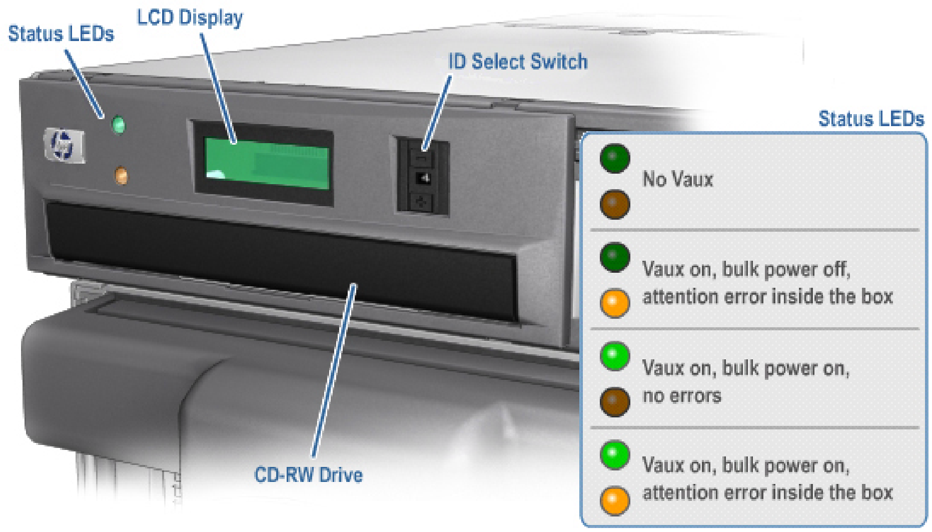

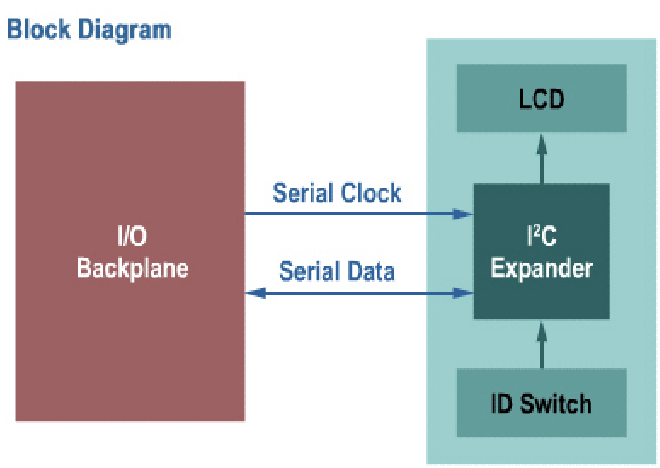

The I/O drawer control panel has a two-line 16 character
LCD display to show system status information.
The control panel has:

- A FAULT LED and an OK LED
- A pushbutton ID select switch used to set the I/O box ID number

Green OK LED	Amber FAULT LED	Description
Off	Off	No Vaux (or no connection to the MBM/PBM
Off	On	Vaux On, bulk power Off, attention error inside drawer
On	Off	Vaux On, bulk power On, no errors
On	On	Vaux On, bulk power On, attention error inside drawer

Storage Shelf

The universal StorageWorks shelf contains up to 14 SCSI disks, available at speeds of 10,000 or 15,000 RPM with capacities of 9.1, 18.2, or 36.4 Gigabytes. Each shelf requires a SCSI adapter.

Front

Rear

Power

Power specifications and illustrations for the system cabinets will help with the design of power capacity and distribution.

> *NOTE: HP recommends designing or preparing a computer site in accordance with NFPA 75: Protection of Electronic Computer / Data Processing Equipment.*
>
> *NOTE: HP recommends designing or preparing the power distribution system in accordance with IEEE 1100-1999: Recommended Practice for Powering and Grounding of Electronic Equipment (IEEE Emerald Book) and any country specific electrical codes.*

The power source should be independent of all other loads; meaning, it should not supply air-conditioners, convenience outlets, lighting, or any other potentially noisy loads. The power source should be an isolation transformer (with electrostatic shield), located in close proximity to the proposed system that has sufficient capacity to support the existing loads, the proposed system, future expansion, and inrush currents.

All aspects of the power distribution system must comply with the minimum standards set forth by the National, State, or local electrical codes.

Today's computer equipment power supplies are tolerant of minor sags and surges on the power line. Many even feature built-in surge suppression. However, on occasion the power is distorted enough to cause operational problems. It is the customer's responsibility to analyze the power quality and determine the most appropriate solution. HP offers services that can assist with this effort.

All receptacles must be derived from dedicated branch circuits that include a grounding conductor. All receptacles must be standard grounding-type receptacles. HP does not recommend the use of isolated grounding (IG) type receptacles except where proven necessary. Branch circuits should not exceed 75 feet in length wherever possible to help minimize ground differential voltages that can upset system operation.

ES47/ES80 System Power

- Three-phase PDU
- Single-phase PDU
- Heat Dissipation
- Airflow

	North America	***Japan***	***Europe***
Three-phase PDU	3X-H7606-AA	3X-H7606-AA	3X-H7606-AB
Nominal Voltage	200-208	200-208	380-415
Rated Current	24A	24A	24A
Frequency	50-60 Hz	50-60 Hz	50-60 Hz
Phases	3W+N+G	3W+N+G	3W+N+G
Kva, Model 2 (1 I/O Expansion Drawer, 9 StorageWorks Shelves)	6.683	6.683	6.683
Kva, Model 4 (2 I/O Expansion Drawers, 6 StorageWorks Shelves)	6.166	6.166	6.166
Kva, Model 6 (2 I/O Expansion Drawers, 5 StorageWorks Shelves)	6.349	6.349	6.349
Kva, Model 8 (2 I/O Expansion Drawers, 3 StorageWorks Shelves)	5.932	5.932	5.932
Line connection	Fixed cord & plug	Fixed cord & plug	Fixed cord & plug
Rating	10/7A per cord	10/7A per cord	10/7A per cord
	3 x 12 AWG	3 x 12 AWG	3 x 2.5mm
Power plug (site)	L21-30P, Hubbell 2811	L21-30P, Hubbell 2811	Hubbell 532P6W
Main breaker	30A	30A	30A

Sub breakers	3x20A(2P, 1x20A(3P)	3x20A(2P, 1x20A(3P)	3x20A(2P, 1x20A(3P)
Single-phase PDU	3X-H7609-EB	3X-H7609-EB	3X-H7609-DB
Voltage	200-240	200-240	200-240
Rated Current	16A	16A	16A
Frequency	50-60 Hz	50-60 Hz	50-60 Hz
Kva	3.2-3.84	3.2-3.84	3.2-3.84
Current per phase	17.5A	17.5A	17.5A
Power outlets (internal)	3 x C19, 24 x C13	3 x C19, 24 x C13	3 x C19, 24 x C13
Line connection	Fixed cord & plug	Fixed cord & plug	Fixed cord & plug
Power cord	3 x 12 AWG	3 x 12 AWG	3 x 2.5 mm^2
	NEMA L6-L20P	NEMA L6-L20P	IEC 309
Main breaker	20A (2)	20A (2)	20A (2)
Sub breakers	N/A	N/A	N/A

Heat Dissipation

Minimally configured system, (Model 2 only)	894W	894W	894W
Btu/hr (Model 2 only)	3,051	3,051	3,051
Fully configured system, (Model 2 only)	1,930W	1,930W	1,930W
Btu/hr (Model 2 only)	6,587	6,587	6,587
Minimally configured system, Model 4 only)	1,788W	1,788W	1,788W
Btu/hr (Model 4 only)	6,102	6,102	6,102

Fully configured system, Model 4 only)	3,860W	3,860W	3,860W
Btu/hr (Model 4 only)	13,174	13,174	13,174
Minimally configured system, Model 6 only)	2,682W	2,682W	2,682W
Btu/hr (Model 6 only)	9,153	9,153	9,153
Fully configured system, Model 6 only)	4,596W	4,596W	4,596W
Btu/hr (Model 6 only)	15,686	15,686	15,686
Minimally configured system, Model 8 only)	3,576W	3,576W	3,576W
Btu/hr (Model 6 only)	12,204	12,204	12,204
Fully configured system, Model 8 only)	5,928W	5,928W	5,928W
Btu/hr (Model 6 only)	20,230	20,230	20,230

Airflow, cfm

Model 2	Minimum	227
	Maximum	492
Model 4	Minimum	554
	Maximum	984
Model 6	Minimum	831
	Maximum	1476
	Minimum	1108
	Maximum	1968

Note: *Maximum heat output is generated with full memory installed in 2P drawer and I/O optimized for full performance. Typical output assumes I/O drawers are optimized for connectivity.*

GS1280 System Power

One power subrack and one AC input box are required for each 8P drawer in the system.

Power Subrack

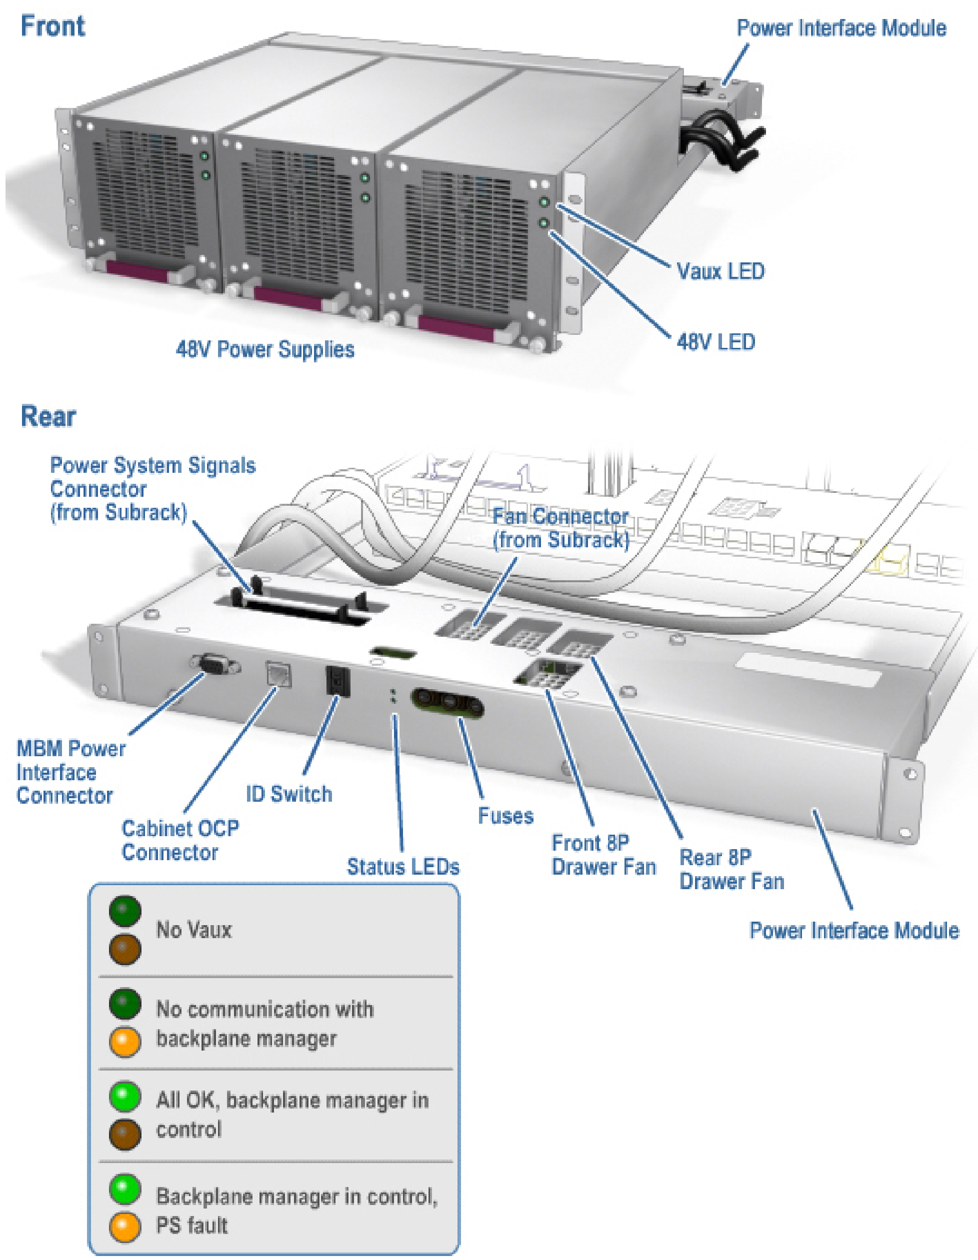

A power subrack supplies power to an 8P drawer. It holds three 48V power supplies; two power supplies provide power to the drawer, the third is redundant. The subrack enables the power supplies to operate in parallel and to share the load. Power is distributed to the 8P drawer by a cable that runs from the subrack to the drawer backplane.

A power interface module monitors the subrack power and provides connectors for system drawer fans, the cabinet control panel, and the MBM power interface. The module also contains status LEDs, an ID switch, and fuses.

AC Input Box

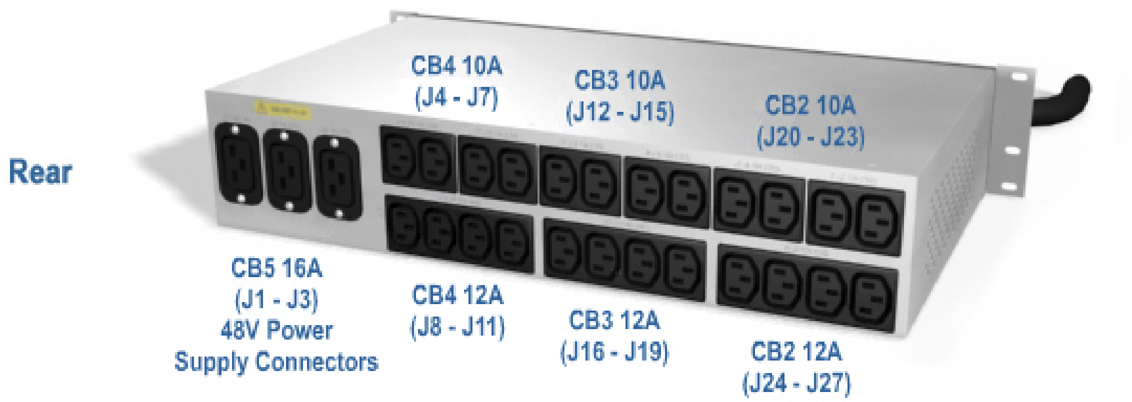

The AC input box (or power distribution unit (PDU)) distributes power to a power subrack. One AC input box supports one power subrack in the system. AC input box(es) are located at the bottom of the cabinet in 8P and 16P systems.

Dual AC

Model 8 Dual AC

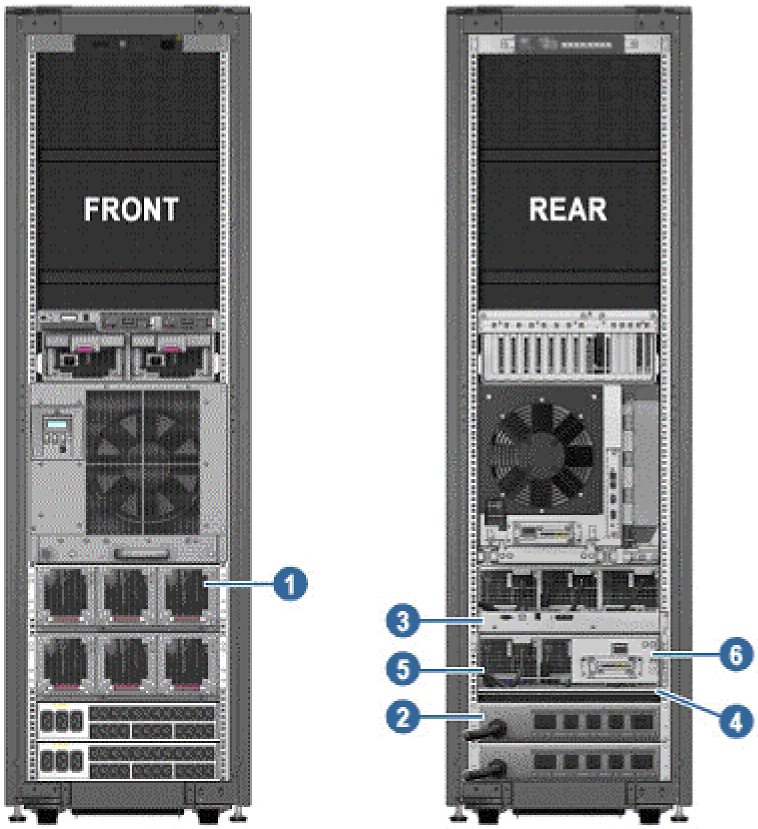

The dual AC power option provides the capability to connect an hp AlphaServer GS1280 (Model 8 or Model 16) to two separate AC feeds. If you did not originally order dual AC, you can order a dual AC upgrade to be installed.

The feeds can be direct from the power utility, or they can be a combination of utility feeds and UPS or generator feeds. During normal operation, the system will receive power from both feeds. If one feed fails, the system will continue to receive power from the remaining feed with no interruption in system operation.

The figure above shows the dual AC set-up for a Model 8. Callouts mark:

1. Second set of 48Vpower supplies
2. Second 30-amp Power Distribution Unit (PDU)
3. Power interface module (WPI) (must be moved up to make room)
4. Cable wire frame
5. Dual AC chassis
6. Dual AC junction box

Model 16 Dual AC

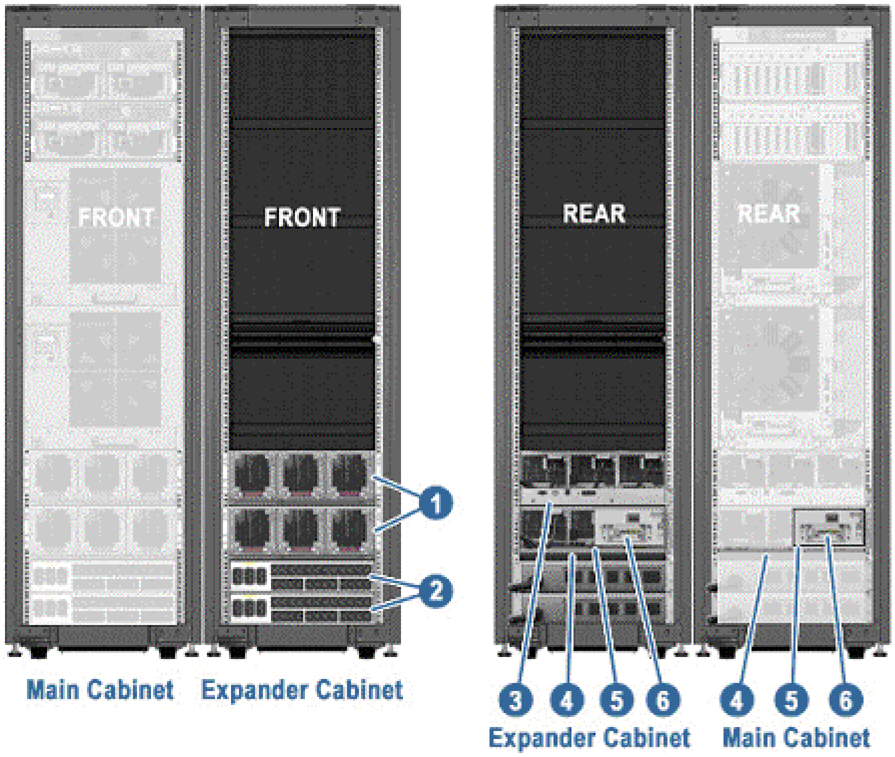

Here is the dual AC set-up for a Model 16. Callouts mark:

1. Two new sets of 48V power supplies
2. Two new 30-amp Power Distribution Units (PDUs)
3. Lower power interface module (PI) in main cab moved to expander cab
4. Two new cable wire frames: one in main cab, one in expander cab
5. Two new dual AC chassis: one in main cab, on in expander cab
6. Two new dual AC junction boxes, one in main cab, one in expander cab

GS1280 Model 8 System Power

- Power
- Heat Dissipation
- Airflow

Power	US/Canada	Japan	Europe
Voltage	120/208	200	380–415
Rated Current	24A	24A	24A
Phase	3W+N+G	3W+G	3W+N+G
Frequency	50-60 Hz	50-60 Hz	50-60 Hz
Kva, Model 8, 1 I/O Drawer, 4 StorageWorks Shelves	6.092	6.092	6.092
Current per phase	17.5A	18A	9.5A
Kva, Model 8, 2 I/O Drawers, 4 StorageWorks Shelves	5.982	5.982	5.982
Current per phase	17A	17.8A	9.5A
Power outlets (Internal)	3 x C19, 24 x C13	3 x C19, 24 x C13	3 x C19, 24 x C13
Line connection	Fixed cord & plug	Fixed cord & plug	Fixed cord & plug
Power cord	5 x 10AWG	5 x 10AWG	5 x 4 mm^2
Power plugs	L21-30P, Hubbell 2811	L21-30P, Hubbell 2811	Hubbell 532P6W
Number of PDUs required	1, 2 for dual AC	1, 2 for dual AC	1, 2 for dual AC
Main breaker	30A	30A	30A
Sub-breakers	3x20A(2p), 1x20A(3p)	3x20A(2p), 1x20A(3p)	3x15A(2p), 1x15A(3p)

Heat Dissipation			
Max heat output,		3,405W	
Max heat output, Btu/hr	11,622	11,622	11,622
Typical heat output	2,430W	2,430W	2,430W
	8,297	8,297	8,297

	6,405W	6,405W	6,405W
	21,861	21,861	21,861
	3,330W	3,330W	3,330W
	11,366	11,366	11,366
Airflow, cfm			
	850	850	850
	1,210	1,210	1,210

Note: *Maximum heat output is generated with a full capacity of CPUs and memory installed in the 8P drawer, and I/O optimized for full performance. Typical heat output assumes one half of the CPU slots are used and the I/O drawer is optimized for connectivity.*

GS1280 Model 16 System Power

Power	US/Canada	Japan	Europe
Voltage	120/208	200	380
Rated Current	24A	24A	24A
Phase	3W + N + G	3W + G	3W + N + G
Frequency	50-60 Hz	50-60 Hz	50-60 Hz
Kva, Model 16, 1 I/O Drawer, 1 StorageWorks Shelf	6.193	6.193	6.193
Current, per phase	19A	20A	10.5A
Kva, Model 16, 2 I/O Drawers, 0 StorageWorks Shelves	18.5A	18.5A	18.5A
Power outlets (Internal)	3 x C19, 24 x C13	3 x C19, 24 x C13	3 x C19, 24 x C13
Line connection	Fixed cord and plug	Fixed cord and plug	Fixed cord and plug
Power cord	5 x 10AWG	5 x 10AWG	5 x 4 mm^2
Power plugs	L21-30P, Hubbell 2811	L21-30P, Hubbell 2811	Hubbell 532P6W
Number of PDUs required	2, 4 for dual AC	2, 4 for dual AC	2, 4 for dual AC
Main breaker	30A	30A	30A
Sub-breakers	3 x 20A(2p), 1x20A(3p)	3 x 20A(2p), 1 x 20A(3p)	3 x 15A(2p), 1 x 15A(3p)
Heat Dissipation			
Maximum heat output	6,810W	6,810W	6,810W
Maximum heat output, Btu/hr	23,243	23,243	23,243
Typical heat output	3,700W	3,700W	3,700W
Typical heat output, Btu/hr	12,632	12,632	12,632
Airflow, cfm			
Airflow, cfm, full rack, minimum	1,500	1,500	1,500
Airflow, cfm, full rack, maximum	2,200	2,200	2,200

Note: *Maximum heat output is generated with a full capacity of CPUs, memory installed in the 8P drawers, and I/O optimized for full performance. Typical heat output assumes one half of the CPU slots are used and the I/O drawer is optimized for connectivity.*

GS1280 Model 32 System Power

Power	US/Canada	Japan	Europe
Voltage	120/208	200	380/415
Current	24A	24A	24A
Phase	3W + N + G	3W + N + G	3W + N + G
Frequency	50-60 Hz	50-60 Hz	50-60 Hz
Kva, Model 32, 1 I/O Drawer, 1 StorageWorks Shelf	10.990	10.990	10.990
Current per phase	11.7A	11.7A	11.7A
Power outlets (internal)	3x C19, 24 x C13	3 x C19, 24 x C13	3 x C19, 24 x C13
Number of PDUs required	4, 8 for dual-AC	4, 8 for dual-AC	4, 8 for dual-AC
Line connection	Fixed cord and plug	Fixed cord and plug	Fixed cord and plug
Power cord	5 x 10AWG	5 x 10AWG	5 x 4mm^2
Power plug	L21-30P, Hubbell 2811	L21-30P, Hubbell 2811	Hubbell 563P6W
Main breaker	30A	30A	30A
Sub-breakers	3x20A(2p), 1x20A(3p)	3x20A(2p), 1x20A(3p)	3x15A(2p), 1x15A(3p)
Heat Dissipation			
Max heat output	14,820W	14,820W	14,820W
Max heat output, Btu/hr	50,581	50,581	50,581
Typical heat output	8,000W	8,000W	8,000W
Typical heat output, Btu/hr	27,304	27,304	27,304
Airflow, cfm			
Airflow cfm, minimum full rack, CPU rack	3,400	3,400	3,400
Airflow cfm, maximum full rack, CPU rack	1,800	1,800	1,800
Airflow cfm, minimum full rack, power rack	1,685	1,685	1,685
Airflow cfm, maximum full rack, power cab	1,455	1,455	1,455

Note: *Maximum heat output is generated with a full capacity of CPUs and memory installed in the 8P drawers, and I/O optimized for full performance. Nominal heat output assumes one half of the CPU slots are used and the I/O drawer is optimized for connectivity.*

GS1280 Model 64 System Power

Power	US/Canada	Japan	Europe
Voltage	120/208	200	380/415
Current	24A	24A	24A
Phase	3W + N + G	3W + N + G	3W + N + G
Frequency	50-60 Hz	50-60 Hz	50-60 Hz
Kva, Model 64, 1 I/O Drawer, 1 StorageWorks Shelf	10.990	10.990	10.990
Current per phase	10.6A	10.9A	6.8A
Power outlets (internal)	3x C19, 24 x C13	3 x C19, 24 x C13	3 x C19, 24 x C13
Number of PDUs required	8	8	8
Line connection	Fixed cord and plug	Fixed cord and plug	Fixed cord and plug
Power cord	5 x 10AWG	5 x 10AWG	5 x 4mm^2
Power plug	L21-30P, Hubbell 2811	L21-30P, Hubbell 2811	Hubbell 563P6W
Main breaker	30A	30A	30A
Sub-breakers	3x20A(2p), 1x20A(3p)	3x20A(2p), 1x20A(3p)	3x15A(2p), 1x15A(3p)
Heat Dissipation			
Max heat output	28,440W	28,440W	28,440W
Max heat output, Btu/hr	97,066	97,066	97,066
Typical heat output	14,800W	14,800W	14,800W
Typical heat output, Btu/hr	50,513	50,513	50,513
Airflow, cfm			
Airflow, minimum full rack, CPU rack	6,800	6,800	6,800
Airflow, maximum full rack, CPU rack	3,600	3,600	3,600
Airflow, minimum full rack, power rack	1,324	1,324	1,324
Airflow, maximum full rack, power cab	1,508	1,508	1,508

Note: Maximum heat output is generated with a full capacity of CPUs and memory installed in the 8P drawers, and I/O optimized for full performance. Nominal heat output assumes one half of the CPU slots are used and the I/O drawer is optimized for connectivity.

Grounding Requirements

Two grounding systems are required: a safety grounding system that meets national, state, and local electrical codes, and a high frequency grounding system for noise reduction.
A safety ground is required. Grounding cabinets together is important in creating a stable electrical environment throughout the system. Safety ground must conform to the following specifications:

- The safety ground must be completely isolated from neutral all the way back to the power transformer.
- The safety ground must never be dependent on a conduit alone. A wire must be run into the electrical feed along with the power wiring and neutral for use as safety ground.
- The safety ground and neutral must be firmly connected together at the power transformer.
- The safety ground must be firmly wired to the unit for personnel safety and to ensure that the AC line filters properly function.

Powering Up

Before powering up the system, make sure the keyswitch is off, and then turn on the circuit breakers in the system cabinet(s) and expansion cabinets, if necessary. Then, set the keyswitch to On to power up the system (the keyswitch must be set to On to power up the system locally or remotely). The power-up display is shown at the system management console and the control panel.

This power-up display shows an 8-processor system, each with 1GB of memory, and one I/O subsystem connected to CPU0.

Cold Start - Power Switch Off

The following is an example of the serial line output of a machine that has its AC power applied and its keyswitch in the off position. Things you might want to notice in the power up are in **bold text.**

<AC power is turned on....>

```
00 01 02 03 04 05 06 07 08 09 Attaching interface lo0...done
                 VxWorks            Embedded OS in the micro processors
that make up the server management network.
Copyright 1984-1998  Wind River Systems, Inc.
           CPU: AMD SC520 CDP
        VxWorks: 5.4.2
    BSP version: 1.2/0
  Creation date: May  7 2003
            WDB: Ready.

GS1280 Server Management Failsafe Loader V2.1-1 Starting up

Image built on May  7 2003 at 10:37:56
Cabinet number: 00
Drawer number: 0
Micro type: MBM
Node IP address: 10.0.0.1                        Address of the MBM in
cabinet 0 drawer 0Press enter to remain in FSL.
boot device          : flash
unit number          : 0
processor number     : 0
host name            : host
file name            : vxWorks
inet on ethernet (e) : 10.250.250.250
host inet (h)        : 10.253.0.254
user (u)             : target
flags (f)            : 0xa0
other (o)            : fei
06 07 08 09 Attaching interface lo0...done
Adding 5756 symbols for standalone.
                 VxWorks
Copyright 1984-1998  Wind River Systems, Inc.
            CPU: AMD SC520 CDP
        VxWorks: 5.4.2
    BSP version: 1.2/0
  Creation date: May  8 2003
            WDB: Ready.

GS1280 Server Management  V2.1-8 Starting up
```

```
        Image built on May  8 2003 at 14:26:45
        Running POST ...
        0A 0B 0C 0D 0E
Cabinet number: 00
Drawer number: 0
Micro type: MBM
Node IP address: 10.0.0.1
        ~GRP-W-(tRootTask) Joined Group ID: 100000a.1    The microprocessor
        polls for other processors on the network and forms a group.

        Forming group~GRP-W-(grp_Monitor) Leaving Group ID: 100000a.1
        ~GRP-W-(grp_Monitor) Joined Group ID: 201fe0a.15
        ........interrupt: ~GRP-I-(interrupt), GROUP HAS FORMED ID:201fe0a.15
        .........interrupt: ~GRP-I-(interrupt), GROUP IS STABLE ID:201fe0a.15

        Warning: No DHCP server address cache! Later entries will not be saved.
        DHCP server started.
        Pco_task started as pco_00
        pco_task started as pco_01
        pco_task started as pco_02
        pco_task started as pco_03
        pco_task started as pco_04
        Welcome – GS1280 Server Manager – V2.1-8         Server Management
        firmware up and running.
        [2003/06/10 11:17:16]
        ~REC-W-(trecTask) Server management group is transitioning.
        [2003/06/10 11:17:16]
        ~REC-W-(trecTask) Server management group is stable.
        Starting telnet port on port:323
        Starting telnet daemon on port:323
        Starting telnet port on port:324
        Starting telnet daemon on port:324
        Starting telnet port on port:325
        Starting telnet daemon on port:325
        MBM Init finished at: TUE JUN 10 11:17:29 2003

        0x186a7e0 (mbm_dhcp): dhcps: read 0 entries from binding and addr-pool
        database.

MBM>    End of serial port output.
```

Power Switch On

In the three examples that follow:

- The first shows hard partition P0 bringing up the console.
- The second and third show two sub-partitions in hard partition P1 bringing up their consoles.

In each case the console is connected through a telnet session to the server management network.

Example 1   Hard partition P0

~NET-I-(td323) Session opened to port 323, fd 62, by host 192.168.2.1, port 2807.

```
ÿû~PCO-I-(pco_02) Powering on partition. HP: WEBSHOOTER89
~PCO-I-(pco_02)
```

(The following map shows the mesh of hard partition H0. Note that the map shows the IP connections between processors, both North–South and East-West, and indicates slots that have processors and slots that have fillers.)

Configuring for 8 CPUs for HP:0 WEBSHOOTER89

```
       0  1  2  3  4  5  6  7  8  9  A  B  C  D  E  F
      .w..w..w..w....................................
  0    .P--F--P--F....................................
      .|..|..|..|....................................   P  Processor
      .|..|..|..|....................................   F  Filler
  1    .P--F--P--F....................................   |  connection
      .|..|..|..|....................................   -  connection
      .|..|..|..|....................................
  2    .P--F--P--F....................................    W <----------> E
      .|..|..|..|....................................  N ^ (0,0)
      .|..|..|..|....................................    |  .
  3    .P--F--P--F....................................    |    .
      .w..w..w..w....................................    |       .
      ...............................................  S v      (ns,ew)
  4    ...............................................
      ...............................................
      ...............................................
  5    ...............................................
      ...............................................
      ...............................................
  6    ...............................................
      ...............................................
      ...............................................
  7    ...............................................
      ...............................................
       0  1  2  3  4  5  6  7  8  9  A  B  C  D  E  F
Running test 10, Initialize RAMBUS ... on 8 EV7s       Diagnostic tests start.
Running test 11, Initialize Memory ... on 8 EV7s
Running test 12, Data Pattern March read/write ... on 8 EV7s
Running test 13, RAID channel Test ... on 8 EV7s
Running test 14, Single Bit Error ... on 8 EV7s
Running test 15, Double Bit Error ... on 8 EV7s
Running test 20, Init IO7 and Start Clocks ... on 3 EV7s
Running test 21, IO7 Data Path (Scratch CSR) ... on 3 EV7s
Running test 22, IO Single Bit Error checkers ... on 3 EV7s
Running test 23, IO Double Bit Error checkers ... on 3 EV7s
Running test 24, IO Timer Expirations ... on 3 EV7s
Running test 25, IO up-hose SBE checkers ... on 3 EV7s
Running test 26, IO up-hose DBE checkers ... on 3 EV7s
Running test 27, IO7 pass2 data mover test ... on 3 EV7s
Running test 30, Configure RBOX Routes ... on 8 EV7s
Running test 31, Clear Errors / Enable Routes ... on 8 EV7s
Running test 32, Route Test: N S E W ... on 8 EV7s (North)
Running test 32, Route Test: N S E W ... on 8 EV7s (South)
```

```
Running test 32, Route Test: N S E W ... on 8 EV7s (East )
Running test 32, Route Test: N S E W ... on 8 EV7s (West )
Running test 33, Inverse Route Setup ... on 8 EV7s
Running test 32, Route Test: N S E W ... on 8 EV7s (North)
Running test 32, Route Test: N S E W ... on 8 EV7s (South)
Running test 32, Route Test: N S E W ... on 8 EV7s (East )
Running test 32, Route Test: N S E W ... on 8 EV7s (West )
Running test 33, Inverse Route Setup ... on 8 EV7s
Running test 34, Single Bit Error checker ... on 8 EV7s
Running test 35, Double Bit Error checker ... on 8 EV7s
Running test 31, Clear Errors / Enable Routes ... on 8 EV7s
Running test 16, Interprocessor Memory Access ... on 8 EV7s
Running test 40, Local I/O Device Interrupts ... on 8 EV7s
Running test 41, Local Interval Timer Interrupts ... on 8 EV7s
Running test 42, Local Interprocess Interrupts ... on 8 EV7s
Running test 43, Software Alerts ... on 1 EV7s
Running test 46, Other Local Interrupt Bits ... on 8 EV7s
~PCO-I-(pco_02) HP:WEBSHOOTER89 SP:Default_SP Primary is NS:0 EW:0 which is cab:00 drw:0
cpu:0  Primary Identified
Running test 50, Loop on Secondary Routine ... on 7 EV7s (SP:Default_SP)
Running test 50, Loop on Secondary Routine ... on 0 EV7s (SP:Free_Pool)

initialized idle PCB
initializing semaphores
initializing heap
initial heap 700c0
memory low limit = 54c000 heap = 700c0, 1fffc0
initializing driver structures
initializing idle process PID
initializing file system
initializing timer data structures
lowering IPL
CPU 0 speed is 1050 MHz
create dead_eater
create poll
create timer
create powerup
entering idle loop
access NVRAM
```

hpcount = 2, spcount = 5, ev7_count = 32, io7_count = 7   We are looking at a 32P system with 2 HP and 5 SP
hard_partition = 0   this is hard partition 0

```
IO7-100 (Pass 2) at PID 8
IO7 North port speed is 210 MHz
Hose 33 - 133 MHz PCI-X
Hose 34 - 133 MHz PCI-X
Hose 35 - 2X AGP
IO7-100 (Pass 2) at PID d
IO7 North port speed is 210 MHz
Hose 53 - 133 MHz PCI-X
Hose 54 - 133 MHz PCI-X
Hose 55 - 2X AGP
IO7-100 (Pass 2) at PID 0
IO7 North port speed is 210 MHz
Hose 0 - 33 MHz PCI
Hose 1 - 33 MHz PCI
Hose 2 - 33 MHz PCI
Hose 3 - 2X AGP
0 sub-partition 0:   start:00000000 00000000   size:00000000 40000000
```
PID 0 console memory base: 0, **1 GB**
```
1 sub-partition 0:   start:00000004 00000000   size:00000000 40000000
```
PID 1 memory: 400000000, **1 GB**
```
2 sub-partition 0:   start:00000020 00000000   size:00000000 40000000
```
PID 4P memory: 2000000000, **1 GB**
```
3 sub-partition 0:   start:00000024 00000000   size:00000000 40000000
```
PID 5 memory: 2400000000, **1 GB**

```
4 sub-partition 0:   start:00000040 00000000   size:00000000 80000000
PID 8 memory: 4000000000, 2 GB
5 sub-partition 0:   start:00000044 00000000   size:00000000 80000000
PID 9 memory: 4400000000, 2 GB
6 sub-partition 0:   start:00000060 00000000   size:00000000 80000000
PID 12 memory: 6000000000, 2 GB
7 sub-partition 0:   start:00000064 00000000   size:00000000 80000000
PID 13 memory: 6400000000, 2 GB
total memory, 12 GB   of sub-partition 0 in hard partition 0
probe I/O subsystem
probing hose 0, PCI
probing PCI-to-PCI bridge, hose 0 bus 2
do not use secondary IDE channel on CMD controller
probing PCI-to-PCI bridge, hose 0 bus 3
bus 2, slot 0, function 0 -- usba -- USB
bus 2, slot 0, function 1 -- usbb -- USB
bus 2, slot 0, function 2 -- usbc -- USB
bus 2, slot 0, function 3 -- usbd -- USB
bus 2, slot 1 -- dqa -- CMD 649 PCI-IDE
bus 2, slot 2 -- pka -- Adaptec AIC-7892
bus 3, slot 2 -- mca -- DEC PCI MC
probing hose 1, PCI
bus 0, slot 1 -- pga -- KGPSA-C
probing hose 2, PCI
probing PCI-to-PCI bridge, hose 2 bus 2
bus 2, slot 4 -- eia -- DE602-AA
bus 2, slot 5 -- eib -- DE602-AA
probing hose 3, PCI
probing hose 33, PCI
probing hose 34, PCI
probing hose 53, PCI
probing hose 54, PCI
starting drivers
Starting secondary CPU 1 at address 400030000
Starting secondary CPU 4 at address 2000030000
Starting secondary CPU 5 at address 2400030000
Starting secondary CPU 8 at address 4000030000
Starting secondary CPU 9 at address 4400030000
Starting secondary CPU 12 at address 6000030000
Starting secondary CPU 13 at address 6400030000
initializing GCT/FRU......................... at 54c000
Initializing dqa eia eib pka pga
AlphaServer Console V6.5-8, built on May  9 2003 at 10:10:56
P00>>>                  Console up and running.
```

Example 2 Hard Partition H1 sub partition P0

```
~NET-I-(td324) Session opened to port 324, fd 59, by host 192.168.2.1, port 2808.
~PCO-I-(pco_03) Powering on partition. HP: WEB89P2
~PCO-I-(pco_03)
```

(The following map shows the mesh of hard partition H1. Note that the map shows the IP connections between processors, both North–South and East-West, and indicates slots that have processors and slots that have fillers. Compare with the map for P0 in example 1)

```
Configuring for 8 CPUs for HP:1 WEB89P2

       0  1  2  3  4  5  6  7  8  9  A  B  C  D  E  F
      ............w..w..w..w........................
  0   ............P--F--P--F........................
      ............|..|..|..|........................   P  Processor
      ............|..|..|..|........................   F  Filler
  1   ............P--F--P--F........................   |  connection
      ............|..|..|..|........................   -  connection
      ............|..|..|..|........................
  2   ............P--F--P--F........................    W <----------> E
      ............|..|..|..|........................  N ^ (0,0)
      ............|..|..|..|........................    |  .
  3   ............P--F--P--F........................    |    .
      ............w..w..w..w........................    |       .
      ..............................................  S v      (ns,ew)
  4   ..............................................
      ..............................................
      ..............................................
  5   ..............................................
      ..............................................
      ..............................................
  6   ..............................................
      ..............................................
      ..............................................
  7   ..............................................
      ..............................................
       0  1  2  3  4  5  6  7  8  9  A  B  C  D  E  F
Running test 10, Initialize RAMBUS ... on 8 EV7s
Running test 11, Initialize Memory ... on 8 EV7s
Running test 12, Data Pattern March read/write ... on 8 EV7s
Running test 13, RAID channel Test ... on 8 EV7s
Running test 14, Single Bit Error ... on 8 EV7s
Running test 15, Double Bit Error ... on 8 EV7s
Running test 20, Init IO7 and Start Clocks ... on 2 EV7s
Running test 21, IO7 Data Path (Scratch CSR) ... on 2 EV7s
Running test 22, IO Single Bit Error checkers ... on 2 EV7s
Running test 23, IO Double Bit Error checkers ... on 2 EV7s
Running test 24, IO Timer Expirations ... on 2 EV7s
Running test 25, IO up-hose SBE checkers ... on 2 EV7s
Running test 26, IO up-hose DBE checkers ... on 2 EV7s
Running test 27, IO7 pass2 data mover test ... on 2 EV7s
Running test 30, Configure RBOX Routes ... on 8 EV7s
Running test 31, Clear Errors / Enable Routes ... on 8 EV7s
Running test 32, Route Test: N S E W ... on 8 EV7s (North)
Running test 32, Route Test: N S E W ... on 8 EV7s (South)
Running test 32, Route Test: N S E W ... on 8 EV7s (East )
Running test 32, Route Test: N S E W ... on 8 EV7s (West )
Running test 33, Inverse Route Setup ... on 8 EV7s
Running test 32, Route Test: N S E W ... on 8 EV7s (North)
Running test 32, Route Test: N S E W ... on 8 EV7s (South)
Running test 32, Route Test: N S E W ... on 8 EV7s (East )
Running test 32, Route Test: N S E W ... on 8 EV7s (West )
Running test 33, Inverse Route Setup ... on 8 EV7s
Running test 34, Single Bit Error checker ... on 8 EV7s
Running test 35, Double Bit Error checker ... on 8 EV7s
Running test 31, Clear Errors / Enable Routes ... on 8 EV7s
```

```
Running test 16, Interprocessor Memory Access ... on 8 EV7s
Running test 40, Local I/O Device Interrupts ... on 8 EV7s
Running test 41, Local Interval Timer Interrupts ... on 8 EV7s
Running test 42, Local Interprocess Interrupts ... on 8 EV7s
Running test 43, Software Alerts ... on 1 EV7s
Running test 46, Other Local Interrupt Bits ... on 8 EV7s
~PCO-I-(pco_03) HP:WEB89P2 SP:89P2 Primary is NS:0 EW:4 which is cab:00 drw:2 cpu:0
Running test 50, Loop on Secondary Routine ... on 3 EV7s (SP:89P2)
~PCO-I-(pco_03) HP:WEB89P2 SP:89P3 Primary is NS:2 EW:4 which is cab:00 drw:3 cpu:0
Running test 50, Loop on Secondary Routine ... on 3 EV7s (SP:89P3)
Running test 50, Loop on Secondary Routine ... on 0 EV7s (SP:Free_Pool)
starting console on CPU 0
initialized idle PCB
initializing semaphores
initializing heap
initial heap 700c0
memory low limit = 54c000 heap = 700c0, 1fffc0
initializing driver structures
initializing idle process PID
initializing file system
initializing timer data structures
lowering IPL
CPU 0 speed is 1050 MHz
create dead_eater
create poll
create timer
create powerup
entering idle loop
access NVRAM
Get Partition DB                    Partition data base already established (see Getting
Started with Partitions)
hpcount = 2, spcount = 5, ev7_count = 32, io7_count = 7
hard_partition = 1      this is hard partition 1
IO7-100 (Pass 2) at PID 0
IO7 North port speed is 210 MHz
Hose 0 - 33 MHz PCI
Hose 1 - 33 MHz PCI
Hose 2 - 33 MHz PCI
Hose 3 - 2X AGP
0 sub-partition 0:   start:00000000 00000000   size:00000000 80000000
PID 0 console memory base: 0, 2 GB
1 sub-partition 0:   start:00000004 00000000   size:00000000 80000000
PID 1 memory: 400000000, 2 GB
2 sub-partition 0:   start:00000020 00000000   size:00000000 40000000
PID 4 memory: 2000000000, 1 GB
3 sub-partition 0:   start:00000024 00000000   size:00000000 40000000
PID 5 memory: 2400000000, 1 GB
0 sub-partition 1:   start:00000040 00000000   size:00000001 00000000
1 sub-partition 1:   start:00000044 00000000   size:00000001 00000000
2 sub-partition 1:   start:00000060 00000000   size:00000000 40000000
3 sub-partition 1:   start:00000064 00000000   size:00000000 40000000
total memory, 6 GB    of sub-partition 0 in HP 1 (note CPUs assigned to this sub-
partition – compare example 1)
probe I/O subsystem
probing hose 0, PCI
probing PCI-to-PCI bridge, hose 0 bus 2
do not use secondary IDE channel on CMD controller
bus 2, slot 0, function 0 -- usba -- USB
bus 2, slot 0, function 1 -- usbb -- USB
bus 2, slot 0, function 2 -- usbc -- USB
bus 2, slot 0, function 3 -- usbd -- USB
bus 2, slot 1 -- dqa -- CMD 649 PCI-IDE
bus 2, slot 2 -- pka -- Adaptec AIC-7892
probing hose 1, PCI
bus 0, slot 1 -- pga -- KGPSA-C
probing hose 2, PCI
probing PCI-to-PCI bridge, hose 2 bus 2
bus 0, slot 1, function 0 -- pkb -- Adaptec AIC-7899
bus 0, slot 1, function 1 -- pkc -- Adaptec AIC-7899
bus 2, slot 4 -- eia -- DE602-AA
bus 2, slot 5 -- eib -- DE602-AA
```

```
probing hose 3, PCI
bus 0, slot 5 -- vga -- 3D Labs OXYGEN VX1 AGP
starting drivers
Starting secondary CPU 1 at address 400030000
Starting secondary CPU 4 at address 2000030000
Starting secondary CPU 5 at address 2400030000
initializing GCT/FRUinitializing keyboard
......................... at 54c000
Initializing dqa eia eib pka pkb pkc pga
AlphaServer Console V6.5-8, built on May  9 2003 at 10:10:56

P00>>>                                  Console up and running.
```

Example 3 Hard Partition H1 sub partition P1

```
~NET-I-(td325) Session opened to port 325, fd 56, by host 192.168.2.1, port 2809.
ÿû~PCO-I-(pco_03) Powering on partition. HP: WEB89P2
~PCO-I-(pco_03)
```

(The following map shows the mesh of hard partition H1. Note that the map shows the IP connections between processors, both North–South and East-West, and indicates slots that have processors and slots that have fillers. Compare with the map for P0 in Example 1 and P1 in Example 2.)

```
Configuring for 8 CPUs for HP:1 WEB89P2

       0  1  2  3  4  5  6  7  8  9  A  B  C  D  E  F
      ............w..w..w..w.......................
  0   ............P--F--P--F.......................
      ............|..|..|..|.......................   P  Processor
      ............|..|..|..|.......................   F  Filler
  1   ............P--F--P--F.......................   |  connection
      ............|..|..|..|.......................   -  connection
      ............|..|..|..|.......................
  2   ............P--F--P--F.......................    W <----------> E
      ............|..|..|..|.......................  N ^ (0,0)
      ............|..|..|..|.......................    |  .
  3   ............P--F--P--F.......................    |    .
      ............w..w..w..w.......................    |       .
      .............................................  S v      (ns,ew)
  4   .............................................
      .............................................
      .............................................
  5   .............................................
      .............................................
      .............................................
  6   .............................................
      .............................................
      .............................................
  7   .............................................
      .............................................
       0  1  2  3  4  5  6  7  8  9  A  B  C  D  E  F
Running test 10, Initialize RAMBUS ... on 8 EV7s
Running test 11, Initialize Memory ... on 8 EV7s
Running test 12, Data Pattern March read/write ... on 8 EV7s
Running test 13, RAID channel Test ... on 8 EV7s
Running test 14, Single Bit Error ... on 8 EV7s
Running test 15, Double Bit Error ... on 8 EV7s
Running test 20, Init IO7 and Start Clocks ... on 2 EV7s
Running test 21, IO7 Data Path (Scratch CSR) ... on 2 EV7s
Running test 22, IO Single Bit Error checkers ... on 2 EV7s
Running test 23, IO Double Bit Error checkers ... on 2 EV7s
Running test 24, IO Timer Expirations ... on 2 EV7s
2Running test 25, IO up-hose SBE checkers ... on 2 EV7s
Running test 26, IO up-hose DBE checkers ... on 2 EV7s
Running test 27, IO7 pass2 data mover test ... on 2 EV7s
Running test 30, Configure RBOX Routes ... on 8 EV7s
Running test 31, Clear Errors / Enable Routes ... on 8 EV7s
Running test 32, Route Test: N S E W ... on 8 EV7s (North)
Running test 32, Route Test: N S E W ... on 8 EV7s (South)
Running test 32, Route Test: N S E W ... on 8 EV7s (East )
Running test 32, Route Test: N S E W ... on 8 EV7s (West )
Running test 33, Inverse Route Setup ... on 8 EV7s
Running test 32, Route Test: N S E W ... on 8 EV7s (North)
Running test 32, Route Test: N S E W ... on 8 EV7s (South)
Running test 32, Route Test: N S E W ... on 8 EV7s (East )
Running test 32, Route Test: N S E W ... on 8 EV7s (West )
Running test 33, Inverse Route Setup ... on 8 EV7s
Running test 34, Single Bit Error checker ... on 8 EV7s
Running test 35, Double Bit Error checker ... on 8 EV7s
```

03/06/10_11:39:43Running test 31, Clear Errors / Enable Routes ... on 8 EV7s
03/06/10_11:39:51Running test 16, Interprocessor Memory Access ... on 8 EV7s
03/06/10_11:39:51Running test 40, Local I/O Device Interrupts ... on 8 EV7s
03/06/10_11:39:51Running test 41, Local Interval Timer Interrupts ... on 8 EV7s
03/06/10_11:39:51Running test 42, Local Interprocess Interrupts ... on 8 EV7s
03/06/10_11:39:52Running test 43, Software Alerts ... on 1 EV7s
03/06/10_11:39:52Running test 46, Other Local Interrupt Bits ... on 8 EV7s
03/06/10_11:39:53~PCO-I-(pco_03) HP:WEB89P2 SP:89P2 Primary is NS:0 EW:4 which is cab:00 drw:2 cpu:0
03/06/10_11:39:53Running test 50, Loop on Secondary Routine ... on 3 EV7s (SP:89P2)
03/06/10_11:39:57~PCO-I-(pco_03) HP:WEB89P2 SP:89P3 Primary is NS:2 EW:4 which is cab:00 drw:3 cpu:0
03/06/10_11:39:57Running test 50, Loop on Secondary Routine ... on 3 EV7s (SP:89P3)
03/06/10_11:39:57Running test 50, Loop on Secondary Routine ... on 0 EV7s (SP:Free_Pool)
starting console on CPU 8
console physical memory base is 4000000000
initialized idle PCB
initializing semaphores
initializing heap
initial heap 700c0
memory low limit = 54c000 heap = 700c0, 1fffc0
initializing driver structures
initializing idle process PID
initializing file system
initializing timer data structures
lowering IPL
CPU 8 speed is 1050 MHz
create dead_eater
create poll
create timer
create powerup
entering idle loop
access NVRAM
Get Partition DB   Partition data base already established (see Getting Started with Partitions)
hpcount = 2, spcount = 5, ev7_count = 32, io7_count = 7
hard_partition = 1
IO7-100 (Pass 2) at PID 8
IO7 North port speed is 210 MHz
Hose 32 - 33 MHz PCI
Hose 33 - 66 MHz PCI
Hose 34 - 66 MHz PCI
Hose 35 - 2X AGP
0 sub-partition 0:   start:00000000 00000000   size:00000000 80000000
1 sub-partition 0:   start:00000004 00000000   size:00000000 80000000
2 sub-partition 0:   start:00000020 00000000   size:00000000 40000000
3 sub-partition 0:   start:00000024 00000000   size:00000000 40000000
0 sub-partition 1:   start:00000040 00000000   size:00000001 00000000
PID 8 console memory base: 4000000000, **4 GB**
1 sub-partition 1:   start:00000044 00000000   size:00000001 00000000
PID 9 memory: 4400000000, **4 GB**
2 sub-partition 1:   start:00000060 00000000   size:00000000 40000000
PID 12 memory: 6000000000, **1 GB**
3 sub-partition 1:   start:00000064 00000000   size:00000000 40000000
PID 13 memory: 6400000000, **1 GB**
total memory, 10 GB of sub-partition 1 in HP 1 (note CPUs assigned to this sub-partition - compare example 2)
waiting for GCT/FRU at 54c000 by CPU 0
waiting for GCT/FRU at 54c000 by CPU 0
waiting for GCT/FRU at 54c000 by CPU 0
waiting for GCT/FRU at 54c000 by CPU 0
probe I/O subsystem
probing hose 32, PCI
probing PCI-to-PCI bridge, hose 32 bus 2
do not use secondary IDE channel on CMD controller
probing PCI-to-PCI bridge, hose 32 bus 3
bus 2, slot 0, function 0 -- usba -- USB
bus 2, slot 0, function 1 -- usbb -- USB
bus 2, slot 0, function 2 -- usbc -- USB
bus 2, slot 0, function 3 -- usbd -- USB
bus 2, slot 1 -- dqa -- CMD 649 PCI-IDE

```
bus 2, slot 2 -- pka -- Adaptec AIC-7892
bus 3, slot 0 -- pkb -- QLogic ISP10x0
bus 3, slot 2 -- mca -- DEC PCI MC
probing hose 33, PCI
bus 0, slot 1, function 0 -- pkc -- Adaptec AIC-7899
bus 0, slot 1, function 1 -- pkd -- Adaptec AIC-7899
probing hose 34, PCI
probing PCI-to-PCI bridge, hose 34 bus 2
bus 2, slot 4 -- eia -- DE602-B*
bus 2, slot 5 -- eib -- DE602-B*
probing hose 35, PCI
starting drivers
Starting secondary CPU 9 at address 4400030000
Starting secondary CPU 12 at address 6000030000
Starting secondary CPU 13 at address 6400030000
initializing GCT/FRU to 54c000
Initializing pkb dqa eia eib pka pkc pkd
AlphaServer Console V6.5-8, built on May  9 2003 at 10:10:56
P08>>>
```

Power-up Flow

Assuming the system is plugged in and its circuit breakers are on, power-up occurs in a given hard partition when the following conditions are met:

- If part or all of the hard partition is in a cabinet, the OCP switch on the cabinet is in the on or secure position.
- The power on command is issued on the internal LAN to an identified hard partition.

Hardware power-up participants on the internal LAN and part of the partition are:

- the Marvel Backplane Manager (MBM) drives the actions of the CMMs
- the CPU Module Manager (CMM), controls the power-up, reset and communication with the CPUs
- PCI Backplane manager (PBM), controls the power and monitors the environment of the PCI drawer
- each CPU performs self-test actions and reports the results to the CMM
- partition primary CPU, a CPU selected to perform actions in a partition which only need to be performed once
- partition primary CMM, the CMM that controls a partition primary CPU
- partition coordinator MBM is the MBM assigned to control a given partition
- group leader, the lowest numbered MBM or PBM in the system

The actions described below in the MBM column are performed by each MBM unless restricted by the prefix partition primary MBM only.
The actions described below in the CMM column are performed by each CMM unless restricted by the prefix partition primary CMM only.
Each CPU performs the actions described below in the CPU column, unless restricted by the prefix partition primary CPU only.

Power-up Flow Table

Step	MBM	CMM	CPU	PBM
1	Execute module POST. Display MBM success/fail on OCP.	Execute module POST. Configure FPGA. Read RIMM speed, size, config. info. (I^2C)		Execute module POST.
2	Discover Server Management LAN Form a group Select the leader			Discover Server Management LAN. Form a group Select the leader
3	Read partition information from non-volatile RAM on itself.			Read partition information from non-volatile RAM on itself
4	If discrepancies found, the group leader reconciles			
5	Poll CMMs			
6		Return POST status, CPU and memory configuration information to the MBM.		
7	Display CMM pass/fail on OCP.			
8	If the power switch is on, the group leader commands all MBMS to power up the backplanes and the CMMs and PBMs to power-up.			
9		Power up the CPUs and put the CPU in RESET, verify DC OK. Load SROM data		Power up the I/O drawer, load hot-plug FPGA code

		into the FPGA. Load shared RAM structures with RIMM SPDs and other config data received in step 1.		
10	Determine IP and IO cable presence and connections by initiating cable tests on MBM & PBM			Determine IO cable presence and connection
11	Start cable ID receivers on S, W. Start cable ID senders on N, S, and IO.			Start cable ID receivers
12	Receive return status			Return receive status
13	Partition coordinator with partition and cable connectivity information compute partition routing configuration and establish PIDs for each CPU in partition			
14		Receive EV7 routing and PIDs. assignments for CPUs CMM resp. for. Store in shared RAM		
15	Take CPUs out of reset to begin EV7 init.			
16		Establish comm link to EV7. Pass MBM command to CPUs.		
17			Execute EV7 BIST. If good, load SROM via the SROM port from FPGA, configure IPRs (except for those for the RAMbus and router). Load and configure the EV7 PID in the CBOX_WHAMI IPR. Configure cache. Init communication to CMM and return self-test status.	
18		Determine good CPUs and assert reset on bad CPUs. Begin XSROM load.		
19			Load XSROM tests via GIO PORT. Respond to CMM	
20		Report XSROM load status to MBM		
21	Display progress on OCP. Command CMMs to begin memory test			
22		Pass MBM command to CPUs		
23			Configure and **test** memory. Return status.	
24		Report memory test status to MBM		
25	If error & FRU EV7, remove CPU from partition and return to step 13. Else command CMMs to run			

	the XSROM tests to configure and test I/O. If memory error, remove resource and keep going. Display progress on OCP		
26		Pass MBM command to CPUs.	
27			Configure the I/O port. Configure IO7 if present. Return status.
28		Return status.	
29	If error & FRU EV7, remove CPU from partition and return to step 13. Else partition coordinator selects a primary CPU for each partition. Partition coordinator computes routing and initiate partition-wide router configuration. If IO error, remove resource and keep going. Display progress on OCP.		
30		Passes MBM initiate router configuration command	
31			Initialize IP ports. Load the partition router config. Return status.
32		Return status.	
33	If error & FRU EV7, remove CPU from partition and return to step 13. Else partition coordinator commands run the router validation XSROM tests.		
34		Passes on MBM router validation XSROM test command	
35			Perform the traversal algorithm for the N, S, E, W router paths. Return status.
36		Return status to partition coordinator.	
37	If error & FRU EV7, remove CPU from partition and return to step 13. Else partition coordinator initiates remote memory/access XSROM tests on each CPU module.		
38		Passes MBM command to CPUs.	
39			Perform memory tests across the EV7 IP network. Return status
40		Return status to partition coordinator	
41	If error & FRU EV7,		

	remove CPU from partition and return to step 13. Else partition coordinator initiate interrupt / error testing on each CPU module.		
42		Passes MBM command to CPUs.	
43			Perform error testing and CPU interrupt handling across the network. Return status
44		Return status to partition coordinator	
45	If error & FRU EV7, remove CPU from partition and return to step 13. Else partition coordinator elects a primary EV7 in each partition. Initiate XSROM test to loop on RBOX_SCRATCH on all secondaries while console loaded by primary.		
46		Execute test on all secondary CPUs.	
47			Execute XSROM test to loop on RBOX_SCRATCH waiting for jump address.
48	Partition coordinator initiates command to load and execute the Console/PAL firmware on the primary EV7.		
49		Partition primary CMM only: Initiate loading of the console/PAL firmware.	
50			Load console/PAL via the CMM. Return Status
51		Initiate transfer of control to console/PAL. Return status.	
52			Transfer control to console. Console performs further initialization for error handling, device interrupt handling, and steps for I/O port configuration.

Booting

Setting Boot Options

You can set a default boot device, boot flags, and network boot protocols for Tru64 UNIX or OpenVMS using the SRM set command with environment variables. Once these environment variables are set, the boot command defaults to the stored values. You can override the stored values for the current boot session by entering parameters on the boot command line.

The SRM boot-related environment variables are:

bootdef_dev

boot_file

boot_osflags

Setting Boot Options: bootdef_dev

Specifies one or more devices from which to boot the operating system. When more than one device is specified, the system searches in the order listed and boots from the first device. Enter the show bootdef_dev command to display the current default boot device. Enter the show device command for a list of all devices in the system.

Syntax

set bootdef_dev *boot_device*

boot_device	The name of the device on which the system software has been loaded. To specify more than one device, separate the names with commas.

Example

In this example, two boot devices are specified. The system will try booting from dkb0 and, if unsuccessful, will boot from dka0.

P00>>> set bootdef_dev dkb0, dka0

NOTE: When you set the bootdef_dev environment variable, it is recommended that you set the operating system boot parameters as well, using the set boot_osflags command.

Setting Boot Options: boot_file

Specifies the default file name to be used for booting when no file name is specified by the boot command.

Syntax

set boot_file filename

Example

In this example, a boot file is specified for booting OpenVMS from the InfoServer. APB_0712 is the file name of the APB program used for the initial system load (ISL) boot program.

P00>>> set boot_file apb_0712

Setting Boot Options: boot_osflags

Sets the default boot flags and, for OpenVMS, a root number.

Boot flags contain information used by the operating system to determine some aspects of a system bootstrap. Under normal circumstances, you can use the default boot flag settings.

To change the boot flags for the current boot only, use the flags_value argument with the boot command.

Examples

P00>>> set boot__osflags a

P00>>> set boot_osflags 2,1

P00>>> set boot_osflags 0,80

Syntax

set boot_osflags flags_value

The flags_value argument is specific to the operating system.

Boot Flag Settings

1	0	Bootstrap conversationally (enables you to modify SYSGEN parameters in SYSBOOT).
2	1	Map XDELTA to a running system
4	2	Stop at initial system breakpoint.
8	3	Perform diagnostic breakpoints.
10	4	Stop at the bootstrap breakpoints.
20	5	Omit header from secondary bootstrap file.
80	7	Prompt for the name of the secondary bootstrap file.
100	8	Halt before secondary bootstrap.
10000	16	Display debug messages during booting.
20000	17	Display user messages during booting.

Booting Tru64 UNIX

For complete instructions on booting *Tru64 UNIX*, see the *Tru64 UNIX Installation Guide*.

Example: Booting over the Internet / RIS boot

```
P00>>> show device
```

dka0.0.0.1.1 DKA0 RZ2DD-LS 0306
dka100.1.0.1.1 DKA100 RZ2DD-LS 0306
dka200.2.0.1.1 DKA200 RZ1CB-CS 0844
dkb0.0.0.3.1 DKB0 RZ25 0900
dqa0.0.0.15.0 DQA0 TOSHIBA CD-ROM XM-6302B 1012 dva0.0.0.1000.0 DVA0
eia0.0.0.4.1 EIA0 00-00-F8-09-90-FF
eib0.0.0.2002.1 EIB0 00-06-2B-00-25-5B
pka0.7.0.1.1 PKA0 SCSI Bus ID 7
pkb0.7.0.3.1 PKB0 SCSI Bus ID 7

```
P00>>> set eia0_protocols bootp

P00>>> set eia0_inet_init bootp

P00>>> boot eia0
```

Booting from a Local SCSI Disk

Perform the following tasks to boot a system from the local SCSI disk:

1. Power up the system. The system stops at the SRM console prompt, P00>>>.
2. Set boot environment variables, if desired.
3. Install the boot medium.
4. Enter the **show device** command to determine the unit number of the drive for your device.
5. Enter the **boot** command and command-line parameters (if you have not set the associated environment variables). In the example, the boot device and boot flags have already been set.

Example: Booting Tru64 UNIX from a local SCSI disk

```
P00>>>show dev
 dka0.0.0.2002.0            DKA0               COMPAQ BF03664664   3B08
 dka100.1.0.2002.0          DKA100             COMPAQ BF03665223   B008
 dkb0.0.0.1.1               DKB0               COMPAQ BF03665223   B008
 dkb100.1.0.1.1             DKB100             COMPAQ BF03665223   B008
 dkd0.0.0.2002.4            DKD0               COMPAQ BF03664664   3B08
 dkd100.1.0.2002.4          DKD100             COMPAQ BF03664664   3B08
 dke0.0.0.1.5               DKE0               COMPAQ BF03665223   B008
 dke100.1.0.1.5             DKE100             COMPAQ BF03665223   B008
 dqa0.0.0.2001.0            DQA0                         CD-W216E   E.0A
 dqb0.0.0.2001.4            DQB0                         CD-W216E   E.0A
 eia0.0.0.2004.2            EIA0               00-08-02-00-D5-7C
 eib0.0.0.2005.2            EIB0               00-08-02-00-D5-7D
 eic0.0.0.2004.6            EIC0               00-08-02-3E-87-F8
```

```
eid0.0.0.2005.6            EID0              00-08-02-3E-87-F9
pka0.7.0.2002.0            PKA0                  SCSI Bus ID 7
pkb0.7.0.1.1               PKB0                  SCSI Bus ID 7
pkc0.7.0.101.1             PKC0                  SCSI Bus ID 7
.
.
.
P00>>>b dka0 -fl a
(boot dka0.0.0.2002.0 -flags A)
block 0 of dka0.0.0.2002.0 is a valid boot block
reading 19 blocks from dka0.0.0.2002.0
bootstrap code read in
base = b80000, image_start = 0, image_bytes = 2600(9728)
initializing HWRPB at 10000
GCT base = 55a000
initializing page table at b6c000
initializing machine state
setting affinity to the primary CPU
jumping to bootstrap code

UNIX boot - Monday September 23, 2002

Loading vmunix ...
Loading text at 0xffffffff00000000
Loading data at 0xffffffff00800000

Sizes:
text =  7053248
data =  1395264
bss  =  1845376
Starting at 0xffffffff00012c00

Loading vmunix symbol table ... [1700592 bytes]
Alpha boot: available memory from 0xbe3a000 to 0x2d00000000
Compaq Tru64 UNIX V5.1B (Rev. 2649); Fri Oct  4 13:03:43 EDT 2002
physical memory = 32768.00 megabytes.
available memory = 31847.62 megabytes.
using 125719 buffers containing 982.17 megabytes of memory
Master cpu at slot 0
.
.
.
Configuring network
hostname: mrqd06.mro.cpqcorp.net
Loading LMF licenses
Combine PRESTOSERVE-OA BIR-PK-94343-2-MLO-ALESI-8559 with PRESTOSERVE-OA BIR-
PK-94299-2-MLO-TETCHY-21024
Multiple Licenses could not be combined for OSF-USR DEC
Combine OSF-BASE BIR-PK-94299-2-MLO-TETCHY-21021 with OSF-BASE BIR-OG-96004-3-
MLO-ALESI-13102
System error logger started
Binary error logger started
binlogd: failed to initialize remote logging. Please make sure the
network is set up properly and then restart binlogd.
add net default: gateway 16.129.104.1
Setting kernel timezone variable
ONC portmap service started
NFS IO service started
Mounting NFS filesystems
Preserving editor files
Clearing temporary files
16-Oct-2002 13:26:17 [700] Correctable error reporting state changed
Unlocking ptys
```

```
 Secure Shell daemon (sshd2) started.
 SMTP Mail Service started.
 Environmental Monitoring Subsystem Configured.
 Using snmp service entry port 161.
 Extensible SNMP master agent started
 Base O/S subagent started
 Server System subagent started
 Server Management subagent started
 CIM SNMP subagent started
 Performance Management subagent started
 Web Based Management Agent started
 ConfigReport Management Module started
 Sysman Management Module started
 Threshold Management Subagent started
 Intelligent Drive Array Subagent started
 The SNMP trap to Event Manager interface is disabled.
 GS Platform View and Discovery V1.3 for Insight Manager is only supported on
Alpha GS series platforms.
 AvFS daemon (advfsd) started.
 Internet services provided.
 Cron service started
 SuperLAT. Copyright 1994 Meridian Technology Corp. All rights reserved.
 LAT started.
 LSM volwatch Service started - mail only
 Printer service started
 Starting DESTA Director process.
 Logging outputs to: /usr/opt/compaq/svctools/logs/desta_dir.log
 The Director process has started successfully.

 SysMan authentication server (smauthd) started
 SysMan Station server (smsd) started


Compaq Tru64 UNIX V5.1B (Rev. 2649) (mrqd06.mro.cpqcorp.net) console

login:
```

Booting OpenVMS

OpenVMS is booted from a local SCSI disk drive or from a CD-ROM drive on the InfoServer. For complete instructions on booting OpenVMS, see the OpenVMS Installation document.

Booting OpenVMS

OpenVMS systems require an ordered pair as the flags_value argument: root_number and boot_flags.

root_number Directory number of the system disk on which OpenVMS files are located. For example:

Root_number	**Root Directory**
0	[SYS0.SYSEXE]
1	[SYS1.SYSEXE]
2	[SYS2.SYSEXE]
3	[SYS3.SYSEXE]

boot_flags The hexadecimal value of the bit number or numbers set. To specify multiple boot flags, add the flag values (logical OR). For example, the flag value 10080 executes both the 80 and 10000 flag settings. See the boot flag settings table.

Booting OpenVMS from an InfoServer

1. Power up the system. The system stops at the P00>>> console prompt.
2. Insert the operating system CD-ROM into the CD-ROM drive connected to the InfoServer.
3. Enter the show device command to determine the unit number of the drive for your device.
4. Enter the boot command and any command-line parameters. In the example the device is EIA0. APB_0721 is the file name of the APB program used for the initial system load (ISL) boot program. The InfoServer ISL program displays a menu.
5. Respond to the menu prompts, using the selections shown in the example.

EXAMPLE of Booting OpenVMS from an InfoServer

```
P00>>> show device
dka0.0.0.1.1 DKA0 RZ2CA-LA N1H0
dka100.1.0.1.1 DKA100 RZ2CA-LA N1H0
dqa0.0.0.15.0 DQA0 TOSHIBA CD-ROM XM-6302B 1012
dva0.0.0.1000.0 DVA0
eia0.0.0.6.1 EIA0 00-00-F8-10-D6-03
pka0.7.0.1.1 PKA0 SCSI Bus ID 7
P00>>>
.
.
.
P00>>> boot -flags 0,0 -file apb_0721 eia0
(boot eia0.0.0.6.1 -file APB_0712 -flags 0,0)
Trying MOP boot.
.............
Network load complete.
Host name: CALSUN
Host address: aa-00-04-00-a4-4e
bootstrap code read inbase = 200000, image_start = 0, image_bytes =
70400
initializing HWRPB at 2000
initializing page table at 3ffee000
initializing machine state
setting affinity to the primary CPU
jumping to bootstrap code
Network Initial System Load Function
Version 1.2

FUNCTION ID FUNCTION
ID
1 - Display Menu2
2 - Help
3 - Choose Service
4 - Select Options
5 - Stop
Enter a function ID Value:
Enter a function ID Value: 3
OPTION OPTION
ID
1 - Find Services
2 - Enter known Service Name

Enter an Option ID value: 2Enter a Known Service Name:
ALPHA_V721_SSB

OpenVMS (TM) Alpha Operating System, Version V7.2-1
```

Booting OpenVMS from a Local SCSI Drive

```
P00>>> show device
 dka0.0.0.2002.0            DKA0              COMPAQ BF03664664  3B08
dka100.1.0.2002.0          DKA100            COMPAQ BF03665223  B008
dkb0.0.0.1.1               DKB0              COMPAQ BF03665223  B008
dkb100.1.0.1.1             DKB100            COMPAQ BF03665223  B008
dkd0.0.0.2002.4            DKD0              COMPAQ BF03664664  3B08
dkd100.1.0.2002.4          DKD100            COMPAQ BF03664664  3B08
dke0.0.0.1.5               DKE0              COMPAQ BF03665223  B008
dke100.1.0.1.5             DKE100            COMPAQ BF03665223  B008
dqa0.0.0.2001.0            DQA0                       CD-W216E  E.0A
dqb0.0.0.2001.4            DQB0                       CD-W216E  E.0A
eia0.0.0.2004.2            EIA0              00-08-02-00-D5-7C
eib0.0.0.2005.2            EIB0              00-08-02-00-D5-7D
eic0.0.0.2004.6            EIC0              00-08-02-3E-87-F8
eid0.0.0.2005.6            EID0              00-08-02-3E-87-F9
pka0.7.0.2002.0            PKA0                  SCSI Bus ID 7
pkb0.7.0.1.1               PKB0                  SCSI Bus ID 7
pkc0.7.0.101.1             PKC0                  SCSI Bus ID 7
.
.
.
 P00>>>b dka0 -fl 0,0
(boot dka0.0.0.2002.0 -flags 0,0)
block 0 of dka0.0.0.2002.0 is a valid boot block
reading 969 blocks from dka0.0.0.2002.0
bootstrap code read in
base = b78000, image_start = 0, image_bytes = 79200(496128)
initializing HWRPB at 10000
GCT base = 55a000
initializing page table at b64000
initializing machine state
setting affinity to the primary CPU
jumping to bootstrap code

    OpenVMS (TM) Alpha Operating System, Version V7.3-1

%SMP-I-SECMSG, CPU #01 message:   START
%SMP-I-CPUTRN, CPU #01 has joined the active set.
%SMP-I-SECMSG, CPU #02 message:   START
%SMP-I-CPUTRN, CPU #02 has joined the active set.
%SMP-I-SECMSG, CPU #05 message:   START
%SMP-I-SECMSG, CPU #03 message:   START
%SMP-I-CPUTRN, CPU #03 has joined the active set.
%SMP-I-SECMSG, CPU #07 message:   START
%SMP-I-SECMSG, CPU #06 message:   START
%SMP-I-SECMSG, CPU #04 message:   START
%SMP-I-CPUTRN, CPU #05 has joined the active set.
%SMP-I-CPUTRN, CP~DBS-W-(trainsrv) Using the free pool to save delta time for
HP: 0, SP: 0
U #04 has joined the active set.
%SMP-I-CPUTRN, CPU #06 has joined the active set.
%SMP-I-CPUTRN, CPU #07 has joined the active set.
%SYSINIT-I- waiting to form or join an OpenVMS Cluster
%VMScluster-I-AUTH_DEFAULT, no authorization file, defaulting to group 0
%EIA0, Auto-negotiation mode set by console
%EIA0, Auto-negotiation started, advertising 100BaseTX Full Duplex
%EIB0, Auto-negotiation mode set by console
%EIB0, Auto-negotiation started, advertising 100BaseTX Full Duplex
%EIA0, Full Duplex 100BaseTX connection selected
%EIC0, Auto-negotiation mode set by console
%EIC0, Auto-negotiation started, advertising 100BaseTX Full Duplex
```

```
%EID0, Auto-negotiation mode set by console
%EID0, Auto-negotiation started, advertising 100BaseTX Full Duplex
%CNXMAN,  Proposing formation of a VMScluster
%CNXMAN,  Now a VMScluster member -- system MRQD11
%CNXMAN,  Completing VMScluster state transition
$!  Copyright ) 2002 Compaq Information Technologies Group, L.P.
%STDRV-I-STARTUP, OpenVMS startup begun at 15-OCT-2002 15:44:05.29
%%%%%%%%%%%  OPCOM  15-OCT-2002 15:44:28.54  %%%%%%%%%%%

The operator console and logfile will not be enabled.
Change OPC$OPA0_ENABLE & OPC$LOGFILE_ENABLE in SYLOGICALS.COM
to enable them.
%RUN-S-PROC_ID, identification of created process is 20200407
%RUN-S-PROC_ID, identification of created process is 20200408
%SET-I-NEWAUDSRV, identification of new audit server process is 2020040E
%RUN-S-PROC_ID, identification of created process is 20200412
%STARTUP-I-AUDITCONTINUE, audit server initialization complete

The OpenVMS Alpha system is now executing the site-specific startup commands.
%RUN-S-PROC_ID, identification of created process is 20200418
%SET-I-INTSET, login interactive limit = 64, current interactive value = 0
Waiting for mouse...    14 seconds remaining
Waiting for mouse...    13 seconds remaining
Waiting for mouse...    12 seconds remaining
Waiting for mouse...    11 seconds remaining
Waiting for mouse...    10 seconds remaining
Waiting for mouse...    9 seconds remaining
Waiting for mouse...    8 seconds remaining
Waiting for mouse...    7 seconds remaining
Waiting for mouse...    6 seconds remaining
Waiting for mouse...    5 seconds remaining
Waiting for mouse...    4 seconds remaining
Waiting for mouse...    3 seconds remaining
Waiting for mouse...    2 seconds remaining
Waiting for mouse...    1 seconds remaining
%DECW$DEVICE-I-NOINPUTDEVICES, no default input devices found.
Starting the DESTA Director Process...
%RUN-S-PROC_ID, identification of created process is 2020041C
%RUN-S-PROC_ID, identification of created process is 2020041D
%NCP-I-NOINFO, No information in database
%RUN-S-PROC_ID, identification of created process is 2020041F

%TCPIP-I-INFO, TCP/IP Services startup beginning at 15-OCT-2002 15:44:58.29
%TCPIP-I-NORMAL, timezone information verified
%RUN-S-PROC_ID, identification of created process is 20200420
%TCPIP-I-SETLOCAL, setting domain and/or local host
%TCPIP-I-STARTCOMM, starting communication
%TCPIP-I-SETPROTP, setting protocol parameters
%TCPIP-I-DEFINTE, defining interfaces
%TCPIP-I-STARTNAME, starting name service
%TCPIP-S-STARTDONE, TCP/IP Kernel startup completed
%TCPIP-I-PROXYLOADED, loaded 0 NFS proxy records
%TCPIP-I-LOADSERV, loading TCPIP server proxy information
%TCPIP-I-SERVLOADED, auxiliary server loaded with 1 proxy records
-TCPIP-I-SERVSKIP, skipped 0 communication proxy records
-TCPIP-I-SERVTOTAL, total of 1 proxy records read
%TCPIP-S-STARTDONE, TCPIP$PROXY startup completed
%TCPIP-S-STARTDONE, TCPIP$FTP startup completed
%TCPIP-S-STARTDONE, TCPIP$FTP_CLIENT startup completed
%TCPIP-S-STARTDONE, TCPIP$REXEC startup completed
%TCPIP-S-STARTDONE, TCPIP$RLOGIN startup completed
%TCPIP-S-STARTDONE, TCPIP$RSH startup completed
%TCPIP-S-STARTDONE, TCPIP$SMTP startup completed
```

```
%TCPIP-S-STARTDONE, TCPIP$TELNET startup completed
%TCPIP-S-STARTDONE, TCP/IP Services startup completed at 15-OCT-2002
15:45:01.52
%MOUNT-I-MOUNTED, SITP mounted on _MRQD11$DKA100:
%MOUNT-I-MOUNTED, CTM mounted on _MRQD11$DKB0:
%MOUNT-I-MOUNTED, PAGE_DUMP mounted on _MRQD11$DKB100:
%MOUNT-I-REBUILD, volume was improperly dismounted; rebuild in progress
%MOUNT-I-MOUNTED, SCRATCH1 mounted on _MRQD11$DKD0:
%MOUNT-I-MOUNTED, SCRATCH2 mounted on _MRQD11$DKD100:
%MOUNT-I-MOUNTED, SCRATCH3 mounted on _MRQD11$DKE0:
%MOUNT-I-MOUNTED, SCRATCH4 mounted on _MRQD11$DKE100:
%DCL-I-SUPERSEDE, previous value of CTM$LOGS has been superseded
Job SITP$SYSTEM_CRASH_ANAL (queue SYS$BATCH, entry 1) started on SYS$BATCH
  SYSTEM       job terminated at 15-OCT-2002 15:45:04.54
  Accounting information:
  Buffered I/O count:               5344      Peak working set size:       7120
  Direct I/O count:                 4059      Peak virtual size:         187216
  Page faults:                      9237      Mounted volumes:                7
  Charged CPU time:        0 00:00:01.70      Elapsed time:       0 00:00:59.45
 Welcome to OpenVMS (TM) Alpha Operating System, Version V7.3-1
Username:
```

Booting an ES47 Tower using the OCP

The ES47 is the only system in this family which comes in a tower configuration.
When you walk up to the system, the default display shows the last 4 alert messages

Hitting the RIGHT button clears the last 4 messages and brings you to the main menu:

```
T1.0-11875
>Show Box
Power On
Power Off
```

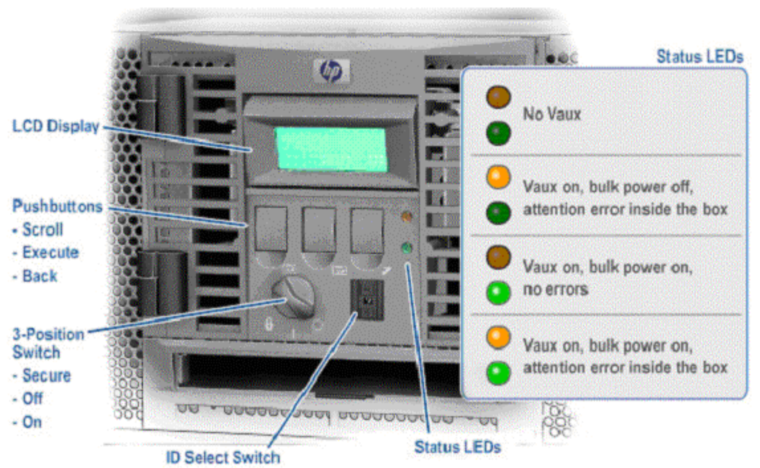

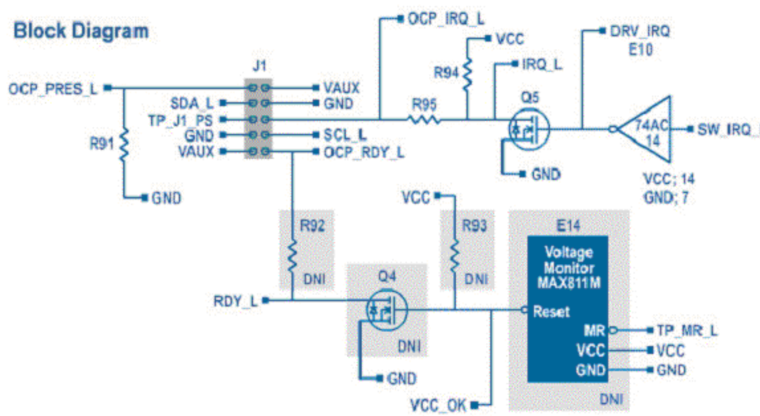

Hit the LEFT button to scroll the “>” prompt to select Power On:

```
T1.0-11875
Show Box
>Power On
Power Off
```

Hit the MIDDLE button to execute the Power On selection which brings you to another display:

```
PARTITIONS
>All Partitions
Default_HP
```

Hit the MIDDLE button to execute the selection, with brings you to the confirmation screen:

```
Are you sure?
>Y
 N
```

Hit the MIDDLE button to confirm that you want to power on all partitions. The screen returns to the main menu:

```
T1.0-11875
Show Box
>Power On
Power Off
```

none of the buttons are pressed for a few seconds, the display automatically returns to the default screen where alerts are displayed. You should see the self-test activity there, indicating powering up:

```
Running Test 10
Running Test 11
Running Test 12
Running Test 13
```

Other Possible Menu Commands are:

```
Reset
Halt in
Halt out
```

Operation

Server Management Overview

Server management consists of a distributed set of microprocessors communicating with each other over either a private LAN or a point-to-point communication. Each dual-processor module has a CPU module manager (CMM). Up to four CMMs are connected via hardwired serial lines to one module backplane manager (MBM). In addition to having all MBMs connected via the private LAN, PCI backplane managers (PBMs) contained within the IO subsystem are connected to the same LAN.

The modular system building block (SBB) is an 8P or a 2P drawer, consisting of dual-processor modules and an MBM. Each dual-processor module consists of two CPUs, one CMM, up to ten voltage regulator modules (VRM), and up to ten memory modules (RIMMs) for each CPU. The CPUs are interconnected at four compass points (North, South, East, and West) via the backplane and between backplanes with interprocessor (IP) cables. Each CPU can optionally be connected to a PCI I/O subsystem via an I/O cable.

Server management software allows the user to boot operating systems, configure systems, monitor system functions, run debugging tools, and update firmware.

There are two graphical interfaces:

1. the Server Platform Manager (SPM), and
2. the AlphaServer Management Utility (AMU).

And there are two command line interfaces:

1. the Command Line Interface (CLI), and
2. the SRM console.

The user can access a system one of three ways:

- directly through the LAN management HUB
- over a local area network (LAN), or
- over a wide area network (WAN).

The Server_Management_Tutorial gives you more information. HP's website has Additional Documentation on each of these interfaces.

Server Platform Manager (SPM)

You use the SPM to manage one or more EV7 systems. The SPM runs on the AlphaServer management station that connects through a LAN or a WAN. The SPM allows you to manage partitions, boot operating systems, and generate information about error logs.

AlphaServer Management Utility (AMU)

The AMU runs on the SPM, or independently from a browser running on a PC or laptop connected to a LAN or WAN.

Command Line Interface (CLI)

You access the CLI using a PC connected directly to the LAN management HUB. Use the CLI (through a telnet session on the private LAN) to perform several server management functions such as displaying configuration and hardware information, updating firmware, powering on and off, providing commands for partitioning and cabling functions, and implementing remote server functions. In addition, the CLI enables connection to the SRM console. The CLI prompt is MBM>.

For more information on the CLI, see *CLI Reference* section on the *Server Management* CD.

SRM Console

The SRM console is used to boot the operating system and to perform other server management functions, such as running tests, turning power on and off, and displaying system status. You enter the SRM console through the SPM, or by using the CLI **`connect`** command.

The SRM console prompt is `P00>>>`.

For more information on the SRM console, see the *SRM Console Reference* on the *Server Management* CD.

System Management Requirements

Each AlphaServer ES47, ES80 and GS1280 system includes System Management software that can significantly enhance and simplify monitoring and control of the system. Use of the System Management software is optional. The software, which runs on a separate Intel or Alpha system, consists of two major components:

1. Alpha Management Station (AMS) - for monitoring and control of multiple ES47, ES80, and GS1280 Alpha Systems. AMS offers the highest level of server management functionality for a single or multi-platform environment. T

2. Alpha Management Utility (AMU) - for monitoring and control of a single ES47, ES80, and GS1280 Alpha System. The AMU is a GUI based application that provides a sophisticated, yet user-friendly graphics interface. The AMU is a Web-based utility, which allows a user local and remote access from a browser.

For more information:

- For up to date requirements, see ***Step 1b "System Management Hardware/Software Requirements"*** in the QuickSpecs for the ES47/ES80 or the GS1280 systems.

- AMS/AMU software kits and instructions may be downloaded from: http://ftp.digital.com/pub/Digital/Alpha/firmware/interim/ams/index.html

System Management Hardware

The following is a list of the modules and devices that make up the system management hardware.

Component	Function
AMS	
multi-server LAN	A LAN controlled by the AMS to which ES47/ES80/GS1280 and GS80/160/320 may sit.
Internal LAN	A LAN internal to ES47/ES80/GS1280 platforms with a distributed platform management utility capable of running on any of the microprocessors attached to it.
MBM	The module backplane manager which controls the 2P or 8P drawer.
PBM	The PCI backplane manager which controls the PCI drawer.
CMM	CPU module manager controls the CPU module at the command of the MBM.
OCP	The three OCPs: the cabinet OCP, the 2P or 8P drawer OCP, and the I/O drawer OCP.
PI	The power interface module that controls and monitors the 48V power supplies in power subracks.
Power supplies	Power supplies monitoring their own environment
I^2C busses and their devices	I^2C busses on backplanes and modules and their devices that monitor the environment and store the module history and serial numbers.

The platform management utility runs on hardware independent of the system powered by auxiliary power that is present when the system is plugged in and breakers are on. If the system management hardware fails, the system continues to run.

The hardware consists of a private local area network known as the internal LAN. This network's hub is cabled to the backplane manager modules, MBMs, in processor drawers, the I/O backplane manager modules, PBMs, in I/O drawers, and a network address translator or NAT box connected to a multi-server LAN. The system management console or SMC, also a node on

the multi-server LAN, controls all systems on this network. Some customers may want to control systems on the multi-server LAN from their offices. To enable this option the System Management Console becomes a node on the corporate LAN to which folks in offices have access.
The MBM controls and monitors the processor drawer. It connects to the CPU management module, to the operator control panel, to the cabinet's OCP, and to the power interface module in the drawer's power subrack.
The I/O backplane manager module monitors the state of the I/O drawer. Each PBM is connected to the I/O drawer's OCP, to the cabinet's OCP, and the LAN HUB.

Environmental conditions and maintenance data is tracked and stored by devices on the I^2C buses located on the backplane managers, CMM, power supplies, backplanes and the PI module. These devices include: LEDs; monitors that sense temperature, power voltage, and fan speed; and registers that control power sequencing, and contain module serial numbers, maintenance history, error information, and configuration information. If the data indicates a serious hazardous condition, the server management firmware can shut down all or part of the system. This information is also available to server management software.

For more information:

- For up to date requirements, see ***Step 1b "System Management Hardware/Software Requirements"*** in the QuickSpecs for the ES47/ES80 or the GS1280 systems.

Environmental Specifications

The physical environment surrounding an operating or stored computer plays an important role in the long-term reliability of the electronic equipment and peripherals. Computer rooms, office areas, and industrial sites present varying environmental conditions that may affect the operation of the computer equipment. Environmental parameters associated with contamination and corrosion may need to be considered, evaluated, and possibly controlled during the computer site preparation process.

Environmental Specifications	GS1280 Systems	ES47/ES80 Systems
Temperature Operating Non-operating	10 to 35ºC (41 to 95 ºF) –40 to 66ºC (-40 to151ºF)	10 to 40ºC (50 to 104 ºF) – 40 to 66ºC (-40 to 151ºF)
Relative humidity (noncondensing) Operating Non-operating	10 to 90% 10 to 95%	10 to 90% 10 to 95%
Maximum altitude Operating Non-operating	3050 m (10,000 ft) 12,200 m (40,000 ft)	3050 m (10,000 ft) 12,200 m (40,000 ft)
Shock	5G 30ms, half sine	5G 30ms, half sine
Vibration	.1 G, 10500 Hz	.1 G, 10500 Hz
Acoustics Sound Power (LwAd B)	7.5 (Model 8) 7.8 (Model 16) 8.1 (Model 32) 8.4 (Model 64)	6.6 (ES47/ES80 Model 2) 6.9 (ES47/ES80 Model 4) 7.2 (ES80 Model 6) 7.5 (ES80 Model 8)
Sound Pressure (LpAm dBA)	57 (Model 8) 60 (Model 16) 63 (Model 32) 66 (Model 64)	48 (ES47.ES80 Model 2) 51 (ES47/ES80 Model 4) 54 (ES80 Model 6) 57 (ES80 Model 8)
Heat Dissipation (Watts) Maximum	3,555W (Model 8) 7,110W (Model 16) 14,820W (Model 32)	3000W (ES47/ES80)

Typical	28,440W (Model 64)	
		2000W (ES47/ES80)
	2,430W (Model 8)	
	3,700W (Model 16)	
	8,000W (Model 32)	
	14,800W (Model 64)	

SRM Console

From the SRM console, you set up and boot the operating system, display the system configuration, and perform other tasks. For complete information on the SRM console, see the *Server Management SRM Console Reference CD.*.

SRM Command Overview

Command	Effect
`Boot`	Boots the operating systems.
`Continue`	Resumes processing after a Ctrl/P is issued (OpenVMS systems).
`Crash`	Forces a crash dump of the operating system.
`Edit`	Invokes the console line editor, which can be used to edit a RAM file or the user power-up script, "nvram," which is always invoked during the power-up sequence.
`help (or man)`	Displays information about all or a specific SRM command.
`Init`	Stores any changes made to environment variables and reinitializes the hardware.
`more [filename]`	Displays a file one screen at a time.
`set envar`	Sets the value of an environment variable.
`show envar`	Displays the state of all or a specified environment variable.
`show config`	Displays the configuration at the last system initialization.
`show device`	Displays the controllers and bootable devices in the system.
`show fru`	Displays the configuration of field-replaceable units (FRUs).
`show memory`	Displays memory module information.
`show pal`	Displays the versions of Tru64 UNIX and OpenVMS PALcode.
`show version`	Displays the version of the SRM console program.
`Test`	Tests the entire system.

Command Line Interface (CLI)

The CLI is the MBM serial interface. The CLI implements remote server management, connects to the SRM console, shows component status, temperatures, and configuration information, and provides partitioning functions. The user can update firmware and turn system power on and off using the CLI. For more information, see the *CLI Reference* on the *Server Management* CD.

Displaying the System Configuration

View the system hardware configuration from the SRM console. It is useful to view the hardware configuration to ensure that the system recognizes all devices, memory configuration, and network connections.

Use the following SRM console commands to view the system configuration.

Parameter	Function
`show config`	Displays the logical configuration of interconnects and buses on the system and the devices found on them.
`show device`	Displays the bootable devices and controllers in the system.
`show memory`	Displays configuration of main memory.

Setting SRM Environment Variables

You may need to set several SRM console environment variables and built-in utilities to configure systems running the *Tru64 UNIX* or *OpenVMS* operating systems. For more information on environment variables, see the *Server Management SRM Console Reference CD.*

Set environment variables at the `P00>>>` prompt.

- To check the setting for a specific environment variable, enter the **`show envar`** command, where the name of the environment variable is substituted for `envar`.

- To reset an environment variable, use the **`set envar`** command, where the name of the environment variable is substituted for `envar`.

Changing the Default Boot Device

You can change the default boot device for *Tru64 UNIX* or *OpenVMS* with the **set bootdef_dev** command.

With the *Tru64 UNIX* or *OpenVMS* operating systems, you can designate a default boot device. You change the default boot device by using the **set bootdef_dev** command. For example, to set the boot device to the IDE CD-ROM, enter commands similar to the following:

```
P00>>> show bootdef_dev
 bootdef_dev dka400.4.0.1.1
 P00>>> set bootdef_dev dqa500.5.0.1.1
 P00>>> show bootdef_dev
 bootdef_dev dqa500.5.0.1.1
```

Partitions

For information on partitions, read the Getting Started with Partitions online tutorial.

A single AlphaServer ES47/ES80or GS1280 can be divided into logical hardware partitions as small as two processors, each running an instance of Tru64 UNIX, OpenVMS, or Linux. Each partition is allocated its own dedicated "shared-nothing" set of hardware resources: CPU module(s), memory module(s), and I/O. Each hardware partition is viewed as a unique node, from a system point-of-view, with its own instance of Tru64 UNIX or OpenVMS operating system and application software, independent system console, and error log.

For example, consider the 32P platform shown here.

- Each of the 8P drawers contains processors and memory.
- Each of the I/O drawers contains a PCI card cage with related I/O capabilities.
- Each partition must have some CPU resources, some memory, and some I/O capability.

In this way, a partition provides the attributes of a separate system, capable of functioning independently.

In the example, consider all the yellow boxes to comprise one partition; all the green boxes, a second partition; and all the blue boxes, a third partition.

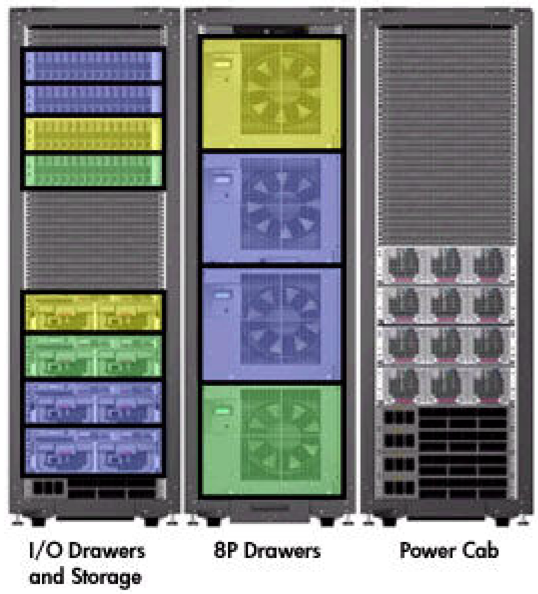

Configuration Requirements for Partitions

- Partitions must contain at least one Dual Processor CPU Module.
- At least one of the processors in a partition must be connected to I/O.
- The two processors on a Dual Processor CPU Module cannot be split between hard partitions.
- The set of processors assigned to a partition must form a continuous rectangle on the system interconnect mesh network.
- Rules for the maximum number of supported options apply to each hard partition.
- **For GS1280 systems:** CPU modules in an 8P Building Block Drawer share the same 48v power feed (from a redundant power supply). If the 8P Building Block Drawer contains all or portions of several hard partitions, a power fault in the drawer could affect all the partitions in that drawer. Hard partitions that consist of whole 8P Building Block Drawers limit the risk of a power fault in one drawer to one partition.
- Check the QuickSpecs for current minimum firmware and software revisions to run hard partitions

Routing the Mesh

There are three basic routing rules followed in the firmware as it brings the system up.

1. The firmware routes as many desired CPUs into a mesh as it can.
2. Messages may pass straight through EV7 system chips. I.E. a message can come in through an East port, pass through the system chip's router, and exit through the West port.
3. A messages may make one, and only one, turn in a system chip on its way from one CPU to another.

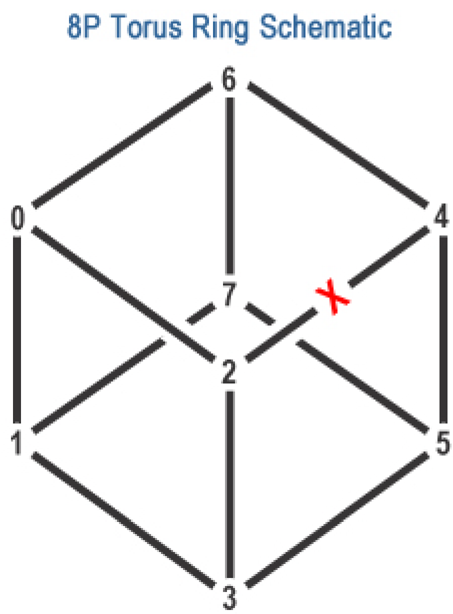

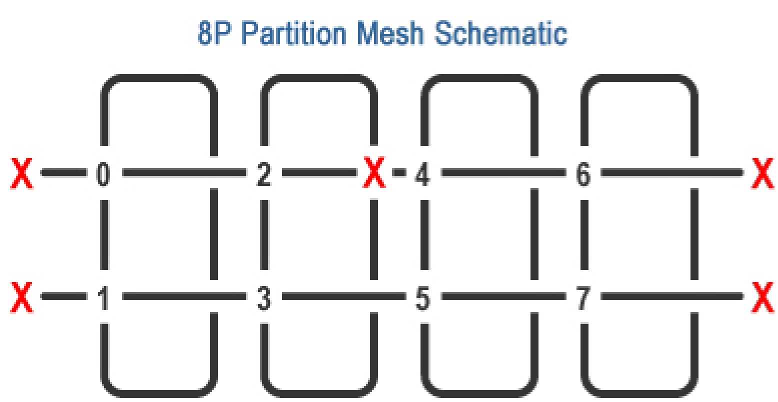

As a consequence of these rules, creating a legal routing system, particularly when the server is partitioned, may lead to a partition that does not use all CPU resources allocated to it. Under normal circumstances, if a CPU in the mesh fails, the system crashes and then automatically reboots. When it reboots an attempt is made to configure the system the way it was before the crash. The firmware discovers that one CPU fails to come up and that CPU is part of a grid like mesh rather than a torus ring. A possible result of rule 3 is that not all CPUs will be able to communicate with each other.

Routing failure in Torus ring

Looking at the torus ring graphic above, if the connection between CPU2 and CPU4 fails in the ring, all eight CPUs in the ring power-up and are part of the partition because, when routing connections wrap, it is possible for any CPU to get to another CPU in the ring by making only one "right angle turn."

Routing failure in a partitioned system

On the other hand, if the same connection between CPU2 and CPU4 fails and the system is partitioned, CPU2 could not get to CPU4 or CPU6 without making two "right turns." Therefore, the system will have a partition of six CPUs instead of the original 8.

Additional Documentation

AlphaServer QuickSpecs contain up to date information on options and new features as they are added. The URL for Quick Specs is http://h18000.www1.hp.com/products/quickspecs/North_America/10410.html.

The most current documentation for hp AlphaServer ES47/ES80/GS1280 systems is on the hp website at http://h18002.www1.hp.com/alphaserver/es.html.

To help you jump right to the page for your system:

- ES40 documents are at http://h18002.www1.hp.com/alphaserver/es47/es47_tech.htm
- ES80 documents are at http://h18002.www1.hp.com/alphaserver/es80/es80_tech.html and
- GS1280 documents are at http://h18002.www1.hp.com/alphaserver/gs1280/gs1280_tech.html

Descriptions are given below for your reference.

Server Management documentation:

- The *Server Management Tutorial* introduces server management features for the hp AlphaServer ES47/ES80/GS1280 systems.

- The *AlphaServer Management Station (AMS) Users Guide* describes the procedures for setting up and managing a management station on these systems.

- The *AlphaServer Management Utility* (AMU) is a Web-based application that provides an access to AlphaServer ES47/ES80/GS1280 systems. AMU can be used in standalone mode or in integration with AMS.

- The Server Management *SRM Console Reference* provides a reference for the SRM console commands.

- The *Server Platform Manager (SPM)* is a Web-based application that provides a single point of access for the AlphaServer management utilities. With the SPM, an administrator can manage multiple GS1280, ES80, and ES47 servers from a single AlphaServer Management Station (AMS).

- The *Server Management Command Line Interface (CLI) Reference manual* provides information on the concepts of partitioning these systems. Specific information is also given on how to configure, plan, and construct hard and soft partitions on these platforms, using the three tools available: the AlphaServer Management Utility (AMU), the AlphaServer Platform Wizard (APW), and the Server Management Command Line Interface (CLI).

Other:

- *Site Preparation* gives you information on maintaining the environment for your systems.

- *Getting Started with Partitions* is an online tutorial on partitioning for the hp AlphaServer ES47/ES80/GS1280 platform operating systems.

- The *AlphaServer Partition Wizard (APW)* is a graphical application that simplifies the creation and management of partitions on AlphaServer ES47/ES80/GS1280 platforms

Glossary

2P drawer A chassis with backplane that supports one dual processor module, and five PCI/PCI-X slots and one AGP slot.

8P drawer A chassis with backplane that supports four dual processor modules.

AGP Accelerated Graphics Processor.

AMS AlphaServer Management Station; a computer, software, and terminal server used to manage the system.

AMU AlphaServer Management Utility. A dedicated utility used to view, monitor and configure a particular AlphaServer ES47/ES80/GS1280. AMU manages both partitions and physical platform connections.

APW AlphaServer Partition Wizard. A tool to assist the system manager to configure partitions.

Backplane manager See MBM.

Cable interconnect module A module in the 2P drawer that provides connectors for a modem port, USB port, and LAN port.

CMM CPU Management Module. A plug-in card on the dual processor module that provides local module power and initialization control.

Corporate LAN A conventional LAN (or WAN) that can be used for remote management by connecting to the Multi-Server LAN.

CPU Building Block module See dual processor module.

CPU module See dual processor module.

Dual processor module A module containing two Alpha EV7 system chips, memory modules, voltage regulator modules (VRMs), and a CPU management module (CMM).

External LAN See Corporate LAN.

EV7 system chip The Alpha chip containing processor logic, second-level cache, two memory controllers, an I/O port, and four interprocessor ports, which function as the system interconnect.

Galaxy OpenVMS software that is used to manage soft partitions.

Hard partition A subset of a system's computing resources that cannot exchange information or resources with any other partition on the system. The boundaries are maintained by a switch in the EV7 system chip. Faults are not propagated across hard partition boundaries.

High-performance I/O drawer An enclosure that has eight high-speed (133 MHz) PCI-X buses, with four I/O riser modules.

I/O drawers Two variants are available. See High-performance I/O drawer, Standard I/O drawer.

I/O expander module A module in the 2P drawer used to provide backplane manager logic and controllers for CD-ROM, SCSI disks, LAN, keyboard, mouse, and modem.

I/O port Logic that provides an interface from the EV7 system chip to the IO7 chip on I/O riser modules.

I/O riser module Module containing the IO7 chip that functions as the interconnect between the EV7 chip and PCI, PCI-X, and AGP buses. The standard I/O drawer has one I/O riser module; the high-performance I/O drawer can have up to four I/O riser modules.

Instance An operating system running in a partition.

Internal LAN A local network within a single system for managing at the lowest level (MBM, CMM). Implemented using an Ethernet hub. A NAT box can interface the Internal LAN to the Multi-Server LAN.

Interprocessor (IP) connections The physical connection between the IP ports of EV7 chips. Interconnect logic is integrated into the EV7 system chip. The interconnect between the processors on a dual processor module is implemented by etch on the module; the interconnect between the EV7 chips on dual-processor modules to chips on other dual-processor modules is implemented either by etch on the 8P backplane or by interprocessor cables. By convention, north ports connect to south ports, and east ports connect to west ports.

Interprocessor (IP) ports Ports in the EV7 system chip designated north, south, east, and west through which the processors communicate.

IO7 chip The interface between an EV7 chip and I/O buses. Each chip can support three PCI/PCI-X buses and one AGP bus.

LAN management hub Hardware hub on the Internal LAN.

LFU Loadable Firmware Update Utility; utility used to update firmware located in various places throughout the system.

Loadable Firmware Update Utility See LFU.

MBM Backplane manager. A module on the backplanes of both the 2P and 8P drawers that controls the CPU management modules (CMMs) and has logic to monitor and control environmental conditions in the drawer.

Memory controller Integrated into the EV7 system chip. Each of the two controllers drives four Rambus data channels and a fifth channel for error detection and correction.

Mesh A grid of connected EV7 system chips. All IP connections at the edges of the grid are not necessarily complete. Compare Torus ring.

Multi-Server LAN An Ethernet connected to each server's router (NAT box), allowing the system manager to manage one or more AlphaServers from the AlphaServer Management Station using high-level tools including SPM and AMU.

NAT box The Network Address Translator box. Found on hp AlphaServer ES47/ES80/GS1280 systems, the NAT box is not part of the Internal LAN. It is programmed to have a unique address on the Multi-Server LAN, and translate requests to this address to specific components within the Internal LAN via the LAN management hub.

OCP Operator Control Panel. One is located in the cabinet, and another is located in each I/O drawer.

Cabinet OCP A control panel with an On, Off, and Secure switch , a switch for setting the cabinet ID, status LEDs, an LCD display, and Scroll, Select, and Clear pushbuttons for navigating system menus and executing commands.

I/O drawer OCP A control panel with a pushbutton ID select switch for setting the I/O drawer ID, two status LEDs, and an LCD display for system status information.

Partition A subset of a system's computing resources, each of which is capable of running a copy, or instance, of an operating system. See also Hard partition, Soft partition.

PBM PCI Backplane Manager. Monitors and controls the activity and environment in the I/O drawers.

PCI Backplane Manager See PBM.

Platform System.

Platform Management LAN See Multi-Server LAN

PCM Platform Console Manager. Used to establish connections to consoles of managed systems and display the status and latest output of each console.

Power distribution unit AC input boxes.

Power subrack A cabinetthat can contain three 48 volt power supplies that provide 48 VDC and Vaux; used with the 8P drawers.

Private LAN See Internal LAN.

RIMMs Rambus in-line memory modules.

Server Management LAN See Internal LAN.

Server Platform Manager See SPM.

SIO Standard I/O module. An optional module in I/O drawers that provides controllers for CD-ROM, SCSI disks, keyboard, mouse, and modem.

SMC LAN See Multi-Server LAN.

Soft partition A subset of a hard partition's computing resources. There are no hardware boundaries between soft partitions. Using resource management software like OpenVMS Galaxy, resources may be moved through agreement of the operating system instances in each soft partition. Hardware faults are propagated throughout the mesh of soft partitions.

SPM Server Platform Manager. The top-level management application that provides management functions to the user of the AlphaServer Management Station.

SRM console Firmware on the backplane manager module that provides a command-line interface for operator control of the system or of a partition. The SRM console is responsible for booting the operating system and passing system configuration data, discovered during power-up, to it.

Standard I/O drawer An enclosure, with eleven PCI/PCI-X slots and one AGP slot, that contains a single I/O riser module. An optional standard I/O module may be present to control an optional CD-ROM drive and SCSI storage drives.

System building block A drawer containing:
- one CPU building block module (including two EV7 system chips)
- four memory option slots
- five PCI-X slots

- one AGP 4X slot
- two redundant power supplies
- two disk drive bays

System management LAN See Internal LAN.

Torus ring A doughnut-shaped mesh in which all interprocessor ports are connected; east ports are connected to west ports, and north ports are connected to south ports. Compare Mesh.

Vaux Low voltage power present in the system whenever the power cord is plugged in and the circuit breakers are on. Vaux powers nodes on the system management LAN.

VRM Voltage regulator modules reside on the dual processor module and convert 48 V DC power into 1.5V, 1/8V and 2.5V DC power for use on the module.