


Overview



This subchapter pertains to all HP 9000 Superdome servers (running PA 8600, PA 8700, PA 8800 or PA-8900 processors) for all markets. With Superdome, HP launches a new strategy to ensure a positive Total Customer Experience is achieved via industry leading HP Services. Our experience has shown us that large solution implementations most often succeed as a result of appropriate skills and attention being applied to the solution design and implementation. To address this on the implementation side, for Superdome, HP is responding to Customer and Industry feedback and delivering Superdome Configurations via three, pre configured Services levels: Critical Service, Proactive Service and Foundation Service. With Superdome, we introduced a new role, the TCE Manager, who manages the fulfillment of an integrated business solution based on customer requirements. For each customer account, the TCE Manager will facilitate the selection of the appropriate configuration. For ordering instructions, please consult the ordering guide.

Superdome Service Solutions

Superdome continues to provide the same positive Total Customer Experience via industry leading HP Services as with existing Superdome servers. The HP Services component of Superdome is described here:

Solution Life Cycle

HP customers have consistently achieved higher levels of satisfaction when key components of their IT infrastructures are implemented using the **Solution Life Cycle**. The Solution Life Cycle focuses on rapid productivity and maximum availability by examining customers' specific needs at each of five distinct phases (plan, design, integrate, install, and manage) and then designing their Superdome solution around those needs.

Overview

Service Solutions

HP offers three pre-configured service solutions for Superdome that provide customers with a choice of lifecycle services to address their own individual business requirements.

- **Foundation Service Solution:** This solution reduces design problems, speeds time-to-production, and lays the groundwork for long-term system reliability by combining pre-installation preparation and integration services, hands-on training and reactive support. This solution includes HP Support Plus 24 to provide an integrated set of 24×7 hardware and software services as well as software updates for selected HP and third party products.
- **Proactive Service Solution:** This solution builds on the Foundation Service Solution by enhancing the management phase of the Solution Life Cycle with HP Proactive 24 to complement your internal IT resources with proactive assistance and reactive support. Proactive Service Solution helps reduce design problems, speed time to production, and lay the groundwork for long-term system reliability by combining pre installation preparation and integration services with hands on staff training and transition assistance. With HP Proactive 24 included in your solution, you optimize the effectiveness of your IT environment with access to an HP-certified team of experts that can help you identify potential areas of improvement in key IT processes and implement necessary changes to increase availability.
- **Critical Service Solution:** Mission Critical environments are maintained by combining proactive and reactive support services to ensure maximum IT availability and performance for companies that can't tolerate downtime without serious business impact. Critical Service Solution encompasses the full spectrum of deliverables across the Solution Lifecycle and is enhanced by HP Critical Service as the core of the management phase. This total solution provides maximum system availability and reduces design problems, speeds time-to-production, and lays the groundwork for long-term system reliability by combining pre-installation preparation and integration services, hands on training, transition assistance, remote monitoring, and mission critical support. As part of HP Critical Service, you get the services of a team of HP certified experts that will assist with the transition process, teach your staff how to optimize system performance, and monitor your system closely so potential problems are identified before they can affect availability.

Other Services

- **HP's Mission Critical Partnership:** This service offering provides customers the opportunity to create a custom agreement with Hewlett-Packard to achieve the level of service that you need to meet your business requirements. This level of service can help you reduce the business risk of a complex IT infrastructure, by helping you align IT service delivery to your business objectives, enable a high rate of business change, and continuously improve service levels. HP will work with you proactively to eliminate downtime, and improve IT management processes.
- **Service Solution Enhancements:** HP's full portfolio of services is available to enhance your Superdome Service Solution in order to address your specific business needs. Services focused across multi-operating systems as well as other platforms such as storage and networks can be combined to compliment your total solution.

Standard Features

NOTE: Given that PA 8600/PA 8700 are single core processors and PA 8800 and PA 8900 are dual core processors, the columns listed in this table refer to 16 processor, 32 processor and 64 processor. This terminology refers to 16 processor, 32 processor and 64 processor for Superdome PA 8600/PA 8700 and 32 processor cores, 64 processor cores and 128 processor cores for Superdome PA 8800 and PA 8900 systems.

System Size	Minimum System		Maximum SPU Capacities	
16 Processors	PA 8600 or PA 8700	PA 8800 or PA 8900	PA 8600 or PA 8700	PA 8800 or PA 8900
Processor cores	1	2	16	32
Memory	2 GB	2 GB	64 GB	256 GB
Cell Boards	1 Cell Board	1 Cell Board	4 Cell Boards	4 Cell Boards
PCI Chassis	1 12 slot chassis	1 12 slot chassis	4 12 slot chassis	4 12 slot chassis
32 Processors				
Processor cores	1	2	32	64
Memory	2 GB	2 GB	128 GB	512 GB
Cell Boards	1 Cell Board	1 Cell Board	8 Cell Boards	8 Cell Boards
PCI Chassis	1 12 slot chassis	1 12 slot chassis	8 12 slot chassis	8 12 slot chassis
64 Processors				
Processor cores	6	6	64	128
Memory	6 GB	6 GB	256 GB	1024 GB (1)
Cell Boards	3 Cell Boards	3 Cell Boards	16 Cell Boards	16 Cell Boards
PCI Chassis	1 12 slot chassis	1 12 slot chassis	16 12 slot chassis	16 12 slot chassis
Standard Features	<ul style="list-style-type: none"> • Redundant Power Supply • Redundant Fans • HP-UX operating system with unlimited user license • Factory Integration of memory and I/O Cards • Installation Guide, Operators Guide, and Architecture Manual • HP Site planning and Installation • One year warranty with same business day on-site service response 			

Configuration

There are three basic building blocks in the Superdome system architecture: the cell, the crossbar backplane and the I/O subsystem. Please note that Superdome with PA 8800 and PA-8900 is based on a different chip set (sx1000) than Superdome with PA 8600 or PA 8700. For more information on the sx1000 chip set, please refer to the following: [HP Integrity Superdome QuickSpec](#).

Cabinets

A Superdome system can consist of up to four different types of cabinet assemblies:

- At least one Superdome left cabinet. The Superdome cabinets contain all of the processors, memory and core devices of the system. They will also house most (usually all) of the system's I/O cards. Systems may include both left and right cabinet assemblies containing, a left or right backplane respectively.
- One or more HP Rack System/E cabinets. These 19-inch rack cabinets are used to hold the system peripheral devices such as disk drives.
- Optionally, one or more I/O expansion cabinets (Rack System/E). An I/O expansion cabinet is required when a customer requires more PCI cards than can be accommodated in their Superdome cabinets.

Superdome cabinets are serviced from the front and rear of the cabinet only. This enables customers to arrange the cabinets of their Superdome system in the traditional row fashion found in most computer rooms. The width of the cabinet accommodates moving it through common doorways in the U.S. and Europe. The intake air to the main (cell) card cage is filtered. This filter can be removed for cleaning/replacement while the system is fully operational.

A status display will be located on the outside of the front and rear doors of each cabinet. The customer and field engineers can therefore determine basic status of each cabinet without opening any cabinet doors.

For PA-8600 and PA-8700 processors (single core per processor):

- Superdome 16 processor: single cabinet (a left cabinet)
- Superdome 32 processor: single cabinet (a left cabinet)
- Superdome 64 processor: dual cabinet (a left cabinet and a right cabinet)

For PA-8800 and PA-8900 processors (dual core per processor):

- Superdome 32 processor cores: single cabinet (a left cabinet)
- Superdome 64 processor cores: single cabinet (a left cabinet)
- Superdome 128 processor cores: dual cabinet (a left cabinet and a right cabinet)

Each cabinet may contain a specific number of cell boards (consisting of processors and memory) and I/O. See the following sections for configuration rules pertaining to each cabinet. The base configuration product numbers for each of the models are as follows:

Configuration

Cells (Processors and Memory)

A cell, or cell board, is the basic building block of a Superdome system. It is a symmetric multi-processor (SMP), containing:

- Four processors that can hold PA 8600, PA 8700, PA 8800 or PA-8900 processors (only one type is allowed per cell)
- Memory (up to 16-GB RAM with 512-MB DIMMs, 32-GB RAM with 1-GB DIMMs and 64-GB with 2-GB DIMMs.
NOTE: 1 GB and 2 GB DIMMs are only supported with Superdome PA 8800 and PA-8900.
- One cell controller (CC)
- Power converters
- Data buses
- Optional link to I/O card cage

Please note the following:

- PA-8600 and PA-8700 processors can be mixed within a Superdome system, but must reside on different cells and in different partitions. PA-8800 and PA-8900 processors cannot be mixed within a Superdome system with PA-8600 or PA-8700 processors.
- A Superdome PA-8800 and PA-8900 system can support up to 1 TB of memory with 2 GB DIMMs. However, due to the limitation of HP-UX 11i, the maximum amount of memory supported in a partition is 512 GB.
- For PA-8600 and PA-8700 processors: The minimum configuration includes one active processor and 2 GB memory per cell board. The maximum configuration includes four active processor and 16 GB memory per cell board. Each cell board ships with four PA-8600 or PA-8700 processors. However, based on the number of active processors ordered, from one through four processors are activated prior to shipment.
- For PA 8800 and PA-8900 processors: The minimum configuration includes four active processors and 2 GB of memory per cell board. The maximum configuration includes eight active processors and 64 GB of memory per cell board using 2 GB DIMMs. Each cell board ships with either four or eight PA 8800 or PA-8900 processors. The minimum number of active processors on each cell board is 4. Additional processors on the cell board can be activated individually (not in pairs).
- 512 MB, 1 GB and 2 GB (future) DIMMs can be mixed on a cell board in Superdome PA 8800 and PA-8900 systems only.

Crossbar Backplane

Each Crossbar backplane contains two sets of two crossbar chips that provide a non blocking connection between eight cells and the other backplane. Each backplane cabinet can support up to eight resulting in a Superdome PA 8600 or Superdome PA 8700 32-processor or Superdome PA 8800 or PA 8900 64-processor system. A backplane supporting four cells would result in a Superdome PA 8600 or Superdome PA 8700 16- processor and a Superdome PA 8800 or PA 8900 32-processor system. Two backplanes can be linked together with flex cables to produce a cabinet that can support up to 16 cells resulting in a Superdome PA 8600 or Superdome PA 8700 64-processor core or Superdome PA 8800 or PA 8900 128-processor core system.

Configuration

I/O Subsystem

Each I/O chassis provides twelve I/O slots. Superdome PA-8600 and Superdome PA-8700 support I/O chassis with 12 PCI 66 capable slots, eight supported via single (1x) ropes (266 MB/s peak) and four supported via dual (2x) ropes (533 MB/s peak). Superdome PA-8800 and PA-8900 supports I/O chassis with 12 PCI-X 133 capable slots, eight supported via single enhanced (2x) ropes (533 MB/s peak) and four supported via dual enhanced (4x) ropes (1066 MB/s peak).

Each Superdome cabinet supports a maximum of four I/O chassis. The optional I/O expansion cabinet can support three I/O chassis enclosures (ICE), each of which supports two I/O chassis for a maximum of six I/O chassis per I/O expansion cabinet.

Each I/O chassis connects to one cell board and the number of I/O chassis supported is dependent on the number of cells present in the system. A Superdome system can have more cells than I/O chassis. For instance, an 8-cell Superdome can have one to eight I/O chassis. Each partition must have at least one I/O chassis with the number of I/O chassis not exceeding the number of cells.

A 4-cell Superdome supports four I/O chassis for a maximum of 48 PCI slots.

An 8-cell Superdome supports eight I/O chassis for a maximum of 96 PCI slots. Since a single Superdome cabinet only supports four I/O chassis, an I/O expansion cabinet and two I/O chassis enclosures are required to support all eight I/O chassis.

A 16 cell Superdome supports 16 I/O chassis for a maximum of 192 PCI slots. Since two Superdome cabinets (left and right) only support eight I/O chassis, two I/O expansion cabinets and four I/O chassis enclosures are required to support all 16 I/O chassis. The four I/O chassis enclosures are spread across the two I/O expansion cabinets, either three ICE in one I/O expansion cabinet and one ICE in the other or two ICE in each.

Core I/O

Superdome's core I/O provides the base set of I/O functions required by every Superdome partition. Each partition must have at least one core I/O card in order to boot. Multiple core I/O cards may be present within a partition (one core I/O card is supported per I/O backplane); however, only one may be active at a time. Core I/O will utilize the standard long card PCI form factor but will add a second card cage connection to the I/O backplane for additional non-PCI signals (USB and utilities). This secondary connector will not impede the ability to support standard PCI cards in the core slot when a core I/O card is not installed.

Any I/O chassis can support a Core I/O card that is required for each independent partition. A system configured with 16 cells, each with its own I/O chassis and core I/O card could support up to 16 independent partitions. Note that cells can be configured without I/O chassis attached, but I/O chassis cannot be configured in the system unless attached to a cell.

The core I/O card's primary functions are:

- Partitions (console support) including USB and RS-232 connections
- 10/100Base-T LAN (general purpose)

Other common functions, such as Ultra/Ultra2 SCSI, Fibre Channel, and Gigabit Ethernet, are not included on the core I/O card. These functions are, of course, supported as normal add in cards.

The unified 100Base-T Core LAN driver code searches to verify whether there is a cable connection on an RJ-45 port or on an AUI port. If no cable connection is found on the RJ-45 port, there is a busy wait pause of 150 ms when checking for an AUI connection. By installing the loopback connector (description below) in the RJ-45 port, the driver would think an RJ 45 cable was connected and would not continue to search for an AUI connection, hence eliminate the 150 ms busy wait state:

Configuration

Product/Option Number	Description
A7108A	RJ-45 Loopback Connector
OD1	Factory integration RJ-45 Loopback Connector

I/O Expansion Cabinet The I/O expansion functionality is physically partitioned into four rack mounted chassis—the I/O expansion utilities chassis (XUC), the I/O expansion rear display module (RDM), the I/O expansion power chassis (XPC) and the I/O chassis enclosure (ICE). Each ICE supports up to two 12-slot I/O chassis.

Factory Integration When an I/O Expansion cabinet is ordered as an upgrade to a Superdome system, it includes the factory testing and integration of any components that are ordered at the same time as the cabinet. This includes any I/O Chassis, PCI or PCI-X cards and peripherals. If it is ordered as an upgrade but not at the time of the Superdome system, additional installation assistance may be required and can be ordered as field installation products.

Field Racking The only field rackable I/O expansion components are the ICE and the 12-slot I/O chassis. Either component would be field installed when the customer has ordered additional I/O capability for a previously installed I/O expansion cabinet.

No I/O expansion cabinet components will be delivered to be field installed in a customer's existing rack other than a previously installed I/O expansion cabinet. The I/O expansion components were not designed to be installed in racks other than Rack System E. In other words, they are not designed for Rosebowl I, pre-merger Compaq, Rittal, or other third-party racks.

The I/O expansion cabinet is based on a modified HP Rack System E and all expansion components mount in the rack. Each component is designed to install independently in the rack. The Rack System E cabinet has been modified to allow I/O interface cables to route between the ICE and cell boards in the Superdome cabinet. I/O expansion components are not designed for installation behind a rack front door. The components are designed for use with the standard Rack System E perforated rear door.

I/O Chassis Enclosure (ICE) The I/O chassis enclosure (ICE) provides expanded I/O capability for Superdome. Each ICE supports up to 24 I/O slots by using two 12-slot Superdome I/O chassis. The I/O chassis installation in the ICE puts the I/O cards in a horizontal position. An ICE supports one or two 12-slot I/O chassis. The I/O chassis enclosure (ICE) is designed to mount in a Rack System E rack and consumes 9U of vertical rack space.

To provide online addition/replacement/deletion access to I/O cards and hot-swap access for I/O fans, all I/O chassis are mounted on a sliding shelf inside the ICE.

Four (N+1) I/O fans mounted in the rear of the ICE provide cooling for the chassis. Air is pulled through the front as well as the I/O chassis lid (on the side of the ICE) and exhausted out the rear. The I/O fan assembly is hot swappable. An LED on each I/O fan assembly indicates that the fan is operating.

Configuration

Cabinet Height and Configuration Limitations Although the individual I/O expansion cabinet components are designed for installation in any Rack System E cabinet, rack size limitations have been agreed upon. IOX Cabinets will ship in either the 1.6-meter (33U) or 1.96-meter (41U) cabinet. In order to allay service access concerns, the factory will not install IOX components higher than 1.6 meters from the floor. Open space in an IOX cabinet will be available for peripheral installation. Refer to the 10000 G2 Series Rack Best Practices Guide for information on rack deployment, stabilization and transportation. Go to <http://HP.com/go/rackandpower> for more information.

Peripheral Support All peripherals qualified for use with Superdome and/or for use in a Rack System E are supported in the I/O expansion cabinet as long as there is available space. Peripherals not connected to or associated with the Superdome system to which the I/O expansion cabinet is attached may be installed in the I/O expansion cabinet.

Server Support No servers except those required for Superdome system management such as Superdome Support Management Station or ISEE may be installed in an I/O expansion cabinet.

Peripherals installed in the I/O expansion cabinet cannot be powered by the XPC. Provisions for peripheral AC power must be provided by a PDU or other means.

Standalone I/O Expansion Cabinet If an I/O expansion cabinet is ordered alone, its field installation can be ordered via option 750 in the ordering guide (option 950 for Superdome Advanced Architect Program Channel partners).

DVD Solution The DVD solution for Superdome requires the following components per partition. External racks A4901A and A4902A must also be ordered with the DVD solution.

NOTE: One DVD and one DAT is required per nPartition.

Superdome DVD Solutions

Description	Part Number	Option Number
PCI Ultra160 SCSI Adapter or PCI Dual Channel Ultra160 SCSI Adapter	A6828A or A6829A	0D1
Surestore Tape Array 5300	C7508AZ	
HP DVD+RW Array Module (one per partition) (NOTE: HP DVD-ROM Array Module for the TA5300 (C7499B) is replaced by HP DVD+RW Array Module (Q1592A) to give customers read capabilities for loading software from CD or DVD, DVD write capabilities for small amount of data (up to 4 GB) and offline hot-swap capabilities.)	Q1592A	0D1
HP DAT 40m Array Module	C7497A	0D1
HP DAT 72 Array Module (carbon)	Q1524B	
DDS-4/DAT40 (one per partition)	C7497B	0D1
PCI LVD/SE SCSI card NOTE: A5149A supports the MSA30 SB/DB as a boot device on Superdome running HP-UX 11i.	A5149A	0D1
Jumper SCSI Cable for DDS-4 (optional) ¹	C2978B	0D1
SCSI Cable 10m VHD/HDTS68 LVD/SE ILT M/M	C7556A	
SCSI cable 5 meter	C7520A	0D1

Configuration

SCSI Terminator	C2364A	0D1
-----------------	--------	-----

¹ 0.5-meter HD-HDTS68 is required if DDS-4 is used.

If using DAT72, it is recommended to use an A6829A dual port SCSI with daisy chaining to connect the DVD and the DAT72 leaving the second port available to connect a SCSI data storage device.

Partitions

Hardware Partitions

A hardware partition (nPar) consists of one or more cells that communicate coherently over a high bandwidth, low latency crossbar fabric. Individual processors on a single cell board cannot be separately partitioned. nPars are logically isolated from each other such that transactions in one nPar are not visible to other nPars within the same complex.

Each nPar runs its own independent operating system. Different nPars may be executing the same or different revisions of an operating system. On HP Integrity Superdome systems, nPars may be executing different operating systems altogether (HP-UX, Windows Server 2003 or Linux). Please refer to the HP Integrity Superdome section for details on these operating systems.

Each nPar has its own independent processors, memory and I/O resources consisting of the resources of the cells that make up the nPar. Resources may be removed from one nPar and added to another without having to physically manipulate the hardware just by using commands that are part of the System Management interface.

Superdome HP-UX 11i supports static nPars. Static nPars imply that any nPar configuration change requires a reboot of that nPar. In a future HP-UX release, dynamic nPars will be supported. Dynamic nPars imply that the nPar configuration changes do not require a reboot of that nPar. Using the related capabilities of dynamic reconfiguration (i.e. on-line addition, on-line removal), new resources may be added to an nPar and failed modules may be removed and replaced while the nPar continues in operation.

NOTE: It is possible for PA 8800 and PA 8900 nPars to co exist with Itanium 2 1.6 GHz nPars in the same Superdome system, but on different partitions. Customers can configure an Itanium 2 1.6 GHz nPar in an HP 9000 Superdome running PA 8800 or PA 8900 (and vice versa) in the field only . Factory orders for mixed Itanium and PA RISC nPars are not allowed.

NOTE: It is possible for PA 8800 and PA 8900 nPars to co exist with Itanium 2 1.6 GHz nPars in the same Superdome system, but on different partitions. Customers can configure an Itanium 2 1.6 GHz nPar in an HP 9000 Superdome running PA 8800 or PA 8900 (and vice versa) in the field only. Factory orders for mixed Itanium and PA RISC nPars are not allowed.

Virtual Partitions

A virtual partition (vPar) provides the capability of dividing a system into different HP UX operating system images. For more details on vPars, please refer to Chapter 6 of this guide.

vPars is available on HP UX 11i and therefore can run on Superdome PA 8600/PA 8700 servers. Support of vPars on Superdome PA 8800 and PA-8900 is also supported.

Mixing of PA-RISC and Itanium Cells in Superdome

1) Which processors can be mixed in a Superdome?

The first step in determining which processors can be mixed within a Superdome is to look at the chipset. A Superdome can only support one type of chipset (legacy chipset or sx1000 chipset) at a time.

The legacy chipset only supports the PA-RISC architecture. The PA-8600, PA-8700, and PA-8700+ processors were supported with this chipset. As a result, they can be mixed

Configuration

within the Superdome but they can not be mixed with processors supported by other chipsets (i.e. Itanium 2 9M with the sx1000 chipset).

With the sx1000 chipset, processors of like architectures (PA RISC and Itanium 2 architectures) that are supported by the sx1000 chipset can be mixed in separate hard partitions. For example, the HP 9000 Superdome supports mixing the PA 8800 and PA 8900 processors in separate hard partitions. The HP Integrity Superdome supports mixing the Itanium 2 6M, and Itanium 2 9M processors, as well as the mx2 dual core processor module in separate hard partitions. In addition, a subset of the PA RISC and Itanium 2 processors (PA 8800, PA 8900 and Itanium 2 9M processors) can be supported at the same time in different hard partitions within a Superdome.

The table below highlights which processors can co-exist on a Superdome in separate hard partitions.

	PA 8600	PA 8700	PA 8700 +	PA 8800	PA 8900	Itanium 2 6M	Itanium 2 9M	mx2 dual-core processor module	Montecito
PA 8600	NA	Yes	Yes	No	No	No	No	No	No
PA 8700	Yes	NA	Yes	No	No	No	No	No	No
PA 8700 +	Yes	Yes	NA	No	No	No	No	No	No
PA 8800	No	No	No	NA	Yes	No	Yes	No	No
PA 8900	No	No	No	Yes	NA	No	Yes	No	Yes
Itanium 2 6M	No	No	No	No	No	NA	Yes	Yes	No
Itanium 2 9M	No	No	No	Yes	Yes	Yes	NA	Yes	No
mx2 dual-core processor module	No	No	No	No	No	Yes	Yes	NA	No
Montecito	No	No	No	No	No	No	No	No	NA

2) In order to run an Itanium-based partition in an Integrity server, what changes are required?

In order to add a new partition with Itanium 2 9M processors on an HP 9000 Superdome, the following steps are required:

Step #1: Upgrade firmware on PA-RISC based partitions

Step #2: Create a new hard partition in the Superdome for Itanium 2-based cell (s)

Step #3: Plug in cell boards for Itanium 2-based cells

Step #4: Some I/O cards may need to be added for that specific hard partition (Windows does not support the identical set of I/O cards that HP-UX 11i supports)

Step #5: Load operating system for Itanium 2-based partition

Upgrading a PA-RISC partition to support Itanium 2 processors would require similar steps:

Step #1: Upgrade firmware on PA-RISC based partitions

Configuration

Step #2: Pull out existing PA-RISC cell boards

Step #3: Swap existing memory into cell boards for Intel Itanium 2 processor (protects investment in current memory)

Step #4: Plug in cell boards for Itanium 2-based cell boards

Step #5: Some I/O cards may need to be added for that specific hard partition (Windows does not support the identical set of I/O cards that HP-UX 11i supports)

Step #6: Load operating system for Itanium 2-based partition

The in-box addition of Itanium 2 processors can be done with no additional hardware, no new chassis and no change to backplane.

3) When will mixing of PA-RISC and Itanium 2 processors be available?

Support for mixing PA-RISC and Itanium 2 processors with the sx1000 chipset in separate hard partitions will be available between July and September of 2005.

4) Is mixing of PA-RISC and Itanium processors factory configurable?

NO. The plan is to enable PA-RISC systems that shipped from the factory to add Itanium partitions later on in the field. In addition, HP offers the ability to enable Itanium-based systems that shipped from the factory to add PA-RISC partitions later on in the field (this will happen less frequently)

The plan is NOT to allow PA-RISC or Intel Itanium-based partitions to be added to systems that shipped from the factory. However, depending on the geographic region, shipment of Superdome systems with a mix of PA-RISC and Intel Itanium-based partitions may be possible.

5) How long does it take to add an Itanium-based partition in an HP 9000 customers?

The effort required to add an Itanium-based partition to an HP 9000 Superdome is slightly more than what is required to add a new PA-RISC-based partition. In both cases you have to create a new hard partition in the Superdome, plug in cell boards for the new processors and load the operating system. There are no hardware changes required to support mixing of processor types. The only known difference at this time is that a customer would need to upgrade the firmware to support the Itanium architecture. The process is currently being tested by R&D. This document will be updated as we learn more through testing.

6) Will mixing impact the overall performance of a Superdome (i.e. would an Itanium-based partition have the same performance in a mixed Superdome as in a Superdome only populated with Itanium 2 processors)?

No. For example, suppose you have a hard partition with 16 Intel Itanium 2 9M processors. This partition would have the same performance in a system with a mix of PA-RISC and Itanium 2-based hard partitions as a server where all of the other partitions are Itanium 2-based. This is assuming that besides the type of processors, all of the other configuration options are the same (memory, cell board locations, number of partitions, etc).

7) Can I add PA-RISC processors to an Itanium-based Integrity Superdome?

Yes. HP offers the ability to enable Itanium-based systems that shipped from the factory to add PA-RISC partitions later on in the field. This serves as a safety net/insurance policy. If the customer determines after

Configuration

the fact that a specific application is not ready for the Itanium technology, they can simply add a PA-RISC hard partition to that Integrity Superdome and continue the execution of the application while it is certified.

8) Why can't PA-8700+ processors be mixed with Itanium 2 9M processors?

For years, HP has had a high level of investment protection with the HP 9000 Superdome. Customers have been provided with the maximum investment protection level by keeping their investment in PA8600 and PA8700 processors while adding the higher speed PA8700+. In addition, each processor type runs at its rated speed so there is no downgrading of the higher speed processors (some competitors are erroneously saying Superdome has to downgrade the higher speed processors). For example, the PA8700 runs at a full 750 MHz while the PA8600 runs at 550 MHz—there is no forcing of the PA8700 down to 550 MHz just to obtain compatibility. This is possible because the processor speeds are independent from the bus speeds.

The Intel Itanium 2 processor family is an advanced architecture featuring exceptional floating point and SSL performance. The Itanium 2 processors have the processing power that is consistent with a Superdome class high-end server. In addition, the Itanium processor accesses cache memory using a data block that is greater than twice the size of the PA8600, PA8700, and PA8700+. Thus, the Itanium processor family and subsequent PA-8800 and PA-8900 processors require a different chipset than the PA8600, PA8700, or PA8700+ processors. One of the key restrictions in mixing processors is that all cell boards in a given Superdome must have the same chipset. The PA-8700+ is supported by a legacy chipset that only supported PA-RISC processors. As a result, the PA-8700+ can only be mixed with PA-8700 and PA-8600 processors. In order to support Itanium 2 9M processors in their current HP 9000 Superdome, the customer would be required first to upgrade all of their PA-8700+ processors to PA-8800 or PA-8900 processors utilizing the sx1000 chipset. Once this is complete, the customer would then have the option of adding the Itanium 2 9M processor in a separate partition because it is also supported by the sx1000 chipset.

9) Why are Itanium 2 6M processors and mx2 dual-core processor modules not supported?

The original plan was to enable PA RISC systems that shipped from the factory to add Itanium based partitions later on. The plan was not to allow PA RISC based partitions to be added to systems that shipped from the factory as Itanium based servers with Itanium 2 6M processors and/or mx2 dual core processor modules. This is the only case where we encounter a problem with no support for Itanium 2 6M processors and mx2 dual core processor modules. If a customer has already made the decision to go with Integrity Superdomes, it is very unlikely that they will take a back step to PA RISC based partitions. Conversely, if a customer wants to upgrade a PA RISC based partition to be Itanium 2 based, they are much more likely to use the Itanium 2 9M processor than the Itanium 2 6M processors or mx2 dual core processor modules because of the added performance and larger cache sizes.

Finally, this feature requires a significant amount of resources to test the different type of configurations supported in a Superdome. In order to provide this functionality in a timely basis to customers, we had to narrow the scope. Thus, the decision was made to not support Itanium 2 6M processors and mx2 dual core processors in a mixed processor type of system.

If a customer with Itanium 2 6M processors or mx2 dual core processor modules wants to add PA RISC modules, they would need to first upgrade the partitions with Itanium 2 6M processors or mx2 dual core processor modules to Itanium 2 9M processors and then they could add a PA RISC based partition to the Superdome.

10) Which versions of the HP-UX 11i operating system will be supported?

Currently HP plans to support HP-UX 11.11 (HP-UX 11i v1) for the PA-RISC architecture as well as 11.23 (HP-UX 11i v2) for Itanium processors and

Configuration

PA-RISC processors in a mixed environment. Based on HP's testing, HP believes that using HP-UX 11.11 (HP-UX 11i v1) in a mixed configuration is going to require a PA-RISC Firmware update.

11) Does a customer have to power down when they add an Itanium-based partition to a PA-RISC-based Superdome?

It depends.

If the existing PA-RISC based partitions are running HP-UX 11i v2 and PDC 22.1 (released in December 2004) then an Itanium-based partition may be added while the PA-RISC partitions are active. If the customer is running WLM it must be patched for proper operation in a mixed environment. We are currently checking with the WLM team to see what impact this would have on the PA-RISC partitions.

HP's shipping firmware for PA-RISC processors does not fully support mixing if the PA-RISC partitions are running HP-UX 11i v1. In this case, downtime for all PA-RISC based partitions would be required to update to the new PDC. The partitions could all be updated independently if desired. HP will release this new PDC when testing is complete for mixing processor types. The above comment regarding WLM also applies.

12) What are the minimum firmware requirements for mixing various operating environments?

The matrix below describes the minimum firmware requirements:

Mixing Mad9M with PA-8800/PA-8900	11i v2 HWE0409 11i v2 HWE0505 (no vPars) Windows Server 2003 Linux	11i v2 HWE0505 (vPars)
11i v1	PDC 22.3 IPF 2.54 MFW 15.14	PDC 22.3 IPF 3.x MFW 15.20
11i v2	PDC 22.1 IPF 2.50 MFW 15.14	PDC 22.1 IPF 3.x MFW 15.20

Mixing Scenarios: operating systems, firmware and management tools

The following table summarizes the various operating environments supported on the PA-8800, PA-8900 and Intel Itanium 2 1.6 GHz processors:

Processor	Operating System
PA8800	HPUX 11.i v1 HWE 0312 (with patches) HPUX 11.i v1 HWE 0406 HPUX 11.i v1 HWE 0412 HPUX 11.i v2 HWE 0409 HPUX 11.i v2 HWE 0505
PA8900	HPUX 11.i v1 HWE 0412 HPUX 11.i v2 HWE 0409 HPUX 11.i v2 HWE 0505
Itanium 2 1.6 GHz	HPUX 11.i v2 HWE 0409 HPUX 11.i v2 HWE 0505 Linux RHEL3 U3 (and higher) Linux SLES 9 (and higher) Windows 2003 Server, Smart Setup 3.1 (and higher)

The following table lists the minimum firmware requirements for mixing:

Configuration

Processor	Supported	Not Supported
PA8800	PDC 22.1 (SMS rel_5.0) or above	PDC 20.8, 21.2
PA8900	PDC 22.1 (SMS rel_5.0) or above	n/a
Itanium 2 1.6 GHz A1	IPF 2.50 (SMS rel_5.0) or above	n/a
Itanium 2 1.6 GHz A2	IPF 2.52 (SMS rel_5.1) or above	n/a

NOTE: PDC 20.8 and PDC 21.2 do not contain the required enhancements for mixing. Customers who are running on either of these PDC versions will be required to upgrade to PDC 22.1 to enable mixing. PDC 22.1 also happens the minimum firmware required to support PA8900.

Any of the above operating environments may be used on a mixed Superdome. In the process of testing all of the above combinations some issues and restrictions were discovered with WLM, gWLM and the GUI version of partition manager on HPUX 11i v1. Firmware, operating system or application updates can resolve these issues. However, such updates may not always be possible or practical.

The following mixing scenarios serve to illustrate the issues that were found during testing and possible ways to work around or resolve them. For simplicity, the first three scenarios assume that all the PA partitions in the Superdome are running the same version of HPUX. The fourth scenario explains the more general case where the PA partitions are running different versions of HPUX.

Scenario 1: PA8800 Partitions running HPUX 11i v1 HWE 0312 or 0406 mixed with Itanium 2 1.6 GHz

Known issue with the GUI version of Parmgr: Both of these HWEs shipped with Parmgr V1. This version of Parmgr and the stack that sits under it are not capable of managing Integrity partitions. Parmgr generates error messages when trying to access the Integrity cells and display information about them.

Alternatives to running Parmgr GUI on the PA partitions:

Integrity partitions (and PA partitions) may be managed from the SMS

Integrity partitions (and PA partitions) may be managed by parcmds from any partition

Integrity partitions (and PA partitions) may be managed by the GUI version of Parmgr on any Integrity partition

The minimum firmware listed in the table above is sufficient for running this configuration.

Possible upgrades: If it is important to be able to use the Parmgr GUI on the PA HPUX 11i v1 you can upgrade HPUX to HWE 0412 and follow the process in Scenario 2. If a customer upgrades to any 11i v1 HWE beyond HWE 0412 (HWE 0509 and later) then the proper nPar provider is already included in the HWE and the firmware upgrades in Scenario 2 are sufficient to get full functionality.

Configuration

Scenario 2: PA-8800 partitions or PA-8900 partitions running 11i v1 HWE 0412 mixed with Itanium 2 1.6 GHz

Known issue with the GUI version of Parmgr: HPUX 11i v1 HWE 0412 introduced Parmgr V2. This version of Parmgr will exhibit the same issues as Scenario 1 with the minimum required firmware from the table. However, upgrades to firmware and the nPar provider will make this version of Parmgr fully functional and capable of managing Integrity partitions. See the upgrade section below for details.

Alternatives to running Parmgr GUI on the PA partitions: These are essentially the same as Scenario 1. Integrity partitions (and PA partitions) may be managed from the SMS
Integrity partitions (and PA partitions) may be managed by parcmds from any partition
Integrity partitions (and PA partitions) may be managed by the GUI version of Parmgr on any Integrity partition
The minimum firmware listed in the table above is sufficient for running this configuration.

Possible upgrades: If it is important to be able to use the Parmgr GUI on the PA partitions to manage the Integrity partitions the following updates must be made to the system:

PDC must be at revision 22.3 or higher in the 11.i v1 partition. This version of PDC contains a single fix on top of PDC 22.1 that compensates for the byte reversal in the cell info structure between PA and Integrity (this is due to the different "endianness" of the processors). PDC 22.3 will be released as part of SMS rel_6.0 in September 2005.

Integrity firmware must be upgraded to the version that supports vPars. This version of Integrity firmware contains a fix that populates I/O slot information in cell info structures used by the PA partition management stack. Currently shipping Integrity firmware does not populate this information and causes the PA management stack to report errors for Integrity cells. This version of Integrity firmware will be released as part of SMS rel_6.0 in September 2005.

The nPartition provider ("NPar" bundle) must be updated to version B.11.11.01.04 or above. This version will first be released on HPUX 11.i v1 HWE 0509. The version of nPar provider that shipped with HWE 0412 did not properly handle the single core Madison 9M processors.

Scenario 3: PA-8800 partition or PA-8900 partition running 11.i v2 (HWE 0409 or 0505) mixed with Itanium 2 1.6 GHz

This is the simplest scenario. The minimum firmware listed in the table is sufficient and all the partition management tools just work. The PA partition can manage other PA partitions and the Integrity partition and vice versa. Enjoy.

Scenario 4: PA-8800 or PA-8900 partitions running a mixture of the above scenarios mixed with Itanium 2 1.6 GHz

The previous scenarios apply to the individual partition pairs. For instance, say you have a Superdome with a PA8800 11.i v1 HWE0406 partition, a PA8900 11.i v1 HWE0412 partition a PA8800 11.i v2 partition and an Integrity partition. As described in Scenario 1 the PA8800 11.i v1 HWE0406 partition will not be able to recognize the Integrity cells. Manage the Integrity partition from the SMS, the PA8800 11.i v2 partition or using parcmds. As described in Scenario 2, the PA8900 11.i v1 HWE0412 partition cannot manage the Integrity partition with out the firmware and nPar provider upgrades outlined above. Use the SMS, the PA 11.i v2 partition or parcmds to manage the Integrity partition.

Configuration

WLM:

WLM (workload manager) currently uses the Unix command "uname -i" to ensure that all the partitions it is managing are on the same Superdome. WLM verifies that all the managed partitions "uname -i" values match. It turns out the Integrity partitions and PA partitions on the same Superdome do not return the same value for "uname -i". This prevents WLM from properly managing all the partitions on a single Superdome in a mixed environment with iCAP. The WLM team has provided patches WLM A.03.00 for both 11.i v1 (PHSS_33499 - s700_800 11.11 WLM A.03.00 Cumulative Patch) and 11.i v2 (PHSS_33477 - s700_800 11.23 WLM A.03.00 Cumulative Patch) to resolve this issue. Customers who wish to use WLM in a mixed environment with iCAP will need to install/upgrade to version A.03.00 and install the appropriate patch. Customers who are using either WLM A.02.x or A.03.x in a non-iCAP environment can use their current version of WLM with no upgrade/patch. Future versions of WLM (A.03.00.01 and later) will not require a patch.

gWLM:

gWLM 1.1.1 (global workload manager) was first released with HPUX 11.i v2 HWE0505. It has the same issue with "uname -i" as WLM. This only affects gWLM iCAP SRDs. VPAR, PSET and FSS SRDs work correctly in mixed environments. The gWLM team plans to remedy this issue in the first maintenance release of gWLM. Contact gwlmfeedback@rsn.hp.com for a workaround if you have a gWLM 1.1.1 customer who needs to deploy an iCAP SRD on a mixed complex before the first maintenance release is available.

High Availability

NOTE: Online addition/replacement for cell boards is not currently supported and will be available in a future HP UX release. Online addition/replacement of individual processors and memory DIMMs will never be supported.)

Superdome high availability offering is as follows:

- **Processor:** The features below nearly eliminate the down time associated with processor cache errors (which are the majority of processor errors).
Dynamic processor resilience w/ Instant Capacity enhancement.
 - Processor cache ECC protection and automatic de allocation
 - Processor bus parity protection
 - Redundant DC conversion
- **Memory:** The memory subsystem design is such that a single SDRAM chip does not contribute more than 1 bit to each ECC word. Therefore, the only way to get a multiple-bit memory error from SDRAMs is if more than one SDRAM failed at the same time (rare event). The system is also resilient to any cosmic ray or alpha particle strike because these failure modes can only affect multiple bits in a single SDRAM. If a location in memory is "bad", the physical page is de-allocated dynamically and is replaced with a new page without any OS or application interruption. In addition, a combination of hardware and software scrubbing is used for memory. The software scrubber reads/writes all memory locations periodically. However, it does not have access to "locked-down" pages. Therefore, a hardware memory scrubber is provided for full coverage. Finally data is protected by providing address/control parity protection.
 - Memory DRAM fault tolerance (i.e., recovery of a single SDRAM failure)
 - DIMM address/control parity protection
 - Dynamic memory resilience (i.e., page de-allocation of bad memory pages during operation)
 - Hardware and software memory scrubbing
 - Redundant DC conversion
 - Cell COD
- **I/O:** Partitions configured with dual path I/O can be configured to have no shared components between them, thus preventing I/O cards from creating faults on other I/O paths. I/O cards in

Configuration

hardware partitions (nPars) are fully isolated from I/O cards in other hard partitions. It is not possible for an I/O failure to propagate across hard partitions. It is possible to dynamically repair and add I/O cards to an existing running partition.

- Full single-wire error detection and correction on I/O links
- I/O cards fully isolated from each other
- Hardware for the prevention of silent corruption of data going to I/O
- On-line addition/replacement (OLAR) for individual I/O cards, some external peripherals, SUB/HUB
- Parity protected I/O paths
- Dual path I/O
- **Crossbar and Cabinet Infrastructure:**
 - Recovery of a single crossbar wire failure
 - Localization of crossbar failures to the partitions using the link
 - Automatic de-allocation of bad crossbar link upon boot
 - Redundant and hotswap DC converters for the crossbar backplane
 - ASIC full burn-in and "high quality" production process
 - Full "test to failure" and accelerated life testing on all critical assemblies
 - Strong emphasis on quality for multiple-nPartition single points of failure (SPOFs)
 - System resilience to Management Processor (MP)
 - Isolation of nPartition failure
 - Protection of nPartitions against spurious interrupts or memory corruption
 - Hot swap redundant fans (main and I/O) and power supplies (main and backplane power bricks)
 - Dual power source
 - Phone-Home capability
- **"HA Cluster-In-A-Box" Configuration:** The "HA Cluster-In-A-Box" allows for failover of users' applications between hardware partitions (nPars) on a single Superdome system. All providers of mission critical solutions agree that failover between clustered systems provides the safest availability-no single points of failures (SPOFs) and no ability to propagate failures between systems. However, HP supports the configuration of HA cluster software in a single system to allow the highest possible availability for those users that need the benefits of a non-clustered solution, such as scalability and manageability. Superdome with this configuration will provide the greatest single system availability configurable. Since no single-system solution in the industry provides protection against a SPOF, users that still need this kind of safety and HP's highest availability should use HA cluster software in a multiple system HA configuration. Multiple Serviceguard or Serviceguard Extension for RAC clusters can be configured within a single Superdome system (i.e., two 4-node clusters configured within a 32-processor Superdome system).

Configuration

Multi-system High Availability

Any Superdome partition with PA-RISC processors that is protected by Serviceguard or Serviceguard Extension for RAC can be configured in a cluster with:

- Another Superdome with PA-RISC processors
- One or more standalone non Superdome systems with PA-RISC processors
- Another partition within the same single cabinet Superdome (refer to "HA Cluster-in-a-Box" above for specific requirements)

Separate partitions within the same Superdome system can be configured as part of different Serviceguard clusters.

Please note that when you add nodes or initially create a cluster, all nodes must be at the same version of the operating system and Serviceguard. This means that you may have to load an operating system update for hardware enablement of the newer hardware, even on older systems. Please refer to the "Compatibility and Feature Matrix" at

http://docs.hp.com/hpux/onlinedocs/4076/SG%20GeRAC%20EMS%20Support%20Matrix_10_3_03.htm

Geographically Dispersed Cluster Configurations

The following Geographically Dispersed Cluster solutions fully support cluster configurations using Superdome systems. The existing configuration requirements for non-Superdome systems also apply to configurations that include Superdome systems. An additional recommendation, when possible, is to configure the nodes of cluster in each datacenter within multiple cabinets to allow for local failover in the case of a single cabinet failure. Local failover is always preferred over a remote failover to the other datacenter. The importance of this recommendation increases as the geographic distance between datacenters increases.

- Extended Campus Clusters (using Serviceguard with Mirrordisk/UX)
- MetroCluster with Continuous Access XP
- MetroCluster with EMC SRDF
- ContinentalClusters

From an HA perspective, it is always better to have the nodes of an HA cluster spread across as many system cabinets (Superdome and non Superdome systems) as possible. This approach maximizes redundancy to further reduce the chance of a failure causing down time.

Management Features

Supportability and management features on HP 9000 Superdome are covered in the next section.

Service Processor (MP) The service processor (MP) utility hardware is an independent support system for nPartition servers. It provides a way for you to connect to a server complex and perform administration or monitoring tasks for the server hardware and its nPartitions. The main features of the service processor include the Command menu, nPartition consoles, console logs, chassis code viewers, and partition Virtual Front Panels (live displays of nPartition and cell states).

Access to the MP is restricted by user accounts. Each user account is password protected and provides a specific level of access to the Superdome complex and service processor commands. Multiple users can independently interact with the service processor because each service processor login session is private. Up to 16 users can simultaneously log in to the service processor through its network (customer LAN) interface and they can independently manage nPartitions or view the server complex hardware states. Two additional service processor login sessions can be supported by the local and remote serial ports. These allow for serial port terminal access (through the local RS-232 port) and external modem access (through the remote RS-232 port).

In general, the service processor (MP) on Superdome servers is similar to the service processor on other HP servers, while providing enhanced features necessary for managing a multiple-nPartition server. For example, the service processor manages the complex profile, which defines nPartition configurations as well as complex-wide settings for the server. The service processor also controls power, reset, and TOC capabilities, displays and records system events (chassis codes), and can display detailed information about the various internal subsystems.

Functional capabilities: The primary features available through the service processor are:

- **The Service Processor Command Menu:** provides commands for system service, status, access configuration, and manufacturing tasks.
- **Partition Consoles:** Each nPartition in a server complex has its own console. Each nPartition's console provides access to Boot Console Handler (BCH) interface and the HP-UX console for the nPartition.
- **Console Logs:** Each nPartition has its own console log, which has a history of the nPartition console's output, including boot output, BCH activity, and any HP-UX console login activity.
- **Chassis Logs Viewers (Live and Recorded Chassis Codes):** Three types of chassis code log views are available: activity logs, error logs, and live chassis code logs.
- **Virtual Front Panels:** Each nPartition's Virtual Front Panel (VFP) displays real-time status of the nPartition boot status and activity, and details about all cells assigned to the nPartition. The VFP display automatically updates as cell and nPartition status changes.

Management Features

Support Management Station (SMS)

The Support Management Station (SMS) runs the Superdome scan tools that enhance the diagnosis and testability of Superdome. The SMS and associated tools also provide for faster and easier upgrades and hardware replacement.

The purpose of the SMS is to provide Customer Engineers with an industry-leading set of support tools, and thereby enable faster troubleshooting and more precise problem root-cause analysis. It also enables remote support by factory experts who consult with and back up the HP Customer Engineer. The SMS complements the proactive role of HP's Instant Support Enterprise Edition (ISEE) that is offered to Mission Critical customers by focusing on reactive diagnosis for both mission-critical and non-mission-critical Superdome customers.

The user of the SMS is the HP Customer Engineer and HP Factory Support Engineer. The Superdome customer benefits from their use of the SMS by receiving faster return to normal operation of their Superdome server and improved accuracy of fault diagnosis, resulting in fewer callbacks. HP can offer better service through reduced installation time.

Functional Capabilities:

The SMS basic functional capabilities are:

- Remote access via customer LAN
- Modem access (PA-8800 and PA-8900 SMS only)
- Ability to be disconnected from the Superdome platform(s) and not disrupt their operation.
- Ability to connect a new Superdome platform to the SMS and be recognized by scan software.
- Support for up to sixteen Superdome systems
- Ability to support multiple, heterogeneous Superdome platforms (scan software capability).
- System scan and diagnostics
- Utility firmware updates
- Enhanced IPMI logging capabilities (Windows-based ProLiant SMS only)

Console Access

The optimal configuration of console device(s) depends on a number of factors, including the customer's data center layout, console security needs, customer engineer access needs, and the degree with which an operator must interact with server or peripheral hardware and a partition (i.e. changing disks, tapes). This section provides a few guidelines. However the configuration that makes best sense should be designed as part of site preparation, after consulting with the customer's system administration staff and the field engineering staff.

Customer data centers exhibit a wide range of configurations in terms of the preferred physical location of the console device. (The term "console device" refers to the physical screen/keyboard/mouse that administrators and field engineers use to access and control the server.) The Superdome server enables many different configurations by its flexible configuration of access to the MP, and by its support for multiple geographically distributed console devices.

Three common data center styles are:

- The secure site where both the system and its console are physically secured in a small area.
- The "glass room" configuration where all the systems' consoles are clustered in a location physically near the machine room.
- The geographically dispersed site, where operators administer systems from consoles in remote offices.

These can each drive different solutions to the console access requirement.

The considerations listed below apply to the design of provision of console access to the server. These must be considered during site preparation.

Management Features

- The Superdome server can be operated from a VT100 or an hpterm compatible terminal emulator. However some programs (including some of those used by field engineers) have a friendlier user interface when operated from an hpterm.
- LAN console device users connect to the MP (and thence to the console) using terminal emulators that establish telnet connections to the MP. The console device(s) can be anywhere on the network connected to either port of the MP.
- Telnet data is sent between the client console device and the MP "in the clear", i.e. unencrypted. This may be a concern for some customers, and may dictate special LAN configurations.
- If an HP UX workstation is used as a console device, an hpterm window running telnet is the recommended way to connect to the MP. If a PC is used as a console device, Reflection¹ configured for hpterm emulation and telnet connection is the recommended way to connect to the MP.
- The MP currently supports a maximum of 16 telnet connected users at any one time.
- It is desirable, and sometimes essential for rapid time to repair to provide a reliable way to get console access that is physically close to the server, so that someone working on the server hardware can get immediate access to the results of their actions. There are a few options to achieve this:
 - Place a console device close to the server.
 - Ask the field engineer to carry in a laptop, or to walk to the operations center.
 - Use a system that is already in close proximity of the server such as the Instant Support Enterprise Edition (ISEE) or the System Management Station as a console device close to the system.
 - The system administrator is likely to want to run X applications or a browser using the same client that they access the MP and partition consoles with. This is because the partition configuration tool, parmgr, has a graphical interface. The system administrator's console device(s) should have X window or browser capability, and should be connected to the system LAN of one or more partitions.

Support

The following matrix describes the supported SMS and recommended console devices for all Superdomes.

SMS and Console Support Matrix

	SMS	Console
PA 8700 (pre March 1, 2004)	Legacy UNIX SMS 1	PC/workstation
PA 8700 (post March 1, 2004)	UNIX rx2600 bundle	TFT5600 + Ethernet switch
PA 8700 upgraded to Integrity or PA-8800/PA-8900	Legacy UNIX SMS with software upgrade2	PC/workstation
	UNIX rx2600 bundle3	TFT5600 + Ethernet switch
	Windows SMS/Console (ProLiant ML350)	
Integrity or PA-8800/PA-8900	Windows SMS/Console (ProLiant ML350)	
	UNIX SMS/Console (rx2600)	

¹ A legacy UNIX SMS could be an A400, A500, rp2430 or rp2470 bundle, depending on when it was ordered

² In order for a legacy SMS to be upgraded to support Integrity or PA8800/PA-8900, it must be running HP-UX 11.0 or later, as sx1000 scan tools are not supported on HP-UX 10.20.

³ rx2600 SMS bundles ordered and installed prior to October 2004 will require a software upgrade in order to support an sx1000-based Superdome. As of October 2004, all rx2600 SMS bundles support PA8800/PA-8900 and Integrity Superdomes without this upgrade.

Management Features

PA-8700

Hardware Requirement Customers ordering an SMS for the first time for a new PA 8700 Superdome should order the rx2600 SMS. The rx2600 SMS can also be used to manage PA8800/PA-8900 and Integrity Superdomes.

Customers using the earlier-released A180 SMS must replace it with the rx2600 if they expect to use it with new Superdome or Integrity servers. Customers may use an existing rp2470 or A500 SMS to manage any new PA 8700 Superdome.

One rx2600 SMS can support up to 16 Superdomes using a switch. Please note, however, that certain datacenters are so large that the networking structure will not permit the sharing of one SMS for the entire datacenter. The SMS is connected to each PA 8700 Superdome system on a private LAN. It is beneficial to have the SMS in close physical proximity to the Superdome(s) because the Customer Engineer (CE) requires SMS access to service the Superdome hardware. The physical connection from the Superdome is a private Ethernet connection and thus, the absolute maximum distance is determined by the Ethernet specification.

The UNIX rx2600 SMS bundle is comprised of:

- HP rx2600 1.0G 1.5MB processor server Solution
- Factory rack kit for rx2600
- 1GB DDR memory
- 36GB 15K HotPlug Ultra320 HDD
- HP-UX 11iv2 Foundation OE
- 1 x HP Tape Array 5300 with DVD-ROM and DAT 40
- HP ProCurve Switch 2124
- CAT 5e Cables

By default, the rx2600 SMS does not come with a display monitor or keyboard unless explicitly ordered to enable console access (the TFT5600 rackmounted display/mouse/keyboard is the recommended solution). See the ordering guide for details on the additional components that are required in order to use the rx2600 SMS as a console.

Software Requirements All SMS software is preloaded in the factory and delivered to the customer as a complete solution. The rx2600 SMS supports only HP UX 11i at this time. Current versions of the SMS software have not been qualified for 64 bit Windows. To ensure only optimal diagnostic solutions are used, an integrated Windows/Linux SMS/Console is not available for PA 8700 Superdomes. All SMS software is preloaded in the factory and delivered to the customer as a complete solution.

Management Features

SMS Connectivity

PA-8700 Superdome requires scan traffic to be isolated from console traffic, therefore two distinct networks are required for the SMS and/or console. The rx2600 SMS has two LAN connections on the integrated multifunction I/O that can support and connect to two LAN interfaces on the Superdome MP: the Private LAN and the Customer LAN. These two LAN connections allow SMS and console operations to be performed remotely.

The 10/100Base TX port on the rx2600 is required, and is connected to the Private LAN on the Superdome MP. This connection is solely used for the various diagnostics supported by the SMS. The 10/100/1000Base TX port on the rx2600 can optionally be connected to the Superdome MP's Customer LAN for console access to the MP (and the Superdome partitions) from the existing management network. More details on console use of the rx2600 is provided later in this chapter.

For use as an SMS only, the rx2600's 10/100Base-TX port is connected to the Private LAN port on the Superdome MP. This can be done with a direct-connect crossover cable, or by using an Ethernet switch. HP recommends the switched connect configuration for the rx2600 SMS in order for the SMS to be shared with other Superdomes, and remotely accessed.

The SMS can be accessed remotely from the Management LAN, or directly via RS 232 on an as needed basis. If the customer chooses to access the SMS from the Customer Management LAN, the SMS traffic is on a distinct and private network. Console traffic goes to Customer LAN port on MP. This diagram assumes that the Customer provides the console infrastructure. Note as well that the UNIX rx2600 SMS is supported for use on all current models of Superdome.

Providing the Ethernet switch is configured with the UNIX rx2600 SMS, additional Superdomes can be easily added into the existing infrastructure with minimal disruption and downtime. Each additional Superdome MP's Private LAN port should be connected to the Ethernet switch. Note this will share SMS scan functionality only, not console access.

Console Connectivity

PA8700 Superdomes require scan and console to be on separate networks.

Existing PA-8700 Superdome customers may have a Legacy UNIX SMS (e.g. A500 or rp2470) that required a separate console device. Typically a Unix workstation or PC was configured in these environments. These customers may continue to use their existing console device to access any new PA 8700 Superdomes.

A Superdome console must meet the following requirements:

- HP Workstation running HP UX 11.0 or 11i v1
- PC running Windows NT, XP, 2000
- C1099A terminal server (limited functionality due to text mode only operation)
- 10Base T Ethernet connection or an RS 232 port or both.
- X windows emulator for console access on the PC (i.e., Reflection for HP with NS/VT).

For new environments, the UNIX rx2600 SMS can also be used as a console for a PA8700 Superdome by creating a distinct network for the console traffic and including a display and keyboard. Separate networks can be created via two Ethernet hubs or one Ethernet switch. The Customer LAN port on the Superdome MP is then connected through the console hub/switch to the 10/100/1000Base TX port on the integrated I/O of the rx2600.

In order to use the UNIX rx2600 SMS as a console for PA8700 Superdome, the following components must be ordered:

- TFT5600 retractable keyboard/display/mouse
- CAT5e cables to create new console network

Management Features

Additional PA 8700 Superdomes must have the Private LAN port connected into the scan network and the Customer LAN port connected into the Console Access network. For graphical access to partitions, the nPar Core I/O cards can be connected into the Console Access switch as well.

PA-8800/PA-8900

Hardware Requirements The UNIX rx2600 can also be used as an SMS for PA8800 and PA-8900 Superdomes. The same hardware requirements as detailed in the previous section for the PA8700 SMS apply.

The Windows ProLiant ML350 SMS/Console solution is supported on sx1000-based (i.e. PA-8800/PA-8900 and Integrity) Superdomes only. It is not supported on PA-8700 Superdomes.

The Windows ProLiant SMS:

- Allows local access to SMS by CE.
- Provides integrated console access, providing hpterm emulation over telnet and web browser, connecting over LAN or serial to a Superdome system
- Provides remote access over a LAN or dialup connection:
- ftp server with capability to ftp the firmware files and logs
- dialup modem access support (i.e., PC Anywhere or VNC)
- Provides seamless integration with data center level management.
- Provides partition logon capability, providing hpterm emulation over telnet, X windows, and Windows Terminal Services capabilities.

Provides following diagnostics tools:

- Runs HP's proven highly effective JTAG scan diagnostic tools, which offer rapid fault resolution to the failing wire.
- Console log storage and viewing
- Event log storage and viewing
- Partition and memory adviser flash applications
- Supports updating platform and system firmware.
- Always on event and console logging for Superdome systems, which captures and stores very long event and console histories, and allows HP specialists to analyze the first occurrence of a problem.
- Allows more than one LAN connected response center engineer to look at SMS logs simultaneously.
- Can be disconnected from the Superdome systems and not disrupt their operation.
- Provides ability to connect a new Superdome system to the SMS and be recognized by scan software.
- Scans one Superdome system while other Superdome systems are connected (and not disrupt the operational systems).
- Supports multiple, heterogeneous Superdome platforms.

The Windows ProLiant SMS/Console is comprised of a ProLiant ML350 G3/G4 and a TFT5600 retractable display monitor/keyboard/mouse to enable console access. This solution also requires (and includes) a switch. This is because scan diagnostics will not work properly if more than one IP address exists on the ProLiant SMS/Console. An important difference between the UNIX rx2600 SMS and the Windows ProLiant SMS/Console is that the ProLiant SMS, by default, provides console functionality whereas the UNIX rx2600 SMS does not. The TFT5600 retractable display/keyboard which is an optional add-on for the UNIX-based rx2600 SMS, is included by default with the Windows ProLiant SMS.

Additionally, the Windows ProLiant SMS/Console includes an internal modem that is intended for connection to a phone line. This is for cases in which the Customer does not want the SMS to be on a public network, and HP Field Services needs to access the SMS (they would then access the SMS via the

Management Features

phone line and PCAnywhere.)

A customer may not substitute any PC running Windows Server 2000 SP4 for the ProLiant SMS/Console due to the specialized software applications that have been qualified on the hardware and OS. Utilizing any other device as the SMS will void the warranty on the Superdome system and degrade the ability to service the system.

The Windows ProLiant SMS bundle is comprised of:

- 1 x HP ProLiant server ML350 G3/G4
- 1 x Intel Xeon™ DP 3.06 GHz processor
- 2 x 256MB
- 36GB 10K U320 HDD
- 1 x internal DVD
- 1 x internal V.90 56K modem with phone cord
- Windows 2000 Server SP4
- 2 x 25' CAT5e cables
- 1 x 4' CAT5e cable
- Ethernet switch and jumper cord
- Retractable display/keyboard/mouse
- Third-party applications

Software Requirements

All SMS software is preloaded in the factory and delivered to the customer as a complete solution.

The UNIX rx2600 SMS supports only HP UX 11i at this time. Current versions of the SMS software have not been qualified for 64 bit Windows.

The Windows ProLiant SMS/Console will run Windows 2000 SP4 as the default operating system. The ProLiant SMS/Console will follow the Windows OS roadmap and support later versions of this operating system as needed. The version of the scan tools used on the sx1000 Superdomes also does not require scan traffic to be isolated from console traffic.

SMS/Console Connectivity

One SMS can support up to 16 Superdomes using a switch (the Windows ProLiant SMS can support Integrity Superdomes only, and the UNIX rx2600 SMS can support both Integrity and HP 9000 Superdomes). Please note, however, that certain datacenters are so large that the networking structure will not permit the sharing of one SMS for the entire datacenter. It is beneficial to have the SMS in close physical proximity to the Superdome(s) because the Customer Engineer (CE) requires SMS access to service the Superdome hardware. The physical connection from the Superdome is a private Ethernet connection and thus, the absolute maximum distance is determined by the Ethernet specification.

Console capabilities are integrated with the Windows ProLiant SMS solution. The ProLiant SMS includes the necessary display, keyboard and mouse, and only one LAN port on the MP is required to be connected. The Windows ProLiant SMS can support and connect to either of the MP LAN interfaces (known as the Private and Customer LAN ports). Both LAN ports on the MP have identical functionality so there is no preference in using one over the other. Only one IP port on the Superdome MP is required to be connected to the ProLiant SMS. Since scan and console traffic can co-exist on the same network, only one IP address exists (and is supported) on the ProLiant SMS.

The PA 8800 and PA-8900 sx1000-based Superdome scan tools use TCP/IP, not UDP, for scan diagnostics. Therefore it is not necessary to isolate scan and console traffic on PA 8800 and PA-8900 Superdomes.

The Core I/O cards from each nPar can optionally be connected to the Ethernet switch in order to

Management Features

facilitate graphical console functionality (i.e., parmgr). However, security concerns may dictate that a partition NIC not be connected to the management LAN. Alternatives are to access from a management station to a partition LAN through a secure router, or to use text mode access to commands via the console.

Additional PA 8800/PA-8900 Superdomes can be added to this configuration. Each new PA 8800/PA-8900 Superdome will require only one CAT5e LAN cable for connection of the Customer/Private LAN port on the Superdome MP to the existing switch.

Customers using the UNIX rx2600 SMS solution for PA 8800/PA-8900 Superdome can also use the rx2600 as a console device. In order to use the UNIX based rx2600 SMS as a console for PA 8800/PA-8900 Superdome, the only component that must be ordered is the TFT5600 retractable keyboard/display/mouse. See the ordering guide for more details.

Additional PA 8800/PA-8900 Superdomes require only a CAT5e LAN cable for connection of the Private/Customer LAN port to the existing switch. PA 8700 Superdomes can also be added to this configuration, but require separate scan and console networks. Details on mixed Superdome environments follow later in this section.

Mixed Superdome Environments

Some important rules regarding SMS/Console support in mixed Superdome environments:

1. PA-8700 Superdomes require scan and console traffic to be isolated on separate networks
2. PA-8800/PA-8900 Superdomes do not require scan and console traffic to be isolated on separate networks
3. The Windows ProLiant SMS cannot be used to manage PA-8700 Superdomes
4. The UNIX rx2600 SMS can be used to manage all current models of Superdome
5. A legacy UNIX SMS can manage PA-8800/PA-8900 Superdomes as long as it undergoes a software upgrade and is running HP-UX 11.0

Upgrading a PA-8700 Superdome to PA-8800

Once a PA-8700 Superdome has been upgraded to PA-8800/PA-8900, a software upgrade must be performed on the SMS in order to have the correct scan tools for the sx1000 chipset. Alternatively, a new Windows ProLiant SMS or a UNIX rx2600 SMS can be ordered.

After performing the software upgrade, scan and console traffic from the SMS to the newly upgraded Superdome can co-exist on the same network. Note that if there are other PA-8700 Superdomes still connected to the SMS, those Superdomes will still require the two separate networks (see Adding a PA-8800/PA-8900 Superdome to a PA-8700 Environment, below).

Adding a PA-8800/PA-8900 Superdome to a PA-8700 Environment

When adding a PA 8800/PA-8900 Superdome to a PA 8700 environment, the Customer can choose to purchase a new SMS/Console (either the Windows or UNIX version) to manage the new PA 8800/PA-8900 Superdome, and continue to use their existing legacy SMS to manage the PA 8700 Superdomes.

A simpler solution is to upgrade the software on the legacy UNIX SMS so that it may manage the new PA 8800/PA-8900 Superdome. The upgraded UNIX SMS is able to determine which Superdomes are PA 8700 and which are PA 8800 or PA-8900, and will separate scan traffic from console traffic on the PA 8700 Superdomes.

To support this, the separate scan and console networks should be maintained for all PA 8700 Superdomes. The new PA 8800/PA-8900 Superdome requires only one connection to the SMS from its Private/Customer LAN port on the MP.

System Management Features

HP-UX

HP-UX Systems Insight Manager is an easy-to-use multi-system management solution with web-enabled and command line interfaces. HP Systems Insight Manager delivers multi system access to all key system administration tools for fault monitoring, configuration, and workload management. HP Systems Insight Manager will replace HP Servicecontrol Manager. It is available for download from the web now and will be included in the system soon. **Service Insight Manager** integrates with many other HP-UX-specific system management tools, including the following:

- **Ignite-UX**-Ignite-UX addresses the need for HP-UX system administrators to perform fast deployment for one or many servers. It provides the means for creating and reusing standard system configurations, enables replication of systems, permits post installation customizations, and is capable of both interactive and unattended operating modes.
- **Software Distributor (SD)** is the HP-UX administration tool set used to deliver and maintain HP-UX operating systems and layered software applications. Delivered as part of HP-UX, SD can help you manage your HP-UX operating system, patches, and application software.
- **System Administration Manager (SAM)** is used to manage accounts for users and groups, perform auditing and security, and handle disk and file system management and peripheral device management. Service Insight Manager enables these tasks to be distributed to multiple systems and delegated using role based security.
- **HP-UX Kernel Configuration**-for self-optimizing kernel changes. The new HP-UX Kernel Configuration tool allows users to tune both dynamic and static kernel parameters quickly and easily from a Web-based GUI to optimize system performance. This tool also sets kernel parameter alarms that notify you when system usage levels exceed thresholds.
- **Partition Manager** creates and manages nPartitions-hard partitions for high-end servers. Once the partitions are created, the systems running on those partitions can be managed consistently with all the other tools integrated into Service Insight Manager. Key features include:
 - Easy to use, familiar graphical user interface.
 - Runs locally on a partition, or remotely. The Partition Manager application can be run remotely on any system running HP-UX 11i remotely manage a complex either by 1) communicating with a booted OS on an nPartition in the target complex via WBEM, or 2) communicating with the service processor in the target complex via IPMI over LAN. The latter is especially significant because a complex can be managed with NONE of the nPartitions booted.
 - Full support for creating, modifying, and deleting hardware partitions.
 - Automatic detection of configuration and hardware problems.
 - Ability to view and print hardware inventory and status.
 - Big picture views that allow system administrators to graphically view the resources in a server and the partitions that the resources are assigned to.
 - Complete interface for the addition and replacement of PCI devices.
 - Comprehensive online help system.
- **Security Patch Check** determines how current a systems security patches are, recommends patches for continuing security vulnerabilities and warns administrators about recalled patches still present on the system.
- **System Inventory Manager** is for change and asset management. It allows you to easily collect, store and manage inventory and configuration information for HP-UX based servers. It provides an easy-to-use, Web based interface, superior performance, and comprehensive reporting capabilities
- **Event Monitoring Service (EMS)** keeps the administrator of multiple systems aware of system operation throughout the cluster, and notifies the administrator of potential hardware or software problems before they occur. HP Service Insight Manager can launch the EMS interface and configure EMS monitors for any node or node group that belongs to the cluster, resulting in increased reliability and reduced downtime.
- **Process Resource Manager (PRM)** controls the resources that processes use during peak system load. PRM can manage the allocation of processor, memory resources, and disk bandwidth. It

System Management Features

allows administrators to run multiple mission critical applications on a single system, improve response time for critical users and applications, allocate resources on shared servers based on departmental budget contributions, provide applications with total resource isolation, and dynamically change configuration at any time—even under load. (fee based)

- **HP-UX Workload Manager (WLM)** A key differentiator in the HP-UX family of management tools, Workload Manager provides automatic processor resource allocation and application performance management based on prioritized service level objectives (SLOs). In addition, WLM allows administrators to set real memory and disk bandwidth entitlements (guaranteed minimums) to fixed levels in the configuration. The use of workload groups and SLOs improves response time for critical users, allows system consolidation, and helps manage user expectations for performance. (Fee based)
- **HP's Management Processor** enables remote server management over the Web regardless of the system state. In the unlikely event that none of the nPartitions are booted, the Management Processor can be accessed to power cycle the server, view event logs and status logs, enable console redirection, and more. The Management Processor is embedded into the server and does not take a PCI slot. And, because secure access to the Management Processor is available through SSL encryption, customers can be confident that its powerful capabilities will be available only to authorized administrators.
- **OpenView Operations Agent**-collects and correlates OS and application events (fee-based)
- **OpenView Performance Agent**-determines OS and application performance trends (fee-based)
- **OpenView GlancePlus**-shows real time OS and application availability and performance data to diagnose problems (fee-based)
- **OpenView Data Protector (Omniback II)**-backs up and recovers data (fee-based)

In addition, the Network Node Manager (NNM) management station automatically discovers, draws (maps), and monitors networks and the systems connected to them.

All other OpenView management tools, such as OpenView Operations, Service Desk, and Service Reporter, will be able to collect and process information from the agents running on the server.

General Site Preparation Rules

AC Power Requirements The modular, N+1 power shelf assembly is called the Front End Power Subsystem (FEPS). The redundancy of the FEPS is achieved with 6 internal Bulk Power Supplies (BPS), any five of which can support the load and performance requirements.

System Management Features

Input Power Options

PDCA Product Number	Source Type	Source Voltage (nominal)	PDCA Required	Input Current Per Phase (200-240 VAC)	Power Required
A5800A Option 006	3-phase	Voltage range 200-240 VAC, phase-to-phase, 50/60 Hz	4-wire	44 A Maximum per phase	2.5 meter UL power cord and OL approved plug provided. The customer must provide the mating in line connector or purchase quantity one A6440A opt 401 to receive a mating in line connector. An electrician must hardwire the in line connector to 60 A/63 A site power. ^{a,b,c}
A5800A Option 007	3-phase	Voltage range 200-240 VAC, phase-to-neutral, 50/60 Hz	5-wire	24 A Maximum per phase	2.5 meter <HAR> power cord and VDE approved plug provided. The customer must provide the mating in line connector or purchase quantity 1 A6440A opt 501 to receive a mating in line connector. An electrician must hardwire the in line connector to 30 A/32 A site power. ^{a,b,d}

a. A dedicated branch is required for each PDCA installed.

b. In the U.S.A, site power is 60 Amps; in Europe site power is 63 Amps.

c. Refer to **Option 006 and 007 Specifics Table** below for detailed specifics related to this option.

d. In the U.S.A. site power is 30 Amps; in Europe site power is 32 Amps.

Console Device

PA-800 Superdome integrates console functionality into the Support Management Station. Note that this SMS/Console is also used on Integrity Superdome.

Console Considerations

The optimal configuration of console device(s) depends on a number of factors, including the customer's data center layout, console security needs, customer engineer access needs, and the degree with which an operator must interact with server or peripheral hardware and a partition (i.e., changing disks, tapes). This section provides a few guidelines. However the configuration that makes best sense should be designed as part of site preparation, after consulting with the customer's system administration staff and the field engineering staff.

Customer data centers exhibit a wide range of configurations in terms of the preferred physical location of the console device. (The term "console device" refers to the physical screen/keyboard/mouse that administrators and field engineers use to access and control the server.) The Superdome server enables many different configurations by its flexible configuration of access to the MP, and by its support for multiple geographically distributed console devices.

Three common data center styles are:

- The secure site where both the system and its console are physically secured in a small area.
- The "glass room" configuration where all the systems' consoles are clustered in a location physically near the machine room.
- The geographically dispersed site, where operators administer systems from consoles in remote offices.

These can each drive different solutions to the console access requirement. The considerations listed below apply to the design of provision of console access to the server. These must be considered during

System Management Features

site preparation.

- The Superdome server can be operated from a VT100 or an hpterm compatible terminal emulator. However some programs (including some of those used by field engineers) have a more friendly user interface when operated from an hpterm.
- LAN console device users connect to the MP (and thence to the console) using terminal emulators that establish telnet connections to the MP. The console device(s) can be anywhere on the network connected to either port of the MP.
- Telnet data is sent between the client console device and the MP "in the clear", i.e. unencrypted. This may be a concern for some customers, and may dictate special LAN configurations.
- If an HP-UX workstation is used as a console device, an hpterm window running telnet is the recommended way to connect to the MP. If a PC is used as a console device, Reflection1 configured for hpterm emulation and telnet connection is the recommended way to connect to the MP.
- The MP currently supports a maximum of 16 telnet-connected users at any one time.
- It is desirable, and sometimes essential, for rapid time-to-repair to provide a reliable way to get console access that is physically close to the server, so that someone working on the server hardware can get immediate access to the results of their actions. There are a few options to achieve this:
 - Place a console device close to the server.
 - Ask the field engineer to carry in a laptop, or to walk to the operations center.
 - Use a system that is already in close proximity of the server such as the Instant Support Enterprise Edition (ISEE) or the System Management Station as a console device close to the system.
- The system administrator is likely to want to run X-applications or a browser using the same client that they access the MP and partition consoles with. This is because the partition configuration tool, *parmgr*, has a graphical interface. The system administrator's console device(s) should have X-window or browser capability, and should be connected to the system LAN of one or more partitions.

Functional Capabilities

- Local console physical connection (RS-232)
- Display of system status on the console (front panel display messages)
- Console mirroring between LAN and RS-232 ports
- System hard and soft (TOC or INIT) reset capability from the console
- Password secured access to the console functionality
- Support of generic terminals (i.e. VT100 compatible).
- Power supply control and monitoring from the console. It will be possible to get power supply status and to switch power on/off from the console.
- Console over the LAN. This means that a PC or HP workstation can become the system console if properly connected on the customer LAN. This feature becomes especially important because of the remote power management capability. The LAN will be implemented on a separate port, distinct from the system LAN, and provide TCP/IP and Telnet access.

System Management Features

Pre-existing SMS Customers

Some customers may have an existing rp2470 or A500 SMS that required a separate console device. Typically a Unix workstation or PC was configured in these environments. These customers may continue to use their existing SMS (with the appropriate software upgrade) along with their existing console device to access the new PA-8800 Superdome. A Superdome console must meet the following requirements:

- HP Workstation running HP-UX 11.0 or 11i v1
- PC running Windows NT, XP, 2000
- C1099A terminal server (limited functionality due to text mode only operation)
- 10Base-T Ethernet connection or an RS-232 port or both.
- X windows emulator for console access on the PC (i.e., Reflection for HP with NS/VT).

Option 006 and 007 Specifics^a

PDCA Product Number	Attached Power Cord	Attached Plug	Customer-Provided Part	
			In-Line Connector	Panel-Mount Receptacle
A5800A Option 006	OLFLEX 190 (PN 600804), four conductor, 6-AWG (16-mm ²), 600-Volt, 60-Amp, 90-degree C, UL, and CSA approved, conforms to CE directives GN/YW ground wire.	Mennekes ME 460P9 3-phase, 4-wire, 60-Amp, 250-Volt, UL approved. Color blue, IEC 309-1, IEC 309-1, grounded at 3:00	Mennekes ME 460C9 3-phase, 4-wire, 60-amp, 250-Volt, UL approved. Color blue, IEC 309-1, IEC 309-1, grounded at 9:00. ^a	Mennekes ME 460R9 3-phase, 4-wire, 60-amp, 250-Volt, UL approved. Color blue, IEC 309-1, IEC 309-1, grounded at 9:00. ^b
A5800A Option 007	Five conductors, 10-AWG (6-mm ²), 450/475-Volt, 32-Amp, <HAR> European wire cordage, GN/YW ground wire.	Mennekes ME 532P6-14 3-phase, 5-wire, 32-Amp, 450/475-volt, VDE certified, color red, IEC 309-1, IEC 309-2, grounded at 6:00.	Mennekes ME 532C6-16 3-phase, 5-wire, 32-Amp, 450/475-Volt, VDE certified, color red, IEC 309-1, IEC 309-2, grounded at 6:00. ^c	Mennekes ME532R6-1276 3-phase, 5-wire, 32-Amp, 450/475-Volt, VDE certified, color red, IEC 309-1, IEC 309-2, grounded at 6:00. ^b

a. In-line connector is available from HP by purchasing A6440A, Option 401.

b. Panel-mount receptacles must be purchased by the customer from a local Mennekes supplier.

c. In-line connector is available from HP by purchasing A6440A, Option 501.

NOTE: A qualified electrician must wire the PDCA in-line connector to site power using copper wire and in compliance with all local codes.

System Management Features

Input Requirements

Reference the Site Preparation Guide for detailed power configuration requirements.

Requirements	Value	Conditions/Comments
Nominal Input Voltage (VAC rms)	200/208/220/230/240	
Input Voltage Range (VAC rms)	200-240	Auto selecting. Measure at input terminals.
Frequency Range (Hz)	50/60	
Number of Phases	3	3-phase 5-wire with power cord; 3-phase 4-wire with power cord
Maximum Input Current (A rms), 3-Phase 5-wire	20	3-phase source with a source voltage of 220 VAC measured phase to neutral
Maximum Input Current (A rms), 3-Phase 4-wire	40	3-phase source with a source voltage of either 208 VAC or 230 VAC measured phase-to-phase
Maximum Inrush Current (A peak)	90	
Circuit Breaker Rating (A), 3-Phase 5-wire	25 A	Per phase
Circuit Breaker Rating (A), 3-Phase 4-wire	45 A	Per phase
Power Factor Correction	0.95 minimum	
Ground Leakage Current (mA)	>3.5 mA, with 6 BPSs installed	Warning label applied to the PDCA at the AC Mains input

Cooling Requirements

- The cooling system in Superdome was designed to maintain reliable operation of the system in the specified environment as shown in the **Superdome Specifications Table**. In addition, the system is designed to provide a redundant cooling (i.e. N+1 fans and blowers) that allows all of the cooling components to be "hot swapped." The typical power dissipation for the PA 8600 is 28,850 BTUs/hour for a fully populated 32-processor cabinet. For PA 8700, the typical power dissipation is approximately 23,660 BTUs/hour for a fully populated 32-processor cabinet.
NOTE: For other configurations see "Site Prep Guide 6th Edition", Table 5. For maximum power dissipation, see **Superdome Specifications Table** under Site Preparation.
- Superdome was designed to operate in all data center environments with any traditional room cooling scheme (i.e. raised floor environments) but in some cases where data centers have previously installed high power density systems, alternative cooling solutions may need to be explored by the customer. Since no such system has been available previously, HP teamed with Liebert to develop an innovative data room cooling solution called DataCool. DataCool is a patented overhead climate system utilizing fluid based cooling coils and localized blowers capable of cooling heat loads of several hundred watts per square foot. Some of DataCool's highlights are listed below:
 - Liebert has filed for several patents on DataCool
 - DataCool, based on Liebert's TeleCool, is an innovative approach to data room cooling
 - Liquid cooling heat exchangers provide distributed cooling at the point of use
 - Delivers even cooling throughout the data center preventing hot spots
 - Capable of high heat removal rates (500 W per square foot)
 - Floor space occupied by traditional cooling systems becomes available for revenue generating equipment.
 - Enables cooling upgrades when installed in data rooms equipped with raised floor cooling

DataCool is a custom-engineered overhead solution for both new data center construction and for data room upgrades for high heat loads. It is based on Liebert's TeleCool product, which has been installed in 600 telecommunications equipment rooms throughout the world. The system utilizes heat exchanger pump units to distribute fluid in a closed system through patented cooling coils throughout the data

System Management Features

center. The overhead cooling coils are highly efficient heat exchangers with blowers that direct the cooling where required. The blowers are adjustable to allow flexibility for changing equipment placement or room configurations. Equipment is protected from possible leaks in the cooling coils by the patented monitoring system and purge function that detects any leak and safely purges all fluid from the affected coils. DataCool has interleaved cooling coils to enable the system to withstand a single point of failure and maintain cooling capability.

Features and Benefits

- Fully distributed cooling with localized distribution
- Even cooling over long distances
- High heat load cooling capacity (up to 500 W per square foot)
- Meets demand for narrow operating temperature for computing systems
- Allows computer equipment upgrade for existing floor cooled data rooms
- Floor space savings from removal of centralized air distribution
- Withstand single point of failures

For More Information

http://www.liebert.com/assets/products/english/products/env/datacool/60hz/bro_8pg/acrobat/sl_16700.pdf

HP has entered into an agreement with Liebert to reference sell the DataCool solution.

- The HP/Liebert business relationship will be managed by HP Complementary Products Division.
- DataCool will be reference by HP. Installation, service and support will be performed by Liebert.
- HP will compensate the HP Sales Representative and District Manager for each DataCool that Liebert sells to a customer referred by HP.
- An HP/Liebert DataCool website will be setup to get more information on the product and to manage the reference sales process. Please go to <http://hpcp.grenoble.hp.com/> for more information.

Environmental

- 20-30 degrees C inlet ambient temperature
- 0-10,000 feet
- 2600 CFM with N+1 blowers. 2250 CFM with N.
- 65 dBA noise level

Uninterruptible Power Supplies (UPS)

HP will be reselling high-end (10 kW and above) three-phase UPS systems from our partners. We will test and qualify a three-phase UPS for Superdome. The UPS is planned to be available Q1 FY01.

- All third-party UPS resold by HP will be tested and qualified by HP to ensure interoperability with our systems
- We plan to include *ups_mond* UPS communications capability in the third-party UPS(s), thus ensuring consistent communications strategy with our PowerTrust UPS(s)
- We will also establish a support strategy with our third party UPS partners to ensure the appropriate level of support our customer have come to expect from HP.
- For more information on the product and to manage the reference sales process. Please go to <http://hpcp.grenoble.hp.com/> for more information.

System Management Features

APC Uninterruptible Power Supplies for Superdome

The Superdome team has qualified the APC Silcon 3-phase 20-kW UPS for Superdome.

There are several configurations that can be utilized depending on the Superdome configuration your customer is deploying. They range from a 64-processor Superdome with dual cord and dual UPS with main tie main to a 32-processor Superdome with single cord and single UPS. In all configurations the APC Silcon SL20KFB2 has been tested and qualified by the Superdome engineers to ensure interoperability.

HP UPS Solutions

Product Number/Description	Quantity/Configuration	Watt	VA	Technology	Family	Package	Output
SL20KFB2 APC Silcon 3-phase UPS	<ul style="list-style-type: none"> Quantity 2/ 8- or 16-cell dual cord/dual UPS with main tie main Quantity 1/ 32- or 64-processor single-cord/single-UPS 	20 kW	20 kVA	Delta conversion on- line double conversion	APC Silcon 3-phase	Standalone rack	Configurable for 200: 208 or 220V 3-phase nominal output voltage
QJB22830 Switch Gear	<ul style="list-style-type: none"> Quantity 1/ 8- or 16-cell dual cord/dual UPS with main- tie-main Quantity 0/ 32- or 64-processor single-cord/single-UPS 	N/A	N/A	N/A	Customer Design for Superdome	N/A	N/A
WSTRUP5X8- SL10 Start-Up Service	<ul style="list-style-type: none"> Quantity 2/ 8- or 16-cell dual cord/dual UPS with main- tie-main Quantity 1/ 32- or 64-processor single-cord/single-UPS 	N/A	N/A	N/A	N/A	N/A	N/A
WONSITENBD- SL10 Next Business Day On site Service	<ul style="list-style-type: none"> Quantity 2/ 8- or 16-cell dual cord/dual UPS with main- tie-main Quantity 1/ 8- or 16-cell single- cord/single-UPS 	N/A	N/A	N/A	N/A	N/A	N/A

NOTE: The APC Silcon 3-phase UPS solutions for Superdome must be ordered directly from APC. Please contact Ron Seredia at rseredia@apcc.com.

Superdome Server Watt Ratings for UPS loading

Class	Models	Watt Rating for UPS loading	UPSs Typically Used
Superdome	8-cell	19 kW	SL20KFB2; 20 kW/20 kVA
Superdome	16-cell	19 kW each cabinet; 38 kW total	SL20KFB2; 20 kW/20 kVA; Quantity 2

System Management Features

Power Protection

Runtimes The UPS will provide battery backup to allow for a graceful shutdown in the event of a power failure. Typical runtime on the APC SL20KFB2 Silcon 3 Phase UPS varies with the kW rating and the load. The APC SL20KFB2 UPS provides a typical runtime of 36.7 minutes at half load and 10.7 at full load. If additional run time is needed please contact your APC representative.

Power Conditioning The APC SL20KFB2 provides unparalleled power conditioning with its Delta-Conversion on-line double conversion technology. This is especially helpful in regions where power is unstable.

Continuous Power during Short Interruptions of Input Power The APC SL20KFB2 will provide battery backup to allow for continuous power to the connected equipment in the event of a brief interruption in the input power to the UPS. Transaction activity will continue during brief power outage periods as long as qualified UPS units are used to provide backup power to the SPU, the Expansion Modules, and all disk and disk array products.

UPS Configuration Guidelines In general, the sum of the "Watt rating for UPS sizing" for all of the connected equipment should not exceed the watt rating of the UPS from which they all draw power. In previous configuration guides, this variable was called the "VA rating for UPS sizing." With Unity Power Factor, the Watt rating was the same as the kVA rating, so it didn't matter which one we used. VA is calculated by multiplying the voltage times the current. Watts, which is a measurement of true power, may be less than VA if the current and voltage are not in phase. APC SL20KFB2 has Unity Power Factor correction, so the kW rating equals the kVA rating. Be sure to add in the needs for the other peripherals and connected equipment. When sizing the UPS, allow for future growth as well. If the configuration guide or data sheet of the equipment you want to protect gives a VA rating, use this as the watt rating. If the UPS does not provide enough power for the additional devices such as system console and mass storage devices, additional UPSs may be required.

Superdome The only qualified UPS available for use with Superdome is the APC SL20KFB2 Silcon 3 Phase 20 kW UPS. The APC SL20KFB2 can provide power protection for the SPU and peripherals. If the system console and primary mass storage devices (such as HP High Availability Disk Array Model 20) also require power protection (which is highly recommended) they may require one or more additional UPSs depending on the total Watts. Make sure that the total watts do not exceed the UPS's voltage rating.

Integration/Installation The APC SL20KFB2 includes both field integration start up service and next day on site service for one year provided by APC.

Power Connections with the APC SL20KFB2

Product Number	Watts	NOM Out	Output Receptacles	Input Receptacles
SL20KFB2	20 kW	115/200 3PH, 120/208 3PH, 127/220 3PHV	Hardwire	Hardwire

Communications Connections A DB-25 RS-232 Contact Closure connection is standard on all APC SL20KFB2 UPS. A WEB/SNMP card is also included.

System Management Features

Power Management

Description	Network interface cards that provide standards based remote management of UPSs
General Features	Boot-P support, Built-in Web/SNMP management, Event logging, Flash Upgradeable, MD5 Authentication Security, Password Security, SNMP Management, Telnet Management, Web Management
Includes	CD with software, User Manual
Documentation	User Manual Installation Guide

Type of UPSs

Some customers may experience chronic "brown out" situations or have power sources that are consistently at the lower spectrum of the standard voltage range. For example, the AC power may come in consistently at 92 VAC in a 110 VAC area. Heavy-load electrical equipment or power rationing are some of the reasons these situations arise. The APC SL20KFB2 units are designed to kick in before the AC power drops below the operating range of the HP Superdome Enterprise Server. Therefore, these UPS units may run on battery frequently if the AC power source consistently dips below the threshold voltage. This may result in frequent system shutdowns and will eventually wear out the battery. Although the on-line units can compensate for the AC power shortfall, the battery life may be shortened. The best solution is to use a good quality boost transformer to "correct" the power source before it enters the UPS unit.

Ordering Guidelines

- The APC SL20KFB2 Silcon 3-phase UPS units may be ordered as part of a new Superdome system order or as a field upgrade to an existing system.
- For new systems order please contact APC at Ron Seredian rseredia@apcc.com during the Superdome pre-consulting phase. APC will coordinate with HP to ensure the UPS is installed to meet the Superdome installation schedule.
- For field upgrades please contact APC at Ron Seredian rseredia@apcc.com when you determine a customer is in need and/or interested in power protection for Superdome. APC will coordinate with the customer to ensure the UPS is installed to meet their requirements.
- Numerous options can be ordered to compliment APC SL20KFB2 Silcon 3-phase UPS units. Your APC consultant can review these option with you are you can visit the APC website at www.apcc.com

Multi-cabinet Configurations

In order to support the maximum number of PCI slots, a 16-cell Superdome system requires 16 I/O chassis. The two Superdome cabinets (left and right) that make up a 16-cell Superdome system only provide eight I/O chassis, therefore four I/O chassis enclosures, each with two I/O chassis are needed. The I/O chassis enclosures are placed in the I/O expansion cabinet. Each I/O expansion cabinet supports up to three I/O chassis, so two I/O expansion chassis are needed. Please note that in order to have 16 I/O chassis, there must also be 16 cells configured.

When configuring Superdome systems that consist of more than one cabinet and include I/O expansion cabinets, certain guidelines must be followed, specifically the I/O interface cabling between the Superdome cabinet and the I/O expansion cabinet can only cross one additional cabinet due to cable length restrictions.

Configuration Rules

System Management Features

SuperDome Configuration Guidelines/Rules		
Category	Rule Index	Rule Description
General	1	Every Superdome complex requires connectivity to a Support Management Station (SMS). For Superdome PA-8800, the PC-based SMS also serves as the system console. For Superdome PA-8600/PA-8700, a separate system console is needed.
	2	Every cell in a Superdome complex must be assigned to a valid physical location.
Processor	3	All processors in a cell are the same type, same Front Side Bus (FSB) frequency and same core frequency. Every cell within a partition must be of the same type, same Front Side Bus (FSB) frequency and same core frequency with no mixing. With PA-8800, the entire cabinet must consist of only PA-8800 cell boards.
Memory	4	Configurations with 8, 16 and 32 DIMM slots are recommended (i.e. are fully qualified and offer the best bandwidth performance).
	5	Configurations with 4 and 24 DIMM slots are supported (i.e. are fully qualified, but don't necessarily offer the best bandwidth performance).
	6	DIMMs can be deallocated in 2 DIMM increments (to support HA) in Superdome PA-8800 systems.). DIMMs are deallocated in 4 DIMM increments for Superdome PA-8600/8700 systems.
	7	Superdome PA-8800 only: Mixed DIMM sizes within a cell board are supported, but only in separate Mbat interleaving groups. At the time of boot, the system firmware gathers information on which DIMMs are present and how to best interleave the memory. Thus, DIMMs of different sizes will not be in the same Mbat interleaving group. From a customer perspective, it is it automatically done by the system firmware.
	8	Superdome PA-8800 only: System orders from the factory provide mixed DIMM sizes in recommended configurations only.
	9	For system orders from the factory, the same memory configuration must be used for all cells within a partition.
	10	DIMMs in the same rank must have SDRAMs with the same number of banks and row and column bits.
	11	Size of memory within an interleave group must be power of 2.
	12	DIMMs within the same interleave group must be same size and have same number of banks, row bits, and column bits.
	13	There are currently no restrictions on mixing DIMMs (of the same type) with different vendor SDRAMs.
I/O	14	One cell in every partition must be connected to an I/O chassis that contains a Core I/O card, a card connected to boot media, a card connected to removable media, and a network card with a connected network.
	15	A partition cannot have more I/O chassis than it has active cells.
	16	Removable media device controller should be in slot 8 of the I/O chassis.
	17	Core I/O card must be in slot 0 of the I/O chassis.
	18	Boot device controller should be in slot 1 of the I/O chassis
	19	PCI X high bandwidth I/O cards should be in the high bandwidth slots in the I/O chassis
	20	Every I/O card in an I/O chassis must be assigned to a valid physical location.
	21	Every I/O chassis in a Superdome complex must be assigned to a valid physical location

Category	Rule Index	Rule Description
Performance	22	NOTE: Only Superdome PA-8800 supports 1 GB and 2 GB DIMM sizes. The amount of memory on a cell should be evenly divisible by 4 GB if using 512-MB DIMMs or 8 GB if using 1-GB

System Management Features

	<p>DIMMs or 16 GB if using 2-GB DIMMs, i.e. 8, 16 or 32 DIMMs. The cell has four memory subsystems and each subsystem should have an echelon (2 DIMMs) populated. The loading order of the DIMMs alternates among the four subsystems. This rule provides maximum memory bandwidth on the cell, by equally populating all four memory subsystems.</p> <p>Superdome PA-8600/8700 supports 512MB DIMMs only. The amount of memory on one of these cells should be evenly divisible by 4GB, i.e. 8, 16 or 32 DIMMs. These cells have two memory subsystems and each subsystem should have a rank (4 DIMMs) populated. The loading order of the DIMMs alternates between the two subsystems. This rule provides maximum memory bandwidth on the cell by equally populating both memory subsystems.</p>
23	All cells in a partition should have the same number of processors.
24	The number of active processors per cell should be balanced across the partition, however minor differences are OK. (Example: 4 active processors on one cell and three active processors on the second cell)
25	If memory is going to be configured as fully interleaved, all cells in a partition should have the same amount of memory (symmetric memory loading). Asymmetrically distributed memory affects the interleaving of cache lines across the cells. Asymmetrically distributed memory can create memory regions that are non optimally interleaved. Applications whose memory pages land in memory interleaved across just one cell can see up to 16 times less bandwidth than ones whose pages are interleaved across all cells.
26	If a partition contains 4 or fewer cells, all the cells should be linked to the same crossbar (quad) in order to eliminate bottlenecks and the sharing of crossbar bandwidth with other partitions. In each Superdome cabinet, slots 0, 1, 2 and 3 link to the same crossbar and slots 4, 5, 6 and 7 link to the same crossbar.
27	A Core I/O card should not be selected as the main network interface to a partition. A Core I/O card is a PCI 1X card that possibly produces lower performance than a comparable PCI 2X 66-MHz or PCI X 2X 133-MHz PCI-X card.
28	The number of cells in a partition should be a power of two, i.e., 2, 4, 8, or 16. Optimal interleaving of memory across cells requires that the number of cells be a power of two. Building a partition that does not meet this requirement can create memory regions that are non optimally interleaved. Applications whose memory pages land in the memory that is interleaved across just one cell can experience up to 16 times less bandwidth than pages which are interleaved across all 16 cells.
29	<p>Before consolidating partitions in a Superdome 32-processor or 64-processor system, the following link load calculation should be performed for each link between crossbars in the proposed partition.</p> <p>Links loads less then 1 are best. As the link load begins to approach 2 performance bottlenecks may occur.</p> <p>For crossbars X and Y Link Load = $Q_x * Q_y / Q_t / L$, where</p> <ul style="list-style-type: none"> - Q_x is the number of cells connected to crossbar X (quad) - Q_y is the number of cells connected to crossbar Y (quad) - Q_t is the total number of cells in the partition - L is the number of links between crossbar X and Y (2 for Superdome 32-processor systems and 1 for Superdome 64-processor systems)
30	Maximum performance for optimal configurations is power of two cells, uniform memory across cells and power of two DIMM ranks per cell.
31	(If rule #30 cannot be met, rule #31 is recommended) Non power of two cells, but still uniform memory across cells, power of two DIMM ranks per cell, uniform type of DIMM.

System Management Features

32	NOTE: Only Superdome PA-8800 supports 1 GB and 2 GB DIMMs. (If rule #30 or #31 cannot be met, rule #32 is recommended) Same amount of memory in each cell, but possibly different memory types in each cell (for instance, a two cell configuration with 8 512MB DIMMs in one cell, and 4 1GB DIMMs in the other). Differences in memory across different cells within the same partition should be minimal for the best performance.
33	NOTE: Only Superdome PA-8800 supports 1 GB and 2 GB DIMMs. Same amount of memory in each cell, but non optimal and/or mixed loading within a cell (for instance, a two cell configuration with 16 512MB DIMMs and 8 1GB DIMMs in each cell).
34	Non-uniform amount of memory across cells (this needs to boot and run, but performance is whatever you get).
35	Superdome PA-8800 only: For the same amount of total memory, best performance is with a larger number of smaller size DIMMs.

Category	Rule Index	Rule Description
Single System High Availability	36	Superdome PA-8600/PA-8700: Each cell should have at least two active processors. Superdome PA-8800: Each cell should have at least four active processors.
	37	NOTE: Only Superdome PA 8800 supports 1-GB DIMMs. Each cell should have at least 4 GB (8 DIMMs) of memory using 512-MB DIMMs and at least 8 GB of memory using 1-GB DIMMs or at least 16 GB of memory using 2-GB DIMMs..
	38	I/O chassis ownership must be localized as much as possible. One way is to assign I/O chassis to partitions in sequential order starting from INSIDE the single cabinet, then out to the I/O expansion cabinet 'owned' by the single cabinet.
	39	I/O expansion cabinets can be used only when the main system cabinet holds maximum number of I/O card cages. Thus, the cabinet must first be filled with I/O card cages before using an I/O expansion cabinet.
	40	Single cabinets connected to form a dual cabinet (using flex cables) should use a single I/O expansion cabinet if possible.
	41	Spread enough connections across as many I/O chassis as it takes to become 'redundant' in I/O chassis'. In other words, if an I/O chassis fails, the remaining chassis have enough connections to keep the system up and running, or in the worst case, have the ability to reboot with the connections to peripherals and networking intact.
	42	All SCSI cards are configured in the factory as unterminated. Any auto termination is defeated. If auto termination is not defeatable by hardware, the card is not used at first release. Terminated cable would be used for connection to the first external device. In the factory and for shipment, no cables are connected to the SCSI cards. In place of the terminated cable, a terminator is placed on the cable port to provide termination until the cable is attached. This is needed to allow HP UX to boot. The customer does not need to order the terminators for these factory integrated SCSI cards, since the customer will probably discard them. The terminators are provided in the factory by use of constraint net logic.
	43	Partitions whose I/O chassis are contained within a single cabinet have higher availability than those partitions that have their I/O chassis spread across cabinets.
	44	A partition's core I/O chassis should go in a system cabinet, not an I/O expansion cabinet.
	45	A partition should be connected to at least two I/O chassis containing Core I/O cards. This implies that all partitions should be at least 2 cells in size. The lowest number cell or I/O chassis is the 'root' cell; the second lowest number cell or I/O chassis combo in the partition is the 'backup root' cell.
46	A partition should consist of at least two cells.	

System Management Features

	47	Not more than one partition should span a cabinet or a crossbar link. When crossbar links are shared, the partition is more at risk relative to a crossbar failure that may bring down all the cells connected to it.
--	----	---

Category	Rule Index	Rule Description
Multi-System High Availability (Please also refer to Multi-System High Availability section following this table)	48	<i>Multi-initiator support is required for Serviceguard.</i>
Traditional Multi-System High Availability	49	To configure a cluster with no SPOF, the membership must extend beyond a single cabinet. The cluster must be configured such that the failure of a single cabinet does not result in the failure of a majority of the nodes in the cluster. The cluster lock device must be powered independently of the cabinets containing the cluster nodes. Alternative cluster lock solution is the Quorum Service, which resides outside the Serviceguard cluster providing arbitration services.
	50	A cluster lock is required if the cluster is wholly contained within two single cabinets (i.e., two Superdome 16- processor or 32-processor systems or two Superdome PA 8800/PA 8900 32-processor or 64-processor systems) or two dual cabinets (i.e. two Superdome 64-processor core systems or two Superdome PA 8800/PA 8900 128 processor core systems). This requirement is due to a possible 50% cluster failure.
	51	Serviceguard only supports cluster lock up to four nodes. Thus a two cabinet configuration is limited to four nodes (i.e., two nodes in one dual cabinet Superdome 64-processor system or Superdome PA 8800/PA 8900 128-processor core system and two nodes in another dual cabinet Superdome 64-processor core system or Superdome PA 8800/PA 8900 128-processor core system). The Quorum Service can support up to 50 clusters or 100 nodes (can be arbitrator to both HP UX and Linux clusters).
	52	Two cabinet configurations must evenly divide nodes between the cabinets (i.e. 3 and 1 is not a legal 4-node configuration).
	53	Cluster lock must be powered independently of either cabinet.
	54	Root volume mirrors must be on separate power circuits.
	55	Redundant heartbeat paths are required and can be accomplished by using either multiple heartbeat subnets or via standby interface cards.
	56	Redundant heartbeat paths should be configured in separate I/O chassis when possible.
	57	Redundant paths to storage devices used by the cluster are required and can be accomplished using either disk mirroring or via LVM's pvlins.
	58	Redundant storage device paths should be configured in separate I/O chassis when possible.
59	Dual power connected to independent power circuits is recommended.	

System Management Features

Category	Rule Index	Rule Description
Heterogeneous Multi-System High Availability	60	Cluster configurations can contain a mixture of Superdome and non-Superdome nodes.
	61	Care must be taken to configure an even or greater number of nodes outside of the Superdome cabinet
	62	If half the nodes of the cluster are within a Superdome cabinet, a cluster lock is required (4-node maximum cluster size)
	63	If more than half the nodes of a cluster are outside the Superdome cabinet, no cluster lock is required (16-node maximum Serviceguard cluster size).
	64	Up to a 4 node cluster is supported within a single cabinet system (Superdome 16-processor or Superdome PA 8800/PA 8900 32-processor)
	65	Up to an 8 node cluster is supported within a single cabinet system* (Superdome 32-processor or Superdome PA 8800/PA 8900 64-processor)
	66	Up to a 16 node cluster is supported within a dual cabinet system* (Superdome 64-processor or Superdome PA 8800/PA 8900 128-processor core)
	67	Cluster lock is required for 2-node configurations
	68	Cluster lock must be powered independently of the cabinet.
	69	Root volume mirrors must be on separate power circuits.
	70	Dual power connected to independent power circuits is highly recommended.

Power Redundancy Superdome servers, by default, provide an additional power supply for N+1 protection. As a result, Superdome servers will continue to operate in the event of a single power supply failure. The failed power supply can be replaced without taking the system down.

Instant Capacity (iCAP) (formerly known as Instant Capacity on Demand [iCOD]) **and Pay Per Use Programs.** For a complete description of how to configure Instant Capacity and Pay Per Use, please refer to the following URL:

http://h18000.www1.hp.com/products/quickspecs/11723_div/11723_div.HTML

Upgrades

For information on Superdome Server Upgrades, please refer to the Superdome Server Upgrades QuickSpec.

Technical Specifications

Superdome Specifications	SPU Description	Superdome 16 processor	Superdome 32 processor	Superdome 64 processor
	SPU Product Number	A6113A	A5201A	A5201A + A5202A
	PA-8600 / PA-8700 Attributes			
	TPC-C disclosure (PA 8700)/ 875 MHz	Not disclosed	Not disclosed	541,674 tpm
	Clock speed (MHz)	552 MHz	552 MHz	552 MHz
	- PA 8600	750 MHz, 875 MHz	750 MHz, 875 MHz	750 MHz, 875 MHz
	- PA 8700			
	Number of processor modules	1-16	1-32	6-64
	Memory (with 512 MB DIMM's)	2-64 GB	2-128 GB	6-256 GB
	Maximum Power Dissipation (watts)	5,820	8,460	16,920
	PA-8800 and PA-8900 Attributes (with sx1000 chip set)			
	PA-8800 Clock Speed (MHz)	1.0 GHz	1.0 GHz	1.0 GHz
	PA-8800 L2 Cache Size	32 MB	32 MB	32 MB
	PA-8900 Clock Speed (MHz)	1.1 GHz	1.1 GHz	1.1GHz
	PA-8900 L2 Cache Size	64 MB	64 MB	64 MB
	Number of cores	1-32	1-64	6-128
	Memory (with 512 MB DIMMs)	2-64 GB	2-128 GB	6-256 GB
	Memory (with 1 GB DIMMs)	4-128 GB	4-256 GB	12-512 GB
	Memory (with 2 GB DIMMs)	8-256 GB	8-512 GB	24-1024 GB ¹
	Maximum Power Dissipation (watts)	5,054	9,038	18,076
	Cell Boards²			
	PA 8800 and PA 8900 are supported in 1-processor module, 2-processor module, 3-processor module, and 4-processor module cell boards. Other PA RISC processors are supported in 4 processor module cell boards.	1-4	1-8	3-16

Expandability / Connectivity

Technical Specifications

12-slot PCI I/O chassis without I/O expansion cabinet (PCI-X supported with PA-8800 and PA-8900 only)	1-4	1-4	3-16
PCI slots without I/O expansion cabinet (PCI-X supported with PA-8800 and PA-8900 only)	12-48	12-48	12-96
PCI slots with I/O expansion cabinet (PCI-X supported with PA-8800 and PA-8900 only)	N/A	24-96	12-192 ³
12-slot PCI I/O chassis with I/O expansion cabinet (PCI-X supported with PA-8800 and PA-8900 only) ⁴	N/A	8	16 ⁵
Number of Partitions without I/O expansion cabinet	1-4	1-4	1-8
Number of Partitions with I/O expansion cabinet	1-4	1-8	1-16
RS-232C Serial Ports	Yes	Yes	Yes
10/100Base-T Ethernet	Yes	Yes	Yes
HP-UX Attributes			
Maximum cores per HP UX 11i v1 partition	16	32	64
Maximum cores per HP UX 11i v2 partition	16	32	128
Maximum memory per HP UX 11i v1 partition	128 GB	256 GB	512 GB
Maximum memory per HP UX 11i v2 partition	128 GB	256 GB	1 TB
HP-UX revisions supported	HP-UX 11i, HP-UX 11i v2	HP-UX 11i, HP-UX 11i v2	HP-UX 11i, HP-UX 11i v2
Maximum I/O Cards See supported I/O table for specific products			
Mass Storage	16 - 48	32 - 96	64 - 192
LAN	8 - 24	16 - 48	32 - 96
WAN	8 - 25	16 - 25	25 - 32
Multi-Function (Mass Storage / LAN)	8 - 16	16 - 32	32 - 64
Additional Interface Cards	8	8 - 16	8 - 32

Electrical Characteristics

Technical Specifications

AC input power-Option 7: 3-phase 5-wire input	200-240 VAC phase-to-neutral, 5-wire, 50/60 Hz
AC input power-Option 6: 3-phase 4-wire input	200-240 VAC phase-to-phase, 4-wire, 50/60 Hz
Maximum Current requirements at 220V 240V:	
Option 7: 3-phase 5-wire input	24 A/Phase
Option 6: 3-phase 4-wire input	44 A/Phase
Required Power Receptacle-Options 6 and 7	None. Cord, plug and included. Receptacle should be ordered separately. Electrician must hard wire receptacle to site power.

Site Preparation

Site planning and installation included	Yes	Yes	Yes
Depth (mm / inches)	1,220 / 48.1	1,220 / 48.1	1,220 / 48.1
Width (mm / inches)	762 / 30	762 / 30	1,524 / 60
Height (mm / inches)	1,960 / 77.2	1,960 / 77.2	1,960 / 77.2
Weight (Kg / lbs)	500 / 1,123	598 / 1,343	1,196 / 2,685

Environmental Characteristics

Acoustics	65 dB
Operating temperature	68° to 86° F (20° to 30° C)
Non-operating temperature	-40° to 158° F (-40° to 70° C)
Maximum rate of temperature change	68° F/hr (20° C/hr)
Operating relative humidity	15% to 80% @ 86° F (30° C)
Operating altitude	0 to 10,000 ft (0 to 3.1 km)
Non-operating altitude	0 to 15,000 ft (0 to 4.6 km)

Regulatory Compliance

Safety	IEC 950:1991+A1, +A2, +A3, +A4; EN60950:1992+A1, +A2, +A3, +A4, +A11; UL 1950, 3rd edition; cUL CSA C22.2 No. 950-95
--------	---

Key Dates

First CPL date	9/00	9/00	9/00
First ship date	4Q /00	4Q /00	4Q /01

Technical Specifications

¹ Due to limitation of HP UX 11i, maximum memory in a partition is 512 GB

² Superdome PA-8600/PA-8700 support only fully populated cells. Superdome PA-8800 and PA-8900 support both half and fully populated cells.

³ A second I/O expansion cabinet is required for if number of PCI slots is greater than 168

⁴ SPU cabinet must be filled first before placing I/O chassis in I/O expansion cabinet

⁵ A second I/O expansion cabinet is required if number of I/O chassis is greater than 14

NOTE: Given that PA 8600/PA 8700 are single core processors and PA 8800 and PA 8900 are dual core processors, the columns listed in this table refer to 16 processor, 32 processor and 64 processor. This terminology refers to 16 cores, 32 cores and 64 cores for Superdome PA 8600/PA 8700 systems and 32 cores, 64 cores and 128 cores for Superdome PA 8800 and PA 8900 systems.

Superdome I/O Expansion SPU Description (IOX) Cabinet Specifications	Superdome 16 processor	Superdome 32 processor	Superdome 64 processor
SPU Product Number	A6113A	A5201A	A5201A+A5202A
Maximum Number of I/O Chassis Enclosures (ICEs)*	3	3	3
Peripherals Supported	All peripherals qualified for use with Superdome and/or for use in an RBII D rack are supported in the I/O expansion cabinet as long as there is available space. Peripherals not connected to or associated with the Superdome system to which the I/O expansion cabinet is attached may be installed in the I/O expansion cabinet.		
Servers Supported	With PA-8600/PA-8700: Superdome 32 processor, Superdome 64 processor With PA-8800 and PA-8900: Superdome 64 cores, Superdome 128 cores		
Electrical Characteristics			
AC input power	200-240 VAC, 50/60 Hz		
Current requirements at 200V-240V	16A		
Typical Maximum Power Dissipation	2,290 watts		
Maximum Power Dissipation	3,200 watts		
Dimensions			
Height (mm / inches)	1.6 meters or 1.96 meters		
Depth (mm / inches)	45.5 in (same depth as 32 W)		
Width (mm / in)	24.0 in		
Environmental	Same as Superdome		
Relevant Product Numbers			
12 slot PCI Chassis for Rack System E Expansion Cabinet	A48564Z	A48564Z	A48564Z
12 slot PCI -X Chassis for Rack System E Expansion Cabinet	A6864Z	A6864Z	A6864Z

Technical Specifications

I/O Expansion Cabinet Power and Utilities Subsystem	A5861A	A5861A	A5861A
I/O Expansion Power and Utilities Subsystem Graphite color	A5861D	A5861D	A5861D
I/O Chassis Enclosure for 12 slot PCI X Chassis	A5862A	A5862A	A5862A

Key Dates

First CPL Date	9 / 00
First Ship Date	2Q / 01

***NOTE:** Each ICE holds two I/O card cages or 24 PCI X I/O slots.

APC SL20KFB2 Specifications

Description	APC Silcon, 20000VA/20000W, Input 115/200 3PH, 120/208 3PH, 127/220 3PHV/ Output 115/200 3PH, 120/208 3PH, 127/220 3PHV, Interface Port DB-25 RS-232, Contact Closure		
General Features	0% to 95% non-condensing, 200% overload capability, Audible Alarms, Built in static bypass switch, Delta Conversion On line Technology, Environmental Protection, Event logging, Extendable Run Time, Full rated output available in kW, Input Power Factor Correction, Intelligent Battery Management, LCD Alphanumeric Display, Overload Indicator, Paralleling Capability, Sine wave output, SmartSlot, Software, Web Management		
Includes	Parallel Card, Triple Chassis for three SmartSlots, User Manual, Web/SNMP Management Card		
Spare parts kits	See APC website http://www.apcc.com		
Documentation	User Manual and Installation Guide		
Input	Nominal input voltage	115/200 3PH, 120/208 3PH, 127/220 3PH V	
	Input frequency	50 Hz programmable +/- 0.5, 1, 2, 4, 6, 8%; 60 Hz programmable +/- 0.5, 1, 2, 4, 6, 8%	
	Input connection type	Hardwire 5-wire (3PH + N + G)	
	Input voltage range for main operations	170-230 (200 V), 177-239 (208 V), 187-242 (220 V) V	
Batteries	Typical backup time at half load	36.7 minutes	
	Typical backup time at full load	10.7 minutes	
	Battery type	Maintenance-free sealed Lead-Acid battery with suspended electrolyte: leak proof	
	Typical recharge time **	2 hours	
Physical	Maximum height dimensions	55.12 in (140.00 cm)	
	Maximum width dimensions	39.37 in (100.00 cm)	
	Maximum depth dimensions	31.50 in (80.01 cm)	
	Net weight	1,290.00 lbs (586.36 kg)	

Technical Specifications

	Shipping Weight	1,340.00 lbs (609.09 kg)
	Shipping Height	66.93 in (170.00 cm)
	Shipping Width	43.31 in (110.00 cm)
	Shipping Depth	35.43 in(90.00 cm)
	Color	Dark green (NCS 7020 B50G), Light gray (NCS 2703 G84Y)
	Units per Pallet	1.0
Communications and Management	Interface port	DB-25 RS-232, Contact Closure
	Smart Slot Interface Quantity	2
	Pre-Installed SmartSlot Cards	AP9606
	Control panel	Multi-function LCD status and control console
	Audible alarm	Beep for each 52 alarm conditions
	Emergency Power Off (EPO)	Yes
	Optional Management Device	See APC website http://www.apcc.com
Environmental	Operating Environment	68° to 86° F (20° to 30 °C) (recommended is 68° to 77° F (20° to 25° C))
	Operating Relative Humidity	15% to 80% at 86°F (30°C)
	Operating Elevation	0 to 10000 ft (0 to 3050 m)
	Storage Temperature	-40° to 158° F (-40° to 70° C)
	Storage Relative Humidity	0% to 90%
	Storage Elevation	0 to 15,000 ft (0 to 4575 m)
	Audible noise at 1 meter from surface of unit	55 dBA
Conformance	Online thermal dissipation	4,094 BTU/hour
	Protection Class	NEMA 1, NEMA 12
	Approvals	EN 55022 Class A, ISO 9001, ISO 14001, UL 1778, UL Listed, cUL Listed
	Standard warranty	One-year repair or replace, optional on-site warranties available, optional extended warranties available
	Optional New Service	See APC website http://www.apcc.com

* Without TAX/VAT

** The time to recharge to 90% of full battery capacity following a discharge to shutdown using a load rated for 1/2 the full load rating of the UPS

Supported I/O Cards with Superdome PA-8600/PA-8700

I/O Card	Product Number	Bootable	Connector Type(s)	Maximum Cards by processors per System			Special Notes
				16	32	64	
Mass Storage Host Bus Adapters							

Technical Specifications

Fast/Wide/Differential SCSI	A4800A	Yes	HD	48	48 ³	96 ³	
PCI 1 port 2x Fibre Channel ⁴	A5158A	Yes	Duplex SC	48	48 ³	96 ³	
Dual port FWD SCSI	A5159B	Yes	VHDCI	24	48	96	
PCI 2 Gb/s Fibre Channel	A6795A		LC	48	96	192	
PCI X 2 channel 2 Gb /s Fibre Channel	A6826A		LC	16	32	64	
PCI 2 channel Ultra320 SCSI	A7173A	Yes	VHDCI	48	96	192	

Local Area Network (LAN) Adapters

PCI 1 port Universal FDDI LAN	A3739B		FDDI SC	16	32	64	
PCI 1 port 1000Base SX ⁴	A4926A		Duplex SC	16	32	64	
PCI 1 port 1000 BaseT ⁴	A4929A			16	32	64	
PCI 1 port 10/100Base TX	A5230A		RJ-45	24	48	96	
PCI 1 port ATM 622 Mb/s (MMF)	A5483A		Duplex SC	8	16	32	
PCI 4 port 100Base TX	A5506B		RJ-45	8	16	32	
PCI 1 port 802.5 Token Ring 4/16/100	A5783A		RJ-45 and DB 9	8	16	32	
PCI 1 port 1000Base T (gigabit copper)	A6825A		RJ-45	16	32	64	
PCI 1 port 1000Base SX (gigabit fiber)	A6847A		Duplex SC	16	32	64	
PCI-X 2-port 1000Base-SX Gigabit Adpt	A7011A			16	32	64	
PCI-X 2-port 1000Base-T Gigabit Adpt	A7012A			16	32	64	

Multi-Function Cards (Mass Storage / LAN)

PCI 2 port 100Base T/ 2 port Ultra2 SCSI	A5838A		VHDCI/RJ-45	8	16	32	
PCI-X 2Gb Fibre Channel / 1000BaseSX	A9782A	Yes	LC	16	32	64	
PCI-X 2Gb Fibre Channel / 1000BaseTX	A9784A	Yes	1 LC, 1 RJ-45	16	32	64	
HP PCI-X Multifunction 2-port 2Gb FC / 2-port 1 Gb Ethernet Adapter	AB465A	Yes	2 LC	16	32	64	3.3 volts only. Supported on HP-UX 11i, HP-UX 11i v2. Factory integration available.

Wide Area Network (WAN) Adapters

PCI 1 port ATM 155 Mbps (MMF)	A5513A		Duplex SC	8	16	32	
2 port Programmable Serial Interface (PSI) X.25 / Frame Relay / SDLC	J3525A		RS-530, RS-232, V.35, RS-449 or X.21	8	16	32	

Technical Specifications

4 port X.25/Frame Relay	J3526A		RS-530, RS 232 V.35, RS-449 or X.21	8	16	32	
PCI 8 port X.25/BX.25 Advance Communications Controller	Z7340A		RS-232, V.35, RS-449 or X.21	25	25	25	

Additional Interface Cards

PCI HyperFabric 4X Fiber ²	A6092A		DB - 37	8	8	8	
PCI HyperFabric 2 Fibre	A6386A		LC Duplex	8	8	8	
PCI 8-port Terminal Multiplexer	A6748A		RS-232	8	14	14	
PCI 64-port Terminal Multiplexer ¹	A6749A		RS-232 or RS-422	8	14	14	
PCI 8-port Terminal Multiplexer ⁴	J3592A		RS-232	8	16	32	
PCI 64-port Terminal Multiplexer ⁴	J3593A		RS-232 or RS-422	8	16	32	

¹ Must order option 001 for Port Module Accessory Kit that enables the cables to connect the MUX to the port modules.

² Discontinued. Upgrade to A6386A PCI HyperFabric2 fiber adapter

³ Supports up to 96 or 192 cards when using an expansion cabinet

⁴ No longer on CPL-not orderable

NOTE: All cards listed have the OLAR capability except for SNA over X.25. Additionally, the Superdome core I/O card in slot 0 does not support OLAR. X.25 dual port (J3525A) needs separate cables for each card port (2 ports per card).

NOTE: The MSA30 SB/DB are supported as a boot disk on Superdome running HP UX 11i with the following cards: A7173A, A6828A, A6829A, A5838A

Supported I/O Cards with Superdome PA-8800 and PA-8900

I/O Card	Product Number	Bootable	Connector Type(s)	Maximum Cards by Processors per System			Special Notes
				16	32	64	
Mass Storage Host Bus Adapters							
Fast/Wide/Differential SCSI 2	A4800A	Yes	HD	48	48	96	
PCI 1 port Ultra2 SCSI 2	A5149A	Yes	VHDCI	48	48 ¹	96 ¹	
PCI 2x Fibre Channel 2	A5158A	Yes	LC	48	48 ¹	96 ¹	
Dual port FWD SCSI	A5159B	Yes	VHDCI	24	48	96	
PCI 2 Gb/s Fibre Channel	A6795A	Yes	LC	48	48 ¹	96 ¹	
PCI X 2 channel 2 Gb /sFibre Channel	A6826A	Yes	LC	48	48 ¹	96 ¹	
PCI 1 channel U160 SCSI	A6828A	Yes	VHDCI	48	48 ¹	96 ¹	
PCI 2 channel U160 SCSI	A6829A	Yes	VHDCI	48	48 ¹	96 ¹	
PCI 2 channel Ultra320 SCSI	A7173A	Yes	VHDCI	48	96	192	
Local Area Network (LAN) Adapters							
PCI 10/100 Base-T	A3738A		AUI, BNC, RJ-45	24	48	96	
PCI 1 port Universal FDDI LAN	A3739B		FDDI SC	16	32	64	

Technical Specifications

PCI 1 port 1000Base SX 2	A4926A		Duplex SC	16	32	64	
PCI 1 port 1000 BaseT 2	A4929A		RJ-45	16	32	64	
PCI 1 port 10/100Base TX	A5230A		RJ-45	24	48	96	
PCI 1 port ATM 622 Mb/s (MMF)	A5483A		Duplex SC	8	16	32	
PCI 4 port 100Base TX	A5506B		RJ-45	8	16	32	
PCI 1 port 802.5 Token Ring 4/16/100	A5783A		RJ-45 and DB-9	8	16	32	
PCI 1 port 1000Base T (gigabit copper)	A6825A		RJ-45	16	32	64	
PCI 1 port 1000Base SX (gigabit fiber)	A6847A		Duplex SC	16	32	64	
PCI-X 2-port 1000Base-SX	A7011A		Duplex SC	16	32	64	
PCI-X 2-port 1000Base-T	A7012A		RJ-45	16	32	64	
PCI -X 2 port 4x Fabric (HPC) Adapter2	AB286A		4x Infiniband Copper	8	8	8	
PCI-X 133MHz 10GbE SR Fiber Adapter	AB287A		Duplex LC	8*	16*	32*	*Max cards are 2 per I/O chassis. Therefore: 16P= 8 max 32P= 16 max 64P=32 max
PCI-X 4-port 1000 Base-T Gigabit Adapter	AB545A	Yes	RJ-45	16	32	64	Stand-alone shipments
PCI -X 24 port 4x Fabric Copper Switch	AB399A		4x Infiniband Copper	8	8	8	
PCI-X 4-port GigE-T	AB545A	Yes	RJ-45	16	32	64	Available July 2005.
PCI-X 10 GigE	AB287A	No	Duplex LC	2	2	2	Available July 2005.

Multi-Function Cards (Mass Storage / LAN)

PCI 2 port 100Base T/ 2 port Ultra2 SCSI	A5838A	Yes	VHDCI/RJ-45	8	16	32	
PCI-X 2Gb Fibre Channel / 1000BaseSX	A9782A	Yes	LC	48	48 ¹	96 ¹	
PCI-X 2Gb Fibre Channel / 1000BaseTX	A9784A	Yes	1 LC, 1 RJ-45	48	96	192	
PCI-X multifunction 2-port 1000BT and dual-port U320 SCSI adapter	AB290A	Yes (SCSI and GigE)	SCSI - LVD / SE LAN - RJ 45	48	96	192	Available July 2005
PCI-X 2-port FC + 2-port GigE-T	AB465A	Yes (FC and GigE)	2 LC 2 RJ-45	48	96	192	Available July 2005

Wide Area Network (WAN) Adapters

PCI 1 port ATM 155 Mbps (MMF)	A5513A		Duplex SC	8	16	32	
2 port Programmable Serial Interface (PSI) X.25 / Frame Relay / SDLC	J3525A		RS-530, RS-232, V.35, RS-449 or X.21	8	16	32	

Technical Specifications

4 port X.25/Frame Relay	J3526A		RS-530, RS 232 V.35, RS-449 or X.21	8	16	32	
PCI 8 port X.25/BX.25 Advance Communications Controller	Z7340A		RS-232, V.35, RS-449 or X.21	25	25	25	
Additional Interface Cards							
PCI HyperFabric 4X Fiber	A6092A		DB - 37	8	8	8	
PCI HyperFabric 2 Fibre	A6386A		LC Duplex	8	8	8	
PCI 8-port Terminal Multiplexer	A6748A		RS-232	8	16	32	
PCI 64-port Terminal Multiplexer ¹	A6749A		RS-232 or RS-422	8	16	32	
¹ Supports up to 96 or 192 cards when using an expansion cabinet ² Not Orderable - No longer on CPL							

NOTE: The MSA30 SB/DB are supported as a boot disk on Superdome running HP UX 11i with the following cards: A7173A, A6828A, A6829A, A5838A.

For HP Storage connectivity, the webpage: <http://spock.corp.hp.com/index.aspx> has all of the detail for HP hardware. Please consult this matrix for HP supported on-line and near-line storage.

Peripherals Required Per Partition (nPar)

	HP-UX 11i version 1 (HWE)
I/O Cards	Core I/O (Slot 0) provides console and LAN
	Default boot device (Slot 1)
	Removable media card (Slot 8)
Peripherals	DVD hard drive (boot disk)
	DDS-4/DAT-40 tape backup drive
	C7508AZ or C7508A (Qualec device)

© Copyright 2006 Hewlett-Packard Development Company, L.P.

The information contained herein is subject to change without notice.

Microsoft and Windows Server 2003 are US registered trademarks of Microsoft Corporation. Intel and Itanium are US registered trademarks of Intel Corporation.

The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.