



meet the
hp superdome
servers

september 2001

a white paper from
Hewlett-Packard
Company

table of contents.

abstract	3
the hp 9000 enterprise server product line	3
the need for “big iron”	4
introducing superdome enterprise servers	4
top-of-the-line performance and capacity	5
availability features for maximum uptime	5
investment protection, with support for PA-RISC or IPF processors	5
better manageability	5
hp-ux 11i, the proven 64-bit operating environment	5
superdome hardware overview	6
superdome hardware	7
the superdome cabinet	7
superdome system architecture	8
modular architecture	9
the cell board	10
processors	11
cell controller ASIC	11
memory controller ASIC	12
the crossbar backplane	12
crossbar ASIC	13
memory and backplane latencies	14
the I/O subsystem	15
superdome: performance and scalability	16
system bandwidth	16
scalability	16
superdome: high availability	16
superdome: system reliability	17
supportability	17
hardware inventory in support tools manager	18
support management station	18
high availability observatory	18
repairability	19
n+1 CPUs with iCOD	19
n+1 fans, power supplies, backplane DC/DC converters	20
dual power sources	20
online addition and replacement (OLAR) of PCI I/O cards and chassis	20
online addition and replacement of cell boards	20
online addition of nPartitions	20
mc/serviceguard- in-a-box	20
superdome: investment protection	20
microprocessor roadmap	21
binary compatibility for across PA-RISC	22
binary compatibility for Itanium Processor Family	22
upgrading a superdome system	22
superdome partitioning	22
the hp partitioning continuum	23
hyperplex: hard partitions with multiple nodes	23
nPartitions: hard partitions within a node	23
virtual partitions	24
resource partitions	24
how nPartitions work	24
uses for nPartitions	24
increased systems utilization	25
increased flexibility via support of multiple environments	26
increased uptime	26
increased scalability	26
superdome: ease of management	26
superdome support management station	26

SMS installation	27
SMS functional capabilities	27
partition manager	27
servicecontrol suite	27
superdome and hp-ux 11i	28
features of hp-ux 11i	28
hp-ux 11i operating environments	28
entry-level hp-ux 11i	29
hp-ux 11i enterprise	29
hp-ux 11i mission-critical	29
the rest of the superdome story	29
the always-on infrastructure and superdome	29
peripherals and superdome	29
mass storage and superdome	29
conclusion	29
availability	29
capacity	30
connectivity	30
security	30
manageability	30
for more information	30

abstract

Superdome is the high-end UNIX® server in the HP Server product family. It's designed to deliver the performance, availability, capacity, security, and manageability needed for the "mission-critical Internet" and other compute-intensive applications. In this white paper, you'll learn about the architecture, system software, and management tools of the Internet-ready HP Superdome Enterprise Server.

the hp 9000 enterprise server product line

First, let's look at the overall HP Server product lineup. These servers address the major computing challenges that customers face today. Running the powerful HP-UX operating environment (Hewlett-Packard's version of UNIX), HP Enterprise Servers are especially suited to the demands of the Internet, online transaction processing (OLTP), system consolidation, large decision support systems, and compute-intensive technical applications.

As shown in **figure 1**, the HP Server line is extremely comprehensive, with products ranging from entry-level products to high-end platforms. And now, added to the existing A-, L-, and N-class servers, Superdome takes its place as the family's most powerful high-end platform.



figure 1: the hp server family

The HP Server product line represents the most scalable range of computing available today. These systems feature:

- *Application compatibility and portability* across the product line, permitting users to select the best platform at the appropriate price point without worry about whether applications will run on it.
- *The highest uniprocessor performance* in the industry, providing reduced time-to-solution through parallelism and increased throughput through multiprocessing, in all environments.
- *Unmatched scalability*, providing the best price/performance over the entire range of products and protecting customers' investments in hardware and software.
- *A consistent programming model*, presenting the same application programming environment regardless of the systems' performance levels; this greatly increases the number of available "off-the-shelf" third-party applications and reduces porting and development costs.
- *Leading-edge utility pricing models* (Pay per Use, Capacity on Demand, and Pay per Forecast), allowing customers to better align their IT costs with revenues as their technology usage fluctuates over time.

the need for "big iron"

Some applications simply need the maximum amount of performance and availability. For instance, in the "Internet powerhouse"—the data center where transaction processing is handled and where the backbone of mission-critical applications runs—you'll want the best servers available to keep those applications running and meeting quality-of-service goals on a 24 x 7 basis. The Internet powerhouse, along with other workload types such as OLTP with database engines, data mining, and customer relationship management, are best handled by "big iron," top-of-the-line computers that also combine security, scalability, flexibility, capacity, and manageability. That's where Superdome fits in.

introducing superdome enterprise servers

Superdome is the latest addition to the HP family of HP-UX Enterprise Servers, and it extends the capabilities of this server family in several important areas:

- Scalability
- Performance
- Capacity
- Availability
- Flexibility
- Investment protection
- Manageability

As a high-end platform for the Internet data center, Superdome is uniquely able to address your requirements, even as those requirements change over time. This capability is based on Service Level Agreement (SLA) features, hard and soft partitions, scalability up to 64 industry-leading processors, and multiple operating environments (OEs). Superdome is also the *first high-end server family to be upgradable to Itanium™ Processor Family (IPF) processors!*

top-of-the-line performance and capacity

Superdome is the best choice for the Internet powerhouse because it greatly extends the capabilities of HP Servers at the high end:

- High-performance RISC processor (PA-8700)
 - 750-MHz, 4-way superscalar (3.0 GFLOPS)
 - 0.18 micron copper silicon-on-insulator CMOS process
 - Data prefetching capability
 - Large on-chip cache (2.25 MB)
 - Upgradable to IPF processors
- Increased memory subsystem performance
 - Increases peak memory bandwidth to 64 GB/s per 64-way system
 - Higher density system packaging provides 256 GB per 64-way system
- Doubled processor capacity compared to V-class
 - Configurations available from 2 to 64 CPUs
- Doubled I/O interface performance; more than doubled I/O bandwidth
 - 64-bit 33 MHz (2x) industry-standard PCI
 - 64-bit 66 MHz (4x) industry-standard PCI
 - Sixteen 1.2 GB/s I/O channels (32 GB/s per 64-way system)
 - Increased number of PCI slots (up to 192)

availability features for maximum uptime

Mission-critical applications are no problem for Superdome. This family of servers has the increased system reliability and availability necessary to provide *maximum uptime*. It has improvements in system quality, resiliency, and fault management; and it provides support for MC/Serviceguard in a Superdome cluster.

flexibility with partitions

Superdome supports the HP Partitioning Continuum, allowing both hard and soft partition. In fact, Superdome can be configured as one large symmetric multiprocessor or as several independent hard partitions, known as nPartitions. These nPartitions can provide hardware isolation and complete software isolation, and you can dynamically resize partitions for the highest degree of flexibility. Partitions allow Superdome to reach new heights in availability and scalability, too.

investment protection, with support for PA-RISC or IPF processors

At first release, Superdome supports PA-RISC processors, with support for IPF processors coming soon. This makes it the logical choice not only for your present and future high-end computing needs, but also for ensuring investment protection in the years to come.

better manageability

Superdome includes all-new management features coupled with tested management software and technologies, all designed to increase your control while reducing administrative overhead. With tools such as the Superdome Support Management Station (SMS), Partition Manager, and Servicecontrol suite at your fingertips, you'll find that a Superdome system is extremely easy to manage.

hp-ux 11i, the proven 64-bit operating environment

HP-UX 11i is an operating environment specifically focused on the needs of both online economy companies and traditional economy concerns. Superdome runs this industry-leading operating environment.

HP-UX 11i is an Internet-focused version of HP's 64-bit operating system that enjoys the industry's greatest support from independent software vendors. In fact, you can choose from more than 16,000 applications, including native 64-bit versions of all major databases and leading enterprise resource planning (ERP) applications. The widely supported 64-bit HP-UX operating system is compatible with the full line of HP Servers.

superdome hardware overview

You can choose Superdome in several different configurations, with capabilities from 2-way multiprocessing all the way to 64-way. Depending on the model, Superdome comes in one or two cabinets. Here's a summary of Superdome hardware:

	superdome 16-way	superdome 32-way	superdome 64-way	superdome I/O expansion (see note 1)
4-CPU cell boards (hot-swap capability offered with HP-UX releases after 11i)	2-4	4-8	8-16	
CPUs	2-16	4-32	8-64	
memory (with 512 MB DIMMs)	4-64 GB	8-128 GB	16-256 GB	
12-slot I/O card cages	4	4/8	8/16	6
hot-swap PCI I/O slots	48 slots (32 33-MHz slots, 16 66-MHz slots)	48/96 slots (64 33-MHz slots, 32 66-MHz slots)	96/192 slots (128 33-MHz slots, 64 66-MHz slots)	72
hot-swap redundant power supplies (N+1 included)	4	6	12	2
I/O fans	6	6	12	4 per I/O chassis enclosure (max 12)
hot-swap redundant fans (N+1 included)	4	4	8	
nPartitions	4	4/8	8/16	
crossbar bandwidth (peak)	16 GB/s	32 GB/s	64 GB/s	
2X PCI I/O bus bandwidth	266 MB/s	266 MB/s	266 MB/s	
4X PCI I/O bus bandwidth	533 MB/s	533 MB/s	533 MB/s	
cell controller to I/O subsystem bandwidth (peak)	2.0 GB/s	2.0 GB/s	2.0 GB/s	
I/O bandwidth (peak)	8 GB/s	16 GB/s	32 GB/s	
cell controller to memory subsystem bandwidth (peak)	4 GB/s	4 GB/s	4 GB/s	
memory bandwidth (peak)	16 GB/s	32 GB/s	64 GB/s	
average memory load-to-use latency:				
4-CPU	235 ns	235 ns	235 ns	
8-CPU	266 ns	266 ns	266 ns	
16-CPU	296 ns	296 ns	296 ns	
32-CPU	N/A	336 ns	336 ns	
64-CPU	N/A	N/A	360 ns	
average idle memory latency (system is idle):				
4 CPU	174 ns	174 ns	174 ns	
8 CPUs	208 ns	208 ns	208 ns	
16 CPUs	228 ns	228 ns	228 ns	
32 CPUs	N/A	261 ns	261 ns	
64 CPUs	N/A	N/A	275 ns	
operating system	HP-UX 11i	HP-UX 11i	HP-UX 11i	HP-UX 11i
availability	1st release	1st release	June/July 01	June/July 01
number of cabinets	1 (left)	1 (left)	2 (1 left, 1 right)	1 Rack System E Expansion
dimensions:				
height	1.96 m (77.2 in)	1.96 m (77.2 in)	1.96 m (77.2 in)	1.96 m (77.2 in) 1.60 m (63.0 in)
width	762 mm (30 in)	762 mm (30 in)	1524 mm (60 in)	610 mm (24 in)
depth	1220 mm (48 in)	1220 mm (48 in)	1220 mm (48 in)	1220 mm (48 in)

Note 1: For I/O expansion up to 168 slots, one Expansion Cabinet (holds up to 6 I/O card cages) is required. For I/O expansion up to 192 slots, two Expansion Cabinets are required. Any remaining space in the I/O Expansion Cabinets can be used to store peripherals. Both the 1.96m and 1.6m heights are available for the I/O Expansion Cabinet.

superdome hardware

For maximum flexibility in configuring the data center, a Superdome system has up to four different types of cabinet assemblies:

- Superdome 16-way and Superdome 32-way systems are single-cabinet configurations with one Superdome left cabinet. The Superdome cabinet contains all of the processors, memory, and core devices of the system. It also houses up to 48 PCI I/O cards or 4 I/O chassis.
- The Superdome 64-way system is a dual-cabinet configuration with one Superdome left cabinet and one Superdome right cabinet. Together, the two cabinets contain all of the processors, memory, and core devices of the system. The dual-cabinet configuration houses up to 96 PCI I/O cards or 8 I/O chassis.
- An optional I/O Expansion Cabinet is added if the required number of PCI I/O cards exceeds the number of PCI I/O cards that can be accommodated in a 32-way or 64-way configuration.

the superdome cabinet

Figure 2 shows the physical layout of a Superdome 32-way cabinet. At the top of the cabinet, hot-swappable main fans are installed after the cabinet arrives at the customer site. Below these fans is a cage for the eight cell boards on which processors and memory DIMMs reside. For all HP-UX 11 releases after 11i, these cell boards support hot-swap capability, so you can change them without bringing down the system.

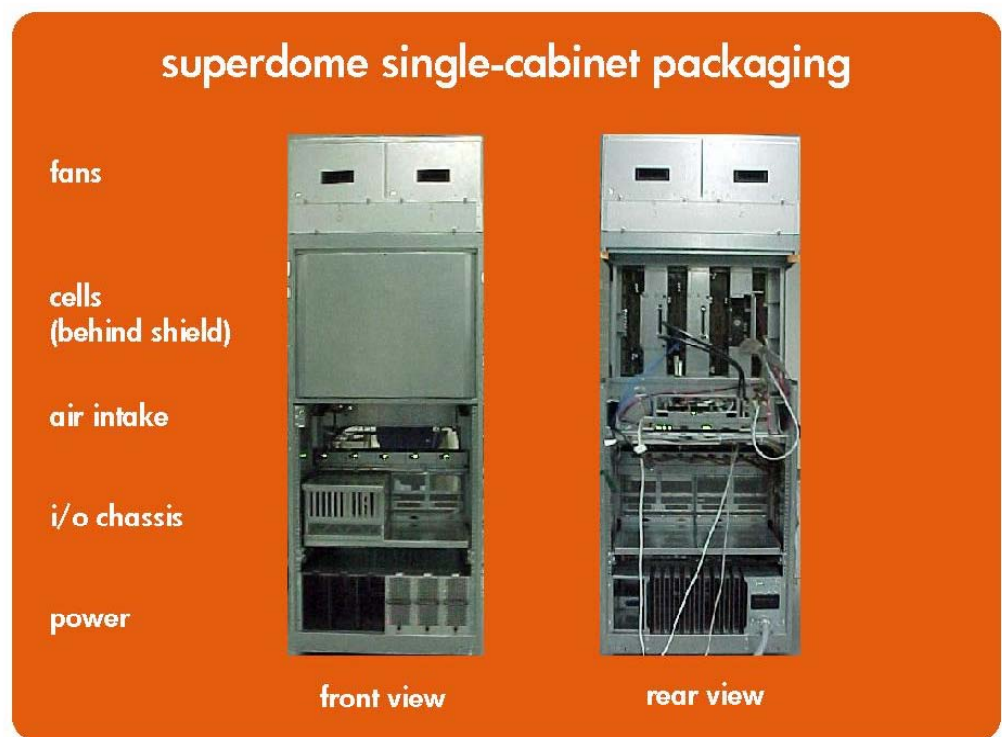


figure 2: front and rear views of superdome 32-way cabinet

Directly below the cell boards is the main air intake, and below that are two I/O chassis. Each I/O chassis holds an I/O card cage. Redundant power supplies are at the bottom of the cabinet. The Superdome family does not use electric plugs; instead, -48 Vdc power is hard-wired to each cabinet. There are two redundant power inputs so that the system can be powered up via two different power grids. An opening in the side of the cabinet allows two Superdome 32-way cabinets to be cabled together to constitute a Superdome 64-way.

superdome system architecture

Features of the Superdome cabinets include:

- **Front and rear servicing:** You can service Superdome cabinets from the front and rear of the cabinet. This lets you arrange the cabinets of your Superdome system in the traditional row fashion found in most computer rooms.
- **Sized for easy installation:** Installation is no problem—the width of the Superdome cabinet allows entry through most doorways without disassembly.
- **Hot-swappable filters:** The intake air to the cell boards is filtered, helping to keep these boards clean and operational longer. You can remove the filter for cleaning or replacement while the system is fully operational.
- **External status display:** A status display on the outside of the front and rear doors of each cabinet allows maintenance personnel and HP field engineers to determine the basic status of each cabinet without opening any cabinet doors.

Superdome is a ccNUMA (cache-coherent, Non-Uniform Memory Access) system. And Superdome presents an SMP (symmetric multi-processing) programming model to the operating system by enabling any processor to have access to any byte of memory anywhere in the system. In fact, Superdome is the first mission-critical UNIX system to exploit distributed processors and memory. Here's why:

- **Usable bandwidth scales with system size, thanks to two important design features:**
 - **Superdome's coherency scheme:** Unlike some UNIX systems, which rely on snoop-based coherency that results in bottlenecks (large, flat latencies and bandwidth starvation) in high-end configurations, Superdome sidesteps the overheads usually attributed to directory coherency.
 - **Superdome's topology:** Superdome implements a point-to-point global packet switch for a communication fabric that is very well balanced across processor, memory, and I/O traffic.
- **Large physical memory with extremely low latency:** At first release, Superdome uses 128Mbit SDRAM technology and provides 2X DRAMs per processor, yielding a maximum memory of 256 GB. In the future, Superdome will be the first system in the industry to surpass 1 TB memory capacity. And even with these high amounts of memory, the latency growth from 4 to 64 processors is only 1.6X (60% growth), which is extremely flat.
- **High I/O bandwidth and connectivity:** Superdome provides a high degree of I/O connectivity while preserving bandwidth. There are 192 PCI cards in the system, each with its own dedicated PCI bus; at the adapter level, this aggregates to 64 GB/s of raw bandwidth. The I/O bandwidth available between the system core (processors and memory) and the I/O controllers is roughly 30 GB/s.
- **Large number of high-performance processors:** Superdome provides 64 PA-8700 processors—and these are the highest-performance processors in the industry. (Compare this system to those from other UNIX system vendors; while the systems may provide up to 64 processors, their uniprocessor performance does not match that of PA-8700.)
- **True hardware isolation of system resources:** Processors, memory, and I/O resources are truly isolated from each other in order to provide flexibility in system usage. This means:
 - Great SMP performance to attack large single workloads.
 - Hardware-enforced isolation of resources, coupled with great single-system high availability and manageability features for a strong consolidation platform.

Superdome's architecture is modular, with components that can be used to construct several different server-class computer products based on PA-8700 processors. (Processor upgrades beyond PA-8700 require a new version of the cell board.) Such modularity helps provide investment protection—you can simply add or change components as system requirements change.

The Superdome platform supports a variety of system configurations from 2 to 64 processors. The system is extremely flexible and has a cell-based hierarchical crossbar architecture that can be configured as one large symmetric multiprocessor or as several independent nPartitions.

You can choose from three models in the Superdome server family. The largest is an 8- to 64-processor system with a scalable memory and I/O subsystem. In a Superdome system, any processor on any nPartition can directly address any byte of memory on any nPartition, through processor-issued load or store instructions.

There are three basic components in the Superdome system architecture: the cell or cell board, the crossbar backplane, and the PCI-based I/O subsystem. **Figure 3** is an illustration of this architecture.

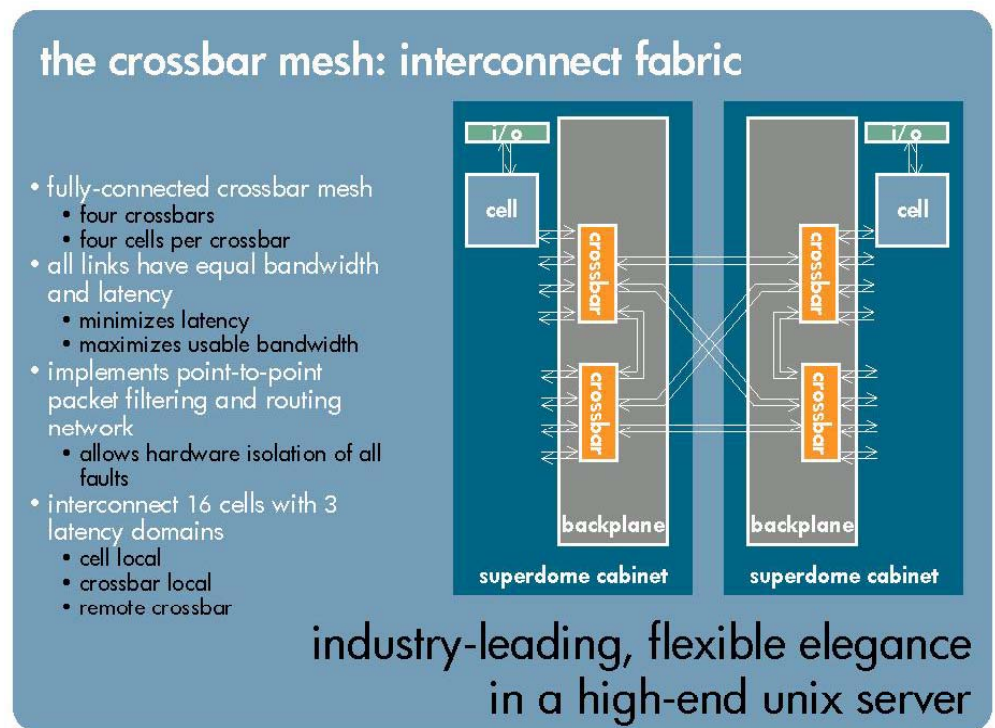


figure 3: superdome's hierarchical crossbar architecture

the cell board

A cell, or cell board, shown in **figure 4**, is the basic building block of a Superdome system. Each cabinet can contain up to eight cell boards, which are plugged into the backplane of the cabinet.

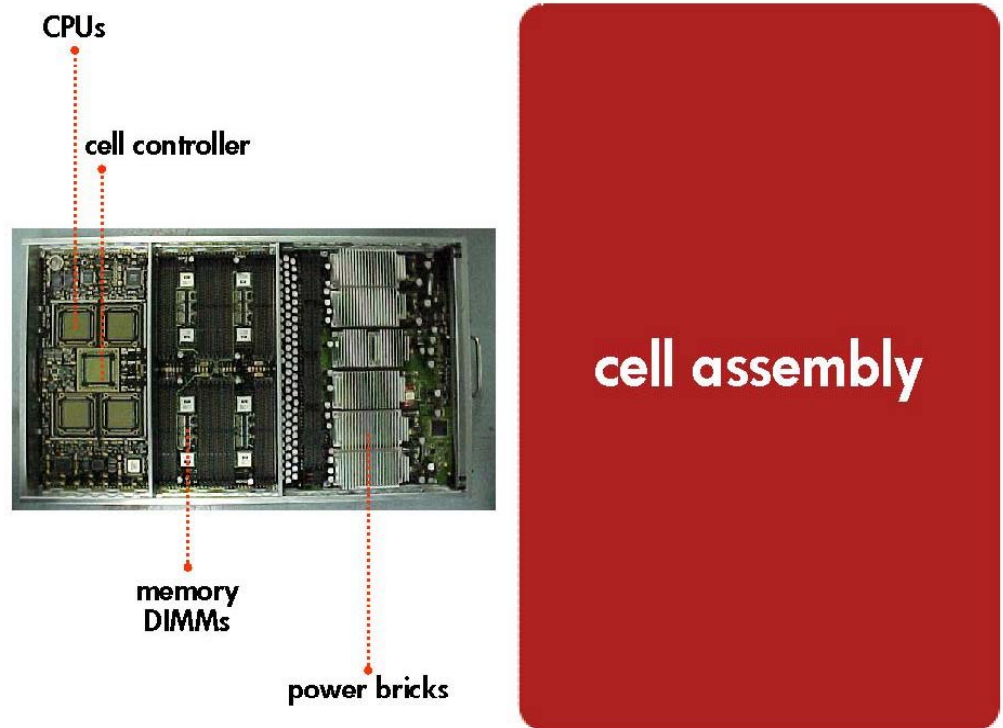


figure 4: superdome cell board layout

Each cell board is a self-contained unit, with a symmetric multiprocessor (SMP), main memory, and all necessary hardware:

- CPUs (up to 4 processors)
- Cell controller ASIC (application-specific integrated circuit)
- Main memory DIMMs (up to 16-GB RAM in 2-GB increments, using 128-Mbit SDRAMs)
- Power bricks (power converters)
- Data buses
- Optional link to 12 PCI I/O slots

Figure 5 illustrates the cell board architecture. The cell has a peak memory bandwidth of 3.2 GB/s (4 GB/s in the next release). A connection to a 12-slot PCI card cage is optional for each cell, and the peak bandwidth of this link is 1.2 GB/s (1.8 GB/s next release). Bandwidth to the crossbar is 6.4 GB/s (8 GB/s next release).

Except for the DIMM connections, most links in the cell are point-to-point connections. All links are protected by parity (processor link, I/O link) as well as ECC—error checking and correcting (all memory data and fabric ports).

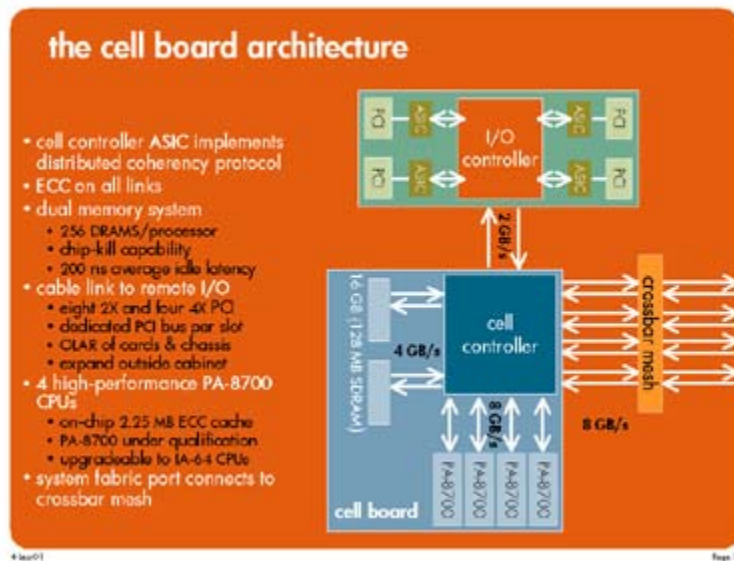


figure 5: superdome cell board and interconnect architecture

In the future, cell boards will support up to 32 GB of memory using 256-Mbit SDRAMs.

processors

The initial release of the Superdome uses ECC-protected PA-8700 processors running with a core frequency of 750 MHz and an on-chip cache of 2.25 MB. Future Superdome implementations will support Itanium Processor Family processors as well as future PA-RISC processors.

cell controller ASIC

Residing on the cell board, the cell controller ASIC (application-specific integrated circuit) is part of the Superdome system chip set. It coordinates memory between the major components of a cell board and determines if a request requires communication with another cell or with the I/O subsystem. HP believes that this cell controller ASIC is the largest ASIC in the world: it has some 24 million transistors and six layers of metal, using copper technology.

The cell controller ASIC has five major interfaces:

- 2 ports to memory controllers (DIMMs)
- 4 ports to processors (one dedicated port per processor for the highest possible performance)
- 1 port to the crossbar, through which all communication to other cells flows
- Processor-dependent hardware (PDH)
- 1 port to the I/O bus, which connects the cell to an I/O subsystem

In addition to the interface logic, the cell controller ASIC maintains cache coherency throughout the system. The cell controller ASIC supports the PA-8700 processor and can support future PA-RISC processors.

The processor-dependent hardware, or PDH, is the module that provides the cell-local resources required to reset a cell and bring it up to a point where it can join other cells and boot the operating system. PDH contains the system boot firmware, which is also used at runtime.

memory controller ASIC

The memory controller ASIC is also part of the Superdome system chip set. Its primary function is to multiplex and demultiplex data between the cell controller ASIC and the SDRAMs in the memory subsystem. When the cell controller ASIC issues transactions on the memory interface command bus, the memory controller ASIC buffers those read transactions and returns the data as soon as it is available.

Note that only the data portion of the memory subsystem goes through the memory controller ASIC. All address and control signals to the DIMMs are generated by the cell controller ASIC and sent directly to the DIMM via the memory interface address bus.

The memory subsystem is a dual-ported implementation. It supports memory DRAM fault tolerance, where a discrete SDRAM chip can fail without compromising data integrity. The memory subsystem provides 4 GB/s of bandwidth to the cell controller ASIC and minimizes the overhead typically associated with directory coherency. What's more, the memory subsystem enjoys a very low latency for cell-to-local-memory access: as low as 200 ns average idle latency at 250 MHz load-to-use.

the crossbar backplane

Each crossbar backplane contains two sets of two crossbar ASICs that provide a non-blocking connection between eight cells and their associated memory and I/O. Each backplane cabinet can support up to eight cells or 32 processors (in a Superdome 32-way). Two backplanes can be linked together with a flex cable to produce a cabinet that can support up to 16 cells or 64 processors (Superdome 64-way).

crossbar ASIC

The crossbar ASIC is yet another part of the Superdome system chip set. It implements a high-performance 8-port non-blocking crossbar and the 400-MHz crossbar link protocol. (The next release will support 500 MHz.) All ports are functionally and electrically identical. Superdome's fabric is a fully connected crossbar mesh with four crossbar ASICs and four cells per crossbar ASIC. Like the cell controller ASIC, the crossbar ASIC has truly impressive scale. HP believes it to be the world's second largest ASIC, with 18 million transistors and six layers of metal, using copper technology.

A very important aspect of the crossbar mesh is that all links have the same bandwidth and latency—which is key to maximizing overall bandwidth and minimizing overall latency. Cell-to-crossbar and crossbar-to-crossbar communication occur at the same speed; there are no excessive latency penalties for going remote. In addition, Superdome's memory interleaves across cells first, then across memory banks. This interleaving scheme tends to balance out memory traffic across all the links.

The crossbar mesh implements a global point-to-point packet filtering network. This mesh features an extremely high level of integrity, with each crossbar port fully independent. The crossbar mesh has dedicated paths for data and control, and each port can be reset, assigned, or reconfigured fully independent of other ports. Superdome's crossbar mesh is an excellent foundation for resource isolation.

The crossbar ASIC offers several features that contribute to Superdome's high performance:

- Support for scaling up to a 256-way coherent shared memory system.
- 250-MHz operation
- 500-MT/s link speed
- Support for two interleaved channels on link protocol.
- Support for double-length data packets for Itanium Processor Family mode.
- Performance counters to enable software tuning.

Each port on the crossbar ASIC is 6.4 GB/s peak (8 GB/s peak in next release).

These ports provide a high-throughput path to cells and other crossbar ASICs:

- Four ports connect to the four cells that reside on the crossbar ASIC (one port per cell).
- Three ports connect to the remaining three crossbar ASICs (in a Superdome 64-way).
The remaining port may connect to another system to form a multi-node configuration (Scalable Computing Architecture) that is currently not supported.

The total crossbar bandwidth for each Superdome model is calculated like this:

$$(Number\ of\ cells\ \times\ peak\ crossbar\ bandwidth\ per\ cell) \div 2\ ports$$

Crossbar bandwidth for the different Superdome systems is excellent:

- For Superdome 16-way, the crossbar bandwidth is 16 GB/s
- For Superdome 32-way, the crossbar bandwidth is 32 GB/s
- For Superdome 64-way, the crossbar bandwidth is 64 GB/s

*memory and backplane
latencies*

Superdome has been designed to reduce memory and backplane latencies, providing optimum performance. There are three types of memory latencies within the Superdome system:

- *Memory latency within the cell* refers to the case where an application runs on an nPartition that consists of a single cell board.
- *Memory latency between cells on the same crossbar* refers to the case where the nPartition consists of two cells that reside on the same crossbar. In this case, 50 percent of the addresses are to the memory of the requesting processor's cell board and the other 50 percent of the addresses are to the memory of the other cell board.
- *Memory latency between cells on different crossbars* refers to the case where the nPartition consists of cell boards that do not reside on the same crossbar. In this case, if there are 16 cells in the nPartition, 1/16 of the requests are to the memory of the cell board in which the processor resides. 3/16 of the requests go to the memory of the other three cell boards on this same crossbar. Finally, the remaining 12/16 of the requests transit two crossbars.

Superdome memory latency depends on the number of CPUs and the location of their corresponding cell board. Assuming that there is equally distributed access to all memory controllers and that cell boards are installed to minimize latency, the average memory idle latency (when nothing is executing on the system) and memory latency (load-to-use) are shown below.

number of CPUs	average idle latency	average memory latency
4-CPU	174 ns	235 ns
8-CPU	208 ns	266 ns
16-CPU	228 ns	296 ns
32-CPU	261 ns	336 ns
64-CPU	275 ns	360 ns

the I/O subsystem

Superdome cells have remote I/O; this enhances modularity and means there are no trade-offs for scaling of processors, memory, and I/O. Each cell can connect to a remote I/O chassis through an I/O cable link.

Superdome's I/O subsystem has plenty of capability for today and expansion for tomorrow. Each I/O module consists of twelve PCI connections, divided among eight 2X and four 4X PCI slots, with an I/O controller ASIC and power. Each PCI slot has its own PCI bus; the 2X PCI slot has 266-MB/s bandwidth and the 4X PCI slot achieves a bandwidth of 533 MB/s. The point-to-point connectivity allows the earliest detection, containment, and recovery from errors.

Any I/O module can support a core I/O card (required for each independent nPartition). Superdome 16-way, 32-way, and 64-way systems can accommodate four, four, and eight I/O modules, respectively, within the system cabinet. This provides a total of 48, 48, and 96 PCI slots, respectively.

An I/O Expansion Cabinet can also be added. On the Superdome 32-way system, this yields an additional 48 PCI slots, giving a maximum connectivity of 96 PCI slots for 8 cells. For a Superdome 64-way system, the I/O Expansion Cabinet provides an additional 96 PCI slots, for a maximum connectivity of 192 PCI slots for 16 cells.

A system configured with 16 cells, each with its own I/O module and core I/O card can support up to 16 independent nPartitions. Note that cells can be configured without I/O modules attached, but I/O modules cannot be configured in the system unless they are attached to a cell.

The I/O subsystem bandwidth is 2.0 GB/s per cell. This makes the total I/O subsystem bandwidth 8.0 GB/s, 16.0 GB/s, and 32 GB/s for Superdome 16-way, Superdome 32-way, and Superdome 64-way systems, respectively.

**superdome:
performance and
scalability**

Superdome offers a host of features that are especially important in the OLTP, business intelligence, and technical markets. And bandwidth and scalability are two important areas of Superdome leadership.

system bandwidth

	superdome 16-way	superdome 32-way	superdome 64-way
crossbar bandwidth (peak)	16 GB/s	32 GB/s	64 GB/s
cell controller to I/O subsystem bandwidth (peak)	2.0 GB/s	2.0 GB/s	2.0 GB/s
I/O bandwidth (peak)	8 GB/s	16 GB/s	32 GB/s
memory bandwidth	16 GB/s	32 GB/s	64 GB/s

scalability

	superdome 16-way	superdome 32-way	superdome 64-way	superdome I/O expansion (see note 1)
4-CPU cell boards (hot-swap capability offered with HP-UX releases after 11i)	2-4	4-8	8-16	N/A
CPUs	2-16	4-32	8-64	N/A
memory (with 512-MB DIMMs)	4-64 GB	8-128 GB	16-256 GB	N/A
12-slot I/O card cages	4	4/8	8/16	8
hot-swap PCI I/O slots	48 slots (32 33-MHz slots, 16 66-MHz slots)	48/96 slots (64 33-MHz slots, 32 66-MHz slots)	96/192 slots (128 33-MHz slots, 64 66-MHz slots)	96
nPartitions	4	4/8	8/16	N/A

Note 1: For I/O expansion up to 168 slots, one Expansion Cabinet (6 I/O card cages) is required. For I/O expansion up to 192 slots, two Expansion Cabinets (2 I/O card cages) are required. Any remaining space in the I/O Expansion Cabinets can be used to store peripherals. Both the 1.96m and 1.6m heights are available for the I/O Expansion Cabinet.

For more information about performance and scalability, please refer to the Superdome Performance White Paper.

**superdome:
high availability**

Superdome offers unprecedented single-system availability in the following areas:

- System reliability
- Supportability
- Repairability
- MC/Serviceguard-in-a-Box

superdome: system reliability

The Superdome product family is bolstered by design enhancements and production techniques that substantially enhance system reliability. Among these enhancements are:

- Memory DRAM fault tolerance—that is, recovery of a single SDRAM failure
- DIMM address parity protection
- Dynamic memory resilience—page de-allocation of bad memory pages during operation
- Dynamic processor resilience
- CPU cache ECC protection and automatic de-allocation
- CPU bus parity protection
- Full single-bit error detection and correction on crossbar and I/O links
- I/O error recovery and system resilience to I/O card failures
- I/O cards fully isolated from each other
- Prevention of silent corruption of data going to I/O
- Recovery of a single I/O-to-cell controller link failure
- 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
- ASIC full burn-in and “high quality” production process, resulting in 10X improvement in ASIC failure rates
- Full “test to failure” and accelerated life testing on all critical Superdome assemblies
- Strong emphasis on quality for multiple-nPartition single points of failure (SPOFs)
- System resilience to Guardian Service Processor (GSP)
- Isolation of nPartition failures; failure of one nPartition will not bring down another partition
- Protection of nPartitions against spurious interrupts or memory corruption

supportability

Supportability is another key feature of the Superdome product line. Some of the supportability features of Superdome include:

- Event Monitoring Service (EMS)
- Improved Support Tools Manager
- Support Management Station
- High Availability Observatory (HAO)

event monitoring service

For Superdome, HP’s Event Monitoring Service aids in system management by keeping track of the system’s “vital signs.” EMS lets you monitor virtually all of the hardware on the system, including:

- Mass storage
- Memory
- Fibre Channel components (MUX, switch, card, fabric, etc.)
- I/O cards
- ECC errors on the main system bus
- ECC errors in the CPU cache
- System temperature
- Support processor problems
- Processor dependent hardware problems
- Processor dependent hardware battery low
- Cabinet fans
- Cabinet power supplies
- Chassis code logging failure
- Processor dependent code (firmware) problems

The system hardware configuration and selected kernel parameters are also monitored.

When a hardware monitor detects a problem, it generates an event that can be reported to you via DMI, SNMP, OpenView Vantagepoint Operations, e-mail, page, syslog, console, or a selected text log file. Each event contains a complete description of the problem, a severity classification (informational, warning, serious, critical), and text that shows the probable cause and the action to take.

When excessive processor cache errors are detected, the processor is automatically removed from use until it is replaced.

*hardware inventory
in support tools
manager*

With Superdome, HP's Support Tools Manager (STM) has been enhanced to give you detailed information about the system. Superdome has been designed with the ability for each individual field replaceable unit (FRU) to report information such as serial number, part number, revision level, etc. This information is available in the "Miscellaneous" information tool in STM. Moreover, all of the hardware inventory information is accessible via desktop management interface (DMI).

*support management
station*

To decrease the possibility of unscheduled downtime still further, the Superdome Support Management Station (SMS) is a separate server that provides diagnosis and test capabilities for the Superdome system. The SMS includes ASIC-level scan tools that can be used by HP Support Engineers to diagnose Superdome systems in your data center. (For more details of the SMS, see "Superdome Support Management Station" on page 26.)

*high availability
observatory*

Superdome systems can be supported by Hewlett-Packard's High Availability Observatory (HAO), a support solution designed to keep mission-critical environments up and running. The HAO lets HP's Mission Critical Support Centers reduce the number, duration, and impact of outages on covered systems.

The HAO is a state-of-the-art suite of tools, processes, and support personnel that combine to provide significant support capabilities for systems covered by mission-critical HP support contracts. The HAO solution improves the availability of Superdome systems by:

- Conducting frequent, automated data collection of the customer's mission-critical environment
- Providing a secure, high-bandwidth link for HP support personnel to remotely access the customer's environment to conduct diagnostic tests and potentially solve reported problems
- Providing reactive and proactive support for HP's mission-critical customers

In addition to covering Superdome and HP-UX 11i, HAO technology covers HP-UX servers running HP-UX 10.20 and 11.0, as well as HP, Dell, and Compaq servers running Windows NT® 4.0 (with Service Pack 3 or later) or Windows® 2000, as these systems are identified in your mission-critical support contract. Network interconnect devices, such as hubs, switches, bridges, and routers manufactured by HP and Cisco, are also covered by HAO tools.

There are three major architectural components of the HAO: the HP Support Node, the high-bandwidth secure link, and the Mission Critical Support Center (MCSC).

- **HP Support Node:** This is an HP-owned and operated HP workstation that resides within your environment. Its job is to act as the depot for various tools that capture your system configurations and network topologies for use by HP support personnel. HAO tools residing on the HP Support Node collect the configuration information and send it to one of HP's Mission Critical Support Centers. The HP Support Node also serves as an HP-managed resource to analyze core dumps and help with other reactive support. The HP Support Node is the entry point for HP support engineers to access your mission-critical systems.
- **High-Bandwidth Secure Link:** The high-bandwidth secure link, or ISDN link, connects HP and the mission-critical customer. This link provides HP with the capability to react quickly when reactive support requires accessing your systems to help resolve a problem. It also is a path used to transport data from the HP Support Nodes to the MCSC for proactive analysis.
- **Mission Critical Support Center:** The MCSC resides inside HP. Its job is to act as the primary support center to collect and track customer information and to perform proactive problem analysis; it also provides a control point in the HAO remote support architecture. To maintain the security of the your environment, the MCSC allows access to the isolated LAN only to authorized HA (high availability) Support Engineers.

A Remote Support Resource (RSR) is placed inside the isolated LAN, providing a secure method for authorized HA Support Engineers to access customer sites.

Configuration analyzer servers in the MCSC take configuration information collected from all HP Support Nodes and continually analyze the data. If an irregularity is identified, a case is opened in the MCSC's workflow management system. MCSC Monitor, a tool at the MCSC, provides a display of customer enterprise configuration data transported nightly from each HP Support Node. When providing reactive support, HP Support Engineers check the MCSC Monitor for configuration changes, and they can view patch summary information when providing proactive support.

repairability

Superdome servers are endowed with a substantial number of features that make them easier to repair with little or no downtime. Among these features are:

- N+1 CPUs with instant capacity on demand (iCOD)
- Hot-swap N+1 fans, power supplies, and backplane DC/DC converters
- Online replacement of PCI I/O cards
- Online addition/replacement of PCI I/O chassis (follow-on release to HP-UX 11i)
- Online addition/replacement of cell boards (follow-on release to HP-UX 11i)
- Dual power sources

n+1 CPUs with iCOD

With HP-UX 11i, Superdome has support for instant capacity on demand (iCOD). This feature lets you bring additional CPUs online without a system reboot. Because no reboot is needed, the system has no loss of availability, even while you're adding capacity.

Another advantage of iCOD is that it allows setting up a system of N+1 CPUs, ensuring maximum single-system availability. If one CPU fails, another is already running and prepared to take its place.

n+1 fans, power supplies, backplane DC/DC converters

Superdome servers are equipped with N+1 fans, power supplies, and backplane DC converters, ensuring maximum availability for these components.

dual power sources

Dual power sources in Superdome servers mean the power supply can be protected against becoming a single point of failure.

online addition and replacement (OLAR) of PCI I/O cards and chassis

PCI card and chassis OLAR enables the online addition and replacement of PCI I/O cards and chassis on Superdome systems. The system hardware uses per-slot power control combined with operating system support for the PCI card OLAR feature to allow you to add a new card and replace an existing card without affecting other components or requiring a reboot.

online addition and replacement of cell boards

Superdome and a future HP-UX release after 11i support online addition and replacement of cell boards, allowing repair and maintenance of these critical components without bringing the system down.

online addition of nPartitions

With HP-UX 11i, Superdome lets you add nPartitions without affecting other running nPartitions. Such dynamic reconfiguration is another reason why Superdome provides superior uptime. And Superdome is enabled for dynamic cell board migration, which will be fully supported with the next release of HP-UX 11i.

mc/serviceguard-in-a-box

In order to increase the uptime of applications within a Superdome system, you can configure nPartitions in an MC/Serviceguard cluster so that the cluster membership is within the system. Upon detection of a failure within an nPartition, MC/Serviceguard fails over the application to another nPartition within the Superdome system.

superdome: investment protection

The Superdome family provides outstanding customer investment protection and lasting value, thanks in part to a system infrastructure designed to accommodate several generations of processor upgrades. For example, Superdome will be board-upgradable to future PA-RISC and IA-64 processors. And all major system components other than the cell board and power conversion card remain the same for PA-8600 or PA-8700 processors; these processors even use the same memory DIMMs. What's even better is the customer's investment in PA-8600 processors is protected, as PA-8600 and PA-8700 processors can be mixed in the same Superdome (but not in the same partition).

microprocessor
roadmap

The microprocessor roadmap in **figure 6** illustrates how the processors used in Superdome evolved. It also illustrates HP's commitment to long-term processor and architecture innovation, which ensures that HP will continue to provide the massive resources needed for future computing requirements.

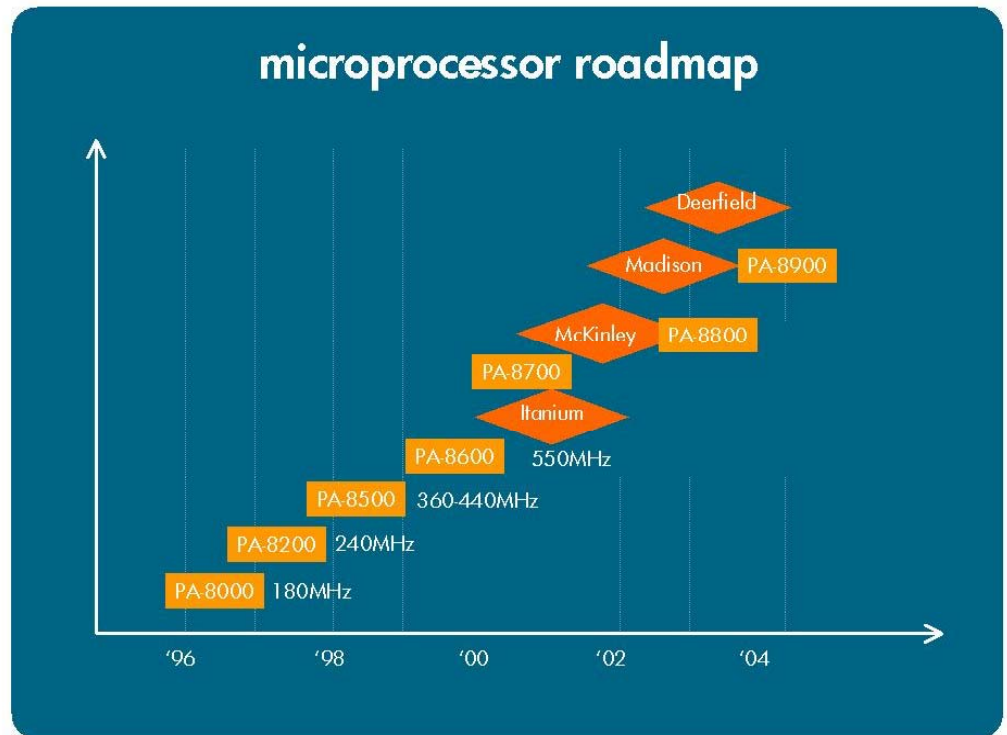


figure 6: hp microprocessor roadmap

The processor roadmap shows not only HP's leading line of RISC processors but also the introduction of binary compatible IPF processors. Notice that HP is investing in several PA-RISC enhancements after the introduction of the IPF; this means you'll be able to move to the new Itanium Processor Family architecture when you're ready, not when forced to by a vendor. Superior microprocessors and binary compatibility help make Superdome the fastest high-end server on the market, and the safest investment for the future.

*binary compatibility for
across PA-RISC*

You can continue to rely on binary compatibility across the PA-RISC family, enabling seamless interoperability with legacy applications on HP systems. Binary compatibility protects your investments, enabling rapid growth and adoption of new technology infrastructures. For performance improvements, you can use existing applications and operating systems with new or more advanced processor technology as it develops.

*binary compatibility for
Itanium Processor Family*

HP will continue to support binary compatibility through the introduction of Itanium Processor Family based systems. As a result of HP's work with Intel® on EPIC (Explicitly Parallel Instruction Computing) architecture, the technology foundation for the Itanium Processor Family, today's HP-UX, Windows 2000, and Linux applications will run unchanged on the Itanium Processor Family. To ensure maximum performance, you can recompile applications without source changes.

*upgrading a
superdome system*

Upgrading an existing Superdome system is as simple as removing all the cell boards and the memory contained on them. You transfer the memory to the new processor cell boards, plug the cell boards into the cabinet, and resume operation.

In the future, Superdome will be upgradable to follow-on PA-RISC processors and to IPF processors. A new cell board and power conversion card are required.

Processor upgrades occur on a cell board basis and can be performed one nPartition at a time in order to minimize downtime for the entire Superdome complex. You can upgrade either PA-RISC or IPF processors in this manner (although mixing PA-RISC and IPF processors in one system is currently not supported).

Today's Superdome is already prepared for the next generation of processors, ensuring that Superdome will stay ahead of tomorrow's performance demands. Whether it's a PA-RISC or an IPF follow-on processor, Superdome delivers the investment protection afforded by in-chassis upgrades.

superdome partitioning

With the addition of Superdome, HP now offers an extremely comprehensive partitioning continuum. Partitions are physical or logical mechanisms for isolating operational environments within single or multiple servers, so that applications can enjoy protection from unrelated events that could otherwise cause disruption, interruption, or performance degradation. Partitioning allows isolating operational environments to ensure privacy and uptime while maintaining the highest degree of flexibility.

the hp partitioning continuum

The HP Partitioning Continuum provides a broad range of partitioning solutions designed to meet the diverse needs of our customers. **Figure 7** illustrates the various types of partitions. This is the only UNIX solution on the market to simultaneously address the often contradictory needs of flexibility and uptime.

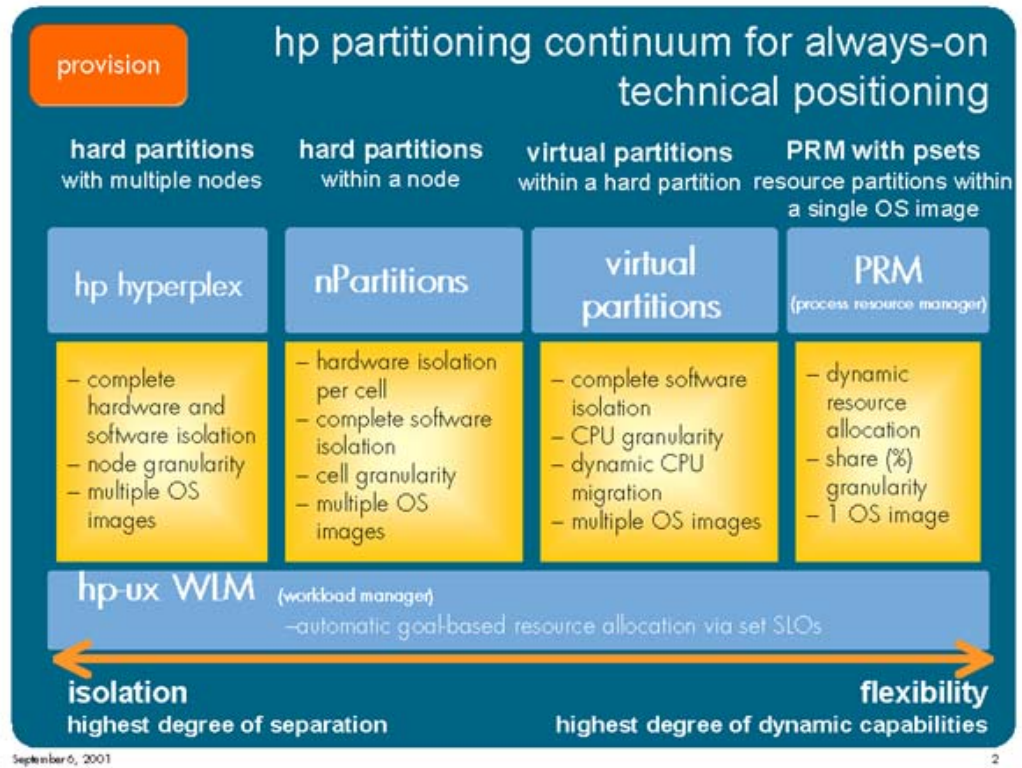


figure 7: hp’s partitioning continuum combines high isolation with excellent flexibility

hyperplex: hard partitions with multiple nodes

Hard partitions with multiple server nodes are called the Hyperplex solution. Hyperplex delivers the optimum capacity at all levels by supporting the complete HP UNIX server product line. And Hyperplex is extremely scalable—in fact, a Hyperplex solution can range in size from as few as two entry-level server nodes up to hundreds of the latest Superdome servers, providing the maximum possible capacity.

The hard partitions with multiple nodes afforded by Hyperplex are designed to allow total isolation from other hard partitions. Multiple applications can run in these partitions, and these applications are completely isolated from other nodes and their respective operating environments.

nPartitions: hard partitions within a node

Hard partitions within a node are called nPartitions. A unique feature of HP Superdome servers, nPartitions form the most powerful HP high-end Hyperplex server nodes. A Superdome server can support anywhere from one to 16 nPartitions, each supporting its own operating system, applications, peripherals, and networks. Currently each nPartition can host HP-UX operating environments. When the Intel IPF versions of Superdome are available, nPartitions will support Linux and Windows environments as well.

virtual partitions

Virtual partitions provide even more portioning granularity, and HP's virtual partitions are unique in the UNIX industry. Superdome supports up to 16 nPartitions and up to 64 virtual partitions. So a system with 16 nPartitions can have as many as four virtual partitions in each nPartition—or, to put it another way, one virtual partition per CPU. Each virtual partition runs its own image of the HP-UX 11i operating system and can fully host its own applications.

This solution offers complete software isolation across virtual partitions. Software isolation guarantees that an application running in one virtual partition is not affected by a software crash in another. Although virtual partitions are susceptible to hardware failures, protection against hardware failures is provided with Hyperplex or nPartitions. Virtual partitions also let you dynamically adjust partition size (without rebooting) by the dynamic addition and deletion of CPUs, effectively moving them from one virtual partition to another.

resource partitions

Very often many applications run on one server at the same time, but each application has different resource needs. (As an analogy, think of electricity use: many homes use electricity from the same generator, but some homes need more power than others.) HP's resource partitions are unique partitions created just for this type of workload management.

Resource partitions run within hard partitions or virtual partitions, and are controlled by workload management functions within the operating environment. HP-UX Workload Manager (WLM) and Process Resource Manager (PRM) software dynamically create resource partitions for applications that need guaranteed amounts of dedicated resources such as CPU cycles, memory, or disk I/O. Applications with specific goals, such as response time, can use HP's goal-based HP-UX WLM to automatically and dynamically allocate the necessary resources to applications or user groups within hard partitions or virtual partitions. With resource partitions and workload management tools, varying service level objectives can be met every time!

how nPartitions work

Within Superdome, cell boards are grouped into physical structures. An nPartition consists of one or more cells that communicate coherently over a high-bandwidth, low-latency crossbar fabric. Special programmable hardware in the cell boards defines the boundaries of an nPartition in such a way that isolation from the actions of other nPartitions is enforced.

Each nPartition runs its own independent operating system, and different nPartitions can be executing the same or different versions of an OS. In a Superdome IPF system, they can even be executing different operating systems altogether (such as HP-UX, Linux, and Windows).

Each nPartition has its own independent set of CPUs, memory, and I/O resources. You can use system management commands to change resources from one nPartition to another without having to physically change the hardware. Moreover, using the dynamic reconfiguration capabilities of HP-UX 11i, such as online addition and removal, you can add new resources to an nPartition and delete resources while the nPartition remains in operation. In addition, dynamic additions of new nPartitions are supported.

uses for nPartitions

With nPartitions, you can configure Superdome as one large symmetric multiprocessor or as several independent nPartitions. Using independent nPartitions has advantages in several areas: increased systems utilization, increased flexibility by supporting multiple environments, increased uptime, and increased scalability.

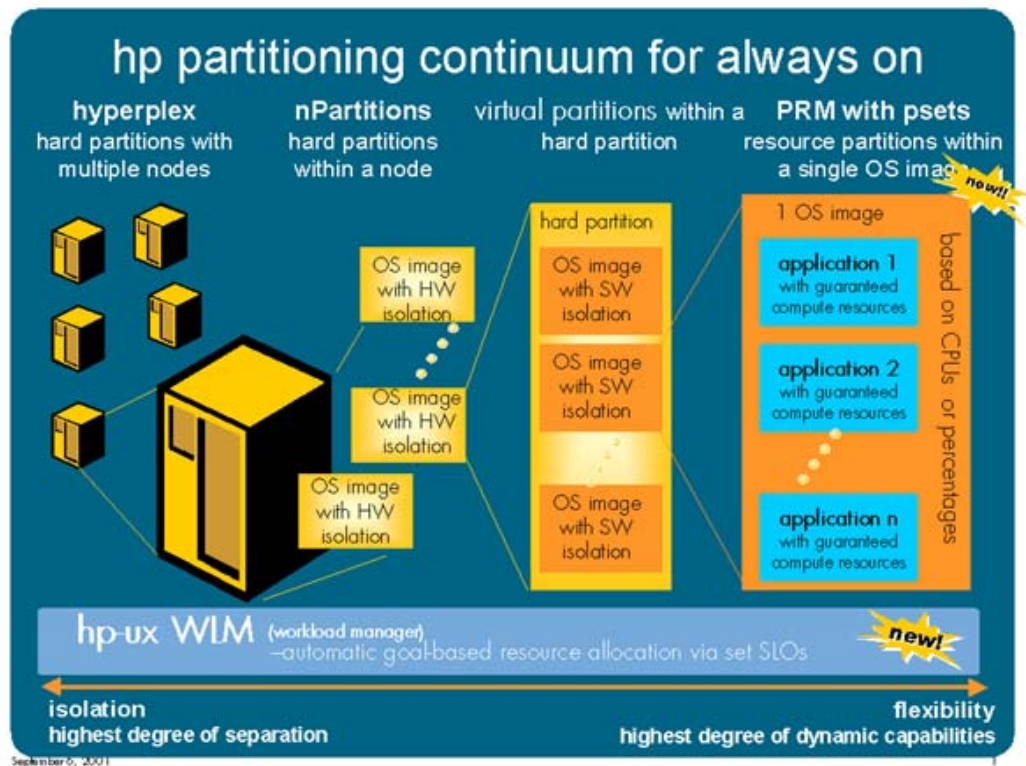


figure 8: using superdome's partitioning continuum

increased systems utilization

Businesses today require that multiple applications running on the same server deliver full performance without conflicting or interfering with one another. It is critical that the actions of one application do not impact other applications or result in denial of service to them. Furthermore, a single failure in the operating system or hardware components on the server supporting the set of applications must not result in the loss of multiple applications running on that server.

Superdome nPartitions are the answer. With Superdome nPartitions, each nPartition is isolated from the others. The dynamically configurable hardware provides an effective firewall between applications so that there is virtually no possibility of interference between nPartitions. The result is increased systems utilization with security—the perfect answer for systems consolidation scenarios.

*increased flexibility
via support of multiple
environments*

Business requirements and competitive pressures can dictate the need for a large set of applications, both old and new. Often applications require different versions or revisions of operating systems, different system parameter settings, and even different patch levels. Performance can suffer from the interactions of applications with incompatible behaviors.

Superdome nPartitions can provide multiple independent environments for these applications. For instance, with nPartitions a portion of the compute and I/O resources can be configured to run an old copy of the operating system, while the rest of the machine is updated to run the newer software. Or you can use nPartitions to set up a test environment where new software and operating systems can be installed while leaving production applications to continue operation without impact.

Future releases will provide the ability to run different operating systems such as HP-UX, Linux, and Windows on the same Superdome platform. And for an even higher degree of flexibility you can run virtual and resource partitions within an nPartition—and change them on the fly—to even further optimize your resource utilization.

increased uptime

nPartitions provide independence of failure or operation. If you have to reboot one nPartition or take it down to perform an operating system upgrade, the other nPartitions in the system are not affected.

Unlike a traditional large SMP, nPartitions allow a server to be configured into a cluster of independent systems. Since failure independence is provided by nPartitions, high availability clusters, consisting of two or more nPartitions, can be configured within the same server to improve single-system availability. In addition, you can configure MC/Serviceguard to fail over applications from one nPartition to another nPartition within the same Superdome system or to another HP Server.

increased scalability

Some applications may not have the ability to scale linearly with large numbers of CPUs. Superdome's nPartitions allow you to tailor the processor, memory, and I/O resources to match application scaling characteristics. What's more, as new application or operating system releases provide improved scaling characteristics, nPartitions allow the resources to be readjusted to optimize application performance.

**superdome: ease of
management**

In spite of all its power and flexibility, the Superdome system is remarkably easy to manage, making your life easier with a host of management features and technologies. These include:

- Superdome Support Management Station
- Partition Manager
- Servicecontrol suite

**superdome support
management station**

The Support Management Station (SMS) is used to run Superdome scan tools. These are the same scan tools used to enhance diagnosis and testability of the platform throughout Superdome development and manufacturing, and they are now available to HP field engineering organizations. This translates into better management, as well as faster and easier upgrades and hardware replacement.

SMS installation

One SMS is installed per customer site (or data center), connected to each Superdome platform via a private LAN. (Ideally, the SMS is installed near the associated Superdomes, because the HP Customer Engineer will run the scan tools and then be able to immediately make any necessary hardware repairs.) The physical connection from the platform is a private Ethernet connection.

The SMS can support two LAN interfaces:

- The dedicated connection to the systems to be supported.
- The connection to interface with the customer's general LAN.

These two LAN connections allow SMS operations to be performed remotely. More than one SMS can exist on a private LAN, but only one SMS should be actively using the LAN.

SMS functional capabilities

Here's what the SMS can do:

- Allows remote access via customer LAN (no modem access).
- Can be disconnected from the Superdome platform(s) and not disrupt their operation.
- Can connect a new Superdome platform to the SMS and be recognized by scan software.
- Supports multiple, heterogeneous Superdome platforms with scan software capability.
- Can scan one Superdome platform while other Superdome platforms are connected, without disrupting the operational platforms.
- Can run the scan software tools.
- Can run up to four scan processes concurrently, each on a different Superdome platform.
- Supports utility firmware updates.

partition manager

Another invaluable management feature for Superdome is Partition Manager. This tool provides control and management of Superdome nPartitions. You can launch it as a GUI from system administration manager (SAM) or directly from the command line, and with it you can:

- Display server status.
- Create and modify nPartitions.
- Display a complete hardware inventory.
- Display status of key server components.
- Check for problem or unusual server conditions.
- Manage power to cells and I/O chassis.
- Toggle attention indicators for cells, I/O chassis, I/O cards, and cabinets.

servicecontrol suite

HP's Servicecontrol suite addresses the requirements for centralized configuration, fault, and workload management, providing effective, efficient management of computing resources.

- HP Servicecontrol provides cost-effective control through truly multi-system management tools allowing rapid deployment, consistency, and asset management.
- Fault management in the Servicecontrol suite adds the full power of monitoring and proactive problem solving.
- Servicecontrol suite's workload management tools enable administrators to optimize Web and application workloads via dynamic resource allocation.

The different components of the Servicecontrol suite are tightly integrated into all the HP-UX 11i operating environments.

**superdome and
hp-ux 11i**

HP-UX 11 is HP's complete 64-bit UNIX operating environment that delivers significant scalability and performance for demanding applications. And the latest version of HP-UX is HP-UX 11i. When teamed with HP's leading server systems, HP-UX 11i provides the power of supercomputing at a fraction of the cost.

features of hp-ux 11i

HP-UX 11i offers:

- Excellent compatibility with Windows and Linux.
- Comprehensive Internet-critical functionality.
- Ultimate performance and scalability, including 64-way SMP support for Superdome.
- Industry-leading investment protection through binary compatibility.
- Superdome management features, including partition management.
- Support for nPartitions and virtual partitions.
- Excellent high availability features such as online addition and replacement for I/O and networking cards, dynamic patching, dynamically tunable kernel parameters, and up to 70 percent faster rebooting.
- Middleware and application software forward compatibility from HP-UX 11.0.

HP-UX 11i delivers an end-to-end scalable, manageable, and secure infrastructure for developing, deploying, and brokering mission-critical e-services. It has features targeted at Internet applications, and it contains functionality for managing Superdome servers.

hp-ux 11i operating environments

With HP-UX 11i, ordering and configuration have been vastly simplified. HP-UX 11i is available in three operating environments, as shown in **figure 9**. Each consists of basic HP-UX plus a collection of layered software products addressing your specific needs for availability, manageability, and security. These environments are simple to install and maintain, and they are a great value to purchase.

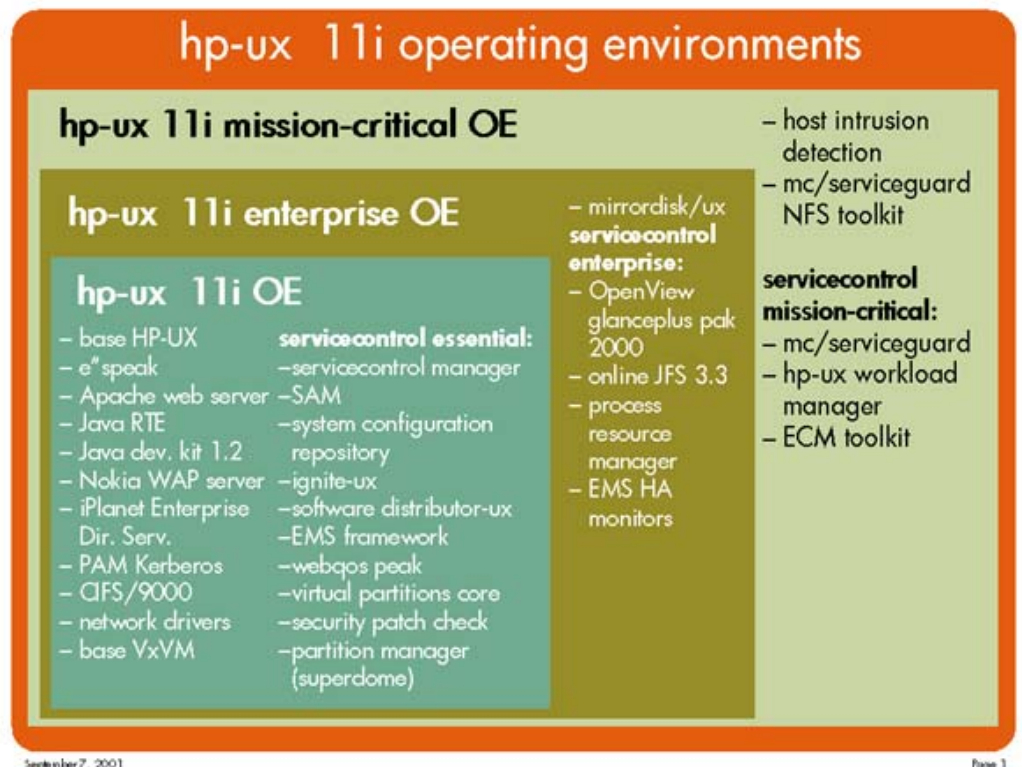


figure 9: hp-ux 11i operating environments

entry-level hp-ux 11i

The entry-level operating environment offered with every server addresses key needs from a Web and small application perspective. And new Superdome-specific functionality, such as Partition Manager, is tightly integrated with the base HP-UX. The core functionality of the new virtual partition capability is also included.

hp-ux 11i enterprise

The HP-UX 11i Enterprise OE is targeted at larger application servers typically served by mid-range systems such as the N-class. Customers in this environment usually need to optimize resources and do proactive capacity planning with higher uptime; so the Enterprise OE includes tools such as Process Resource Manager, Glanceplus Pak, and Mirrordisk/UX.

*hp-ux 11i
mission-critical*

The mission-critical OE is highly recommended for most Superdome systems because it specifically addresses the needs of large application and database server environments. This operating environment includes host intrusion detection to protect the system from outside attackers, MC/Serviceguard for the highest degree of HA, and HP-UX Workload Manager for automatic resource allocation based on set service levels. Goal-based resource allocation is a must in today's Internet world, and HP is today the *only* UNIX vendor offering such sophisticated resource allocation capability.

**the rest of the
superdome story**

This white paper has barely touched the surface of Superdome's capabilities. Here are some other areas of interest.

the always-on
infrastructure and
superdome

Because no Superdome stands in isolation, you can expect even more from HP. Superdome can be part of your always-on infrastructure, meeting the needs of your total IT environment throughout the life of the solution.

Superdome is delivered with a Foundation Configuration of services that ensure you get the right configuration for your solution the first time, and in record time. The services begin with up-front assessments of business/IT alignment, IT technical skills, and environment and mission-critical readiness. You'll also get a detailed architecture design based on your solution; factory pre-integration for software, peripherals, and middleware; and an HP solution manager assigned to ensure your satisfaction. You'll also be able to add mission-critical support services to meet the demanding requirements for stability and availability so characteristic of the Internet powerhouse.

peripherals and
superdome

Superdome peripherals are supported through industry-standard PCI. The PCI FWD SCSI host adapter provides a fast and wide (FWD) SCSI interconnect currently available at 40 MB/s. This adapter provides an ideal interface to high-performance disk array subsystems and high-performance tape drives.

mass storage and
superdome

Superdome supports a wide range of mass storage options, including several varieties of SCSI-attached RAID, Fibre Channel-attached RAID, SCSI-attached tape drives and tape libraries, and Fibre Channel-to-SCSI multiplexer connections.

conclusion

In this era of 24 x 7 applications and mission-critical Internet operations, businesses need a bulletproof infrastructure. That infrastructure is powered by HP's industry-leading, high-performing, enterprise servers and the robust HP-UX 11i operating environment. And Superdome represents another giant step forward for the HP server line.

availability

With its industry-leading multi-system availability, single-system reliability, and superior nPartition resilience, Superdome is designed to meet the needs of large enterprises, service providers, and established online economy companies.

capacity	Superdome provides a computing utility infrastructure unmatched by any competitor. Superdome's superior system and processor architecture provides scalability that outstrips the competition.
connectivity	Superdome offers massive amounts of I/O connectivity. And it has high-speed connectivity, too.
security	Thanks to HP Virtualvault and other leading security products, Superdome can offer unparalleled levels of security. Virtualvault today has been entrusted to protect some \$7 trillion in assets, testimony to its industry-wide acceptance.
manageability	HP's suite of management products are all ready to make Superdome management easier. And with its dynamic new capabilities as part of the HP Partitioning Continuum, Superdome provides the best combination of system flexibility and convenient management.
for more information	<p>Looking for more information about HP's Superdome?</p> <ul style="list-style-type: none"> • For more information about performance and scalability, please refer to the Superdome Performance White Paper. • For more information about high availability, please refer to the Superdome High Availability White Paper. • For more information on Superdome services, go to: www.hp.com/go/superdome

For more information, contact any of our worldwide sales offices or HP Channel.

For the location of the nearest sales office call:

United States of America: +1 800 637 7740

Canada:

Hewlett-Packard Ltd.
5150 Spectrum Way
Mississauga, Ontario L4W 5G1
+1 905 206 4725

Japan:

Hewlett-Packard Japan, Ltd.
Japan Country H.Q.
3-29-21, Takaido-Higashi, Suginami-ku,
Tokyo, 160-8585 Japan
+81 3 3331 6111

Latin America:

Hewlett-Packard
Latin American Region Headquarters
Waterford Building, 9th Floor
5200 Blue Lagoon Drive
Miami, Florida 33126 USA
+1 305 267 4220
Refer to country phone numbers

Australia/New Zealand:

Hewlett-Packard Australia Ltd.
31-41 Joseph Street
Blackburn, Victoria 3130
Australia (A.C.N. 004 394 763)
+61 3 9272 2895

Asia Pacific:

Hewlett-Packard Asia Pacific Ltd.
17-21/F, Shell Tower
Times Square
1 Matheson Street
Causeway Bay
Hong Kong
+8522 599 7777

Europe/Africa/Middle East:

Hewlett-Packard S.A.
150, Route du Nant-d'Avril
CH-1217 Meyrin 2
Geneva, Switzerland
+41 22 780 81 11
European Multicountry: +41 22 780 81 11
Middle East and Africa: +41 22 780 71 11
European Headquarters: +41 22 780 81 81
Refer to country phone numbers

For direct country contact call:

Argentina: +541 787 7145

Austria: +43 1 25 000 0

Belgium and Luxembourg: +32 2 778 31 11

Brazil: +5511 7296 8000

Chile: +562 203 3233

Colombia: +571 629 5030

Denmark: +45 45 99 10 00

East Central Europe, CIS, and Yugoslavia:
+43 1 25 000 0

Finland: +358 9 887 21

France: +33 1 69 82 60 60

Germany: +49 7031 140

Greece: +30 1 689 644

Hungary: +36 1 252 7300

Iceland: High Performance Systems hf.
+354 1 67 10 00

Ireland: +353 1 615 8200

Israel: Computation and Measurement Systems
(CMS) Ltd. +972 3 5380 333

Italy: +39 2 92122770

Mexico: +525 326 4600

Netherlands: +31 20 547 6911

Norway: +47 22 7356 00

Poland: +48 22 608 77 00

Portugal: +351 1301 7343

Russia and the CIS, excl. Ukraine:
+7 095 923 5001

Slovenia: +38 61 55 84 72

Spain: +34 1 631 1600

Sweden: +46 8 444 2000

Switzerland: +411 735 7111

South Africa: Hewlett-Packard South Africa
(Pty) Ltd.+27 11 806 1000

Turkey: +90 212 224 5925

United Kingdom: +44 1344 369231

Venezuela: +582 239 4133

Full information on HP products is available at www.hp.com

Intel and Itanium are trademarks or registered trademarks of Intel Corporation in the U.S. and other countries and are used under license. Java is a U.S. trademark of Sun Microsystems, Inc. UNIX is a registered trademark of The Open Group. Windows and Windows NT are U.S. registered trademarks of Microsoft Corporation.

The information contained in this document is subject to change without notice.

© Copyright Hewlett-Packard Company 2001.

M0901

5980-7679EN

